

Adaptive Sampling of River Plume Fronts:
Integrating Statistical Modeling and Autonomous
Path Planning for Enhanced Oceanographic
Exploration

Yaolin Ge

December 14, 2023

Abstract

The ocean remains largely unexplored and presents a great challenge for scientific research. Ocean fronts have shown importance for understanding both biological and physical oceanographic phenomena, with river plume fronts being particularly intriguing. These fronts are a complex combination of freshwater and oceanic systems and provide a unique perspective on the dynamics of frontal systems.

Despite the acknowledged significance of ocean fronts, their comprehensive sampling is still a goal not yet achieved. To address this issue, this study tries to develop an approach that utilizes statistical techniques, robotics, and oceanographic knowledge to create an intelligent and ASS tailored to investigate river plume fronts or other similar frontal systems. This system is designed to address the spatial and temporal complexities of ocean fronts, allowing for more efficient and representative sampling.

The basis of our approach is the incorporation of Gaussian random fields. This modeling technique provides a robust proxy to the intricate field dynamics, allowing us to capture underlying patterns and forecast evolving behaviors. This surrogate modeling approach enables a nuanced understanding of the river plume front, reducing the computational burden associated with real-time decision-making, this means that it can be conducted on-board a robotic agent such as an autonomous underwater vehicle.

Building upon this proxy model, we develop and implement both myopic and non-myopic path planning algorithms. Myopic planning guides autonomous agents in immediate decision-making, capitalizing on localized data to direct sampling efforts. In contrast, non-myopic planning offers a broader perspective, considering the entirety of

the available information to optimize sampling across the entire field. The harmonization of these algorithms presents an unprecedented opportunity to balance immediate responsiveness with long-term strategic sampling.

We engage in a rigorous evaluation of our system through a series of simulation studies, modeling different scenarios and conditions that represent real-world complexities. Furthermore, we conduct experimental validations in authentic marine environments, leveraging our system's adaptability to capture the transient and spatially heterogeneous nature of river plume fronts. Autonomous underwater vehicles are used in our field deployments.

The results of our investigation demonstrate the robustness and versatility of our approach and its potential to improve oceanographic sampling. By seamlessly weaving together statistical modeling and strategic path planning, we present an adaptive sampling framework for exploring our ocean more efficiently and intelligently. This work showcases the possibility of adaptive ocean sampling and contributes to our understanding of river plume fronts.

Preface

This dissertation is submitted to the Norwegian University of Science and Technology (NTNU) in partial fulfillment of the requirements for the degree of Doctor of Philosophy (PhD). The work contained herein unveils an interdisciplinary investigation into the adaptive sampling of river plume fronts within the realm of ocean science. This exploration harmonizes disciplines including statistics, robotics, and oceanography, presenting a comprehensive approach to a complex subject. This dissertation is edited as a collection of papers.

The research was carried out at the Department of Mathematical Sciences, with invaluable support from the Research Council of Norway (RCN) through the MASCOT project (grant number 305445). The field experiment component of this study was made possible by an esteemed collaboration with the Applied Underwater Robotics Laboratory (AURLab) in Trondheim, Norway, and the Laboratório de Sistemas e Tecnologia Subaquática (LSTS) in Porto, Portugal. This cooperative endeavor has significantly enriched the research process, fusing expertise and resources to propel our understanding of adaptive sampling in oceanic applications.

Acknowledgment

I wish to express my deepest appreciation to my main supervisor, Professor Jo Eidsvik, for entrusting me with this opportunity and guiding me with unwavering enthusiasm and dedication. His passion for life and research, coupled with his constant encourage-

ment, has been an enduring source of inspiration throughout my doctoral journey.

I am grateful to my co-supervisors, Geir-Arne Fuglstad and Kanna Rajan, for their invaluable support. Mr. Fuglstad's expertise in statistics has greatly enhanced this research, while Mr. Rajan's broad assistance has not only shaped many aspects of the project, but also significantly improved my presentation skills.

The collaborative spirit of my fellow PhD researchers, Martin Outzen Berild and André Julius Hovd Olaisen, has been indispensable. Together, we forged great memories onboard boat trips, engaged in thought-provoking discussions to unravel complex problems, and developed a fellowship that enriched our shared experience over the past three years.

Special acknowledgment is extended to Per Kristian Hove and Torstein Fiskvik from the IT group at the department of mathematical sciences at NTNU. Their expertise in computer skills and server setup was essential for the simulation runs, providing robust technical support that was critical to the success of my research endeavors.

I extend my sincere gratitude to Tore Mo-Bjørkelund and Renato Mendes, collaborators from the field trips in Trondheim, Norway and Porto, Portugal, respectively. Their dedication, readiness to assist, and good-natured companionship have not only facilitated the technical aspects of the project but also made the long hours on the boat an enjoyable experience.

Special thanks are due to our collaborators in Trondheim and Porto, Karoline Barstein, Pedro De La Torre, Kay Arne Skarpnes, Paulo Dias, Manuel Ribeiro, Marie Costa, and Wang Dzak Choi for their assistance, ranging from tests in Nidelva and Douro to hardware and software support, has been invaluable.

My gratitude extends to Professors João Sousa, Tor Arne Johansen, and post-doctoral researcher R. Praveen Jain for their scholarly insights, which significantly enriched this research.

I also want to express my heartfelt gratitude to the Norwegian Research Council (RCN) for their generous grant, and the Applied Underwater Robotics Laboratory (AURLab) and Laboratório de Sistemas e Tecnologia Subaquática (LSTS) for supplying the Autonomous Underwater Vehicles (AUVs) and other equipment essential for the tests.

I would be remiss if I did not acknowledge the support from Ingrid Ellingsen for the SINMOD data and many other valuable insights.

Lastly, I must express my profound gratitude to my family, friends, and colleagues. Their unwavering support, love, and encouragement have been the bedrock upon which this research journey was built.

YaoLin Ge

Trondheim, December 14, 2023

Contents

Abstract	3
Preface	5
1 Introduction	11
1.1 Background	11
1.2 Project description	12
1.3 Outline	13
2 Problem Statement	15
2.1 River plumes	15
2.2 Adaptive AUV sampling	19
3 Adaptive Sampling System	21
3.1 Sense	22
3.1.1 Gaussian random field	23
3.1.2 Illustrative study	25
3.2 Plan	28
3.2.1 Boundary classification	28
3.2.2 Planning metrics	29
3.2.3 Path planning	30
Myopic path planning	30

Non-myopic path planning	31
3.3 Act	32
4 Evaluation and Testing	35
4.1 Simulation	35
4.2 Experiments	36
5 Summary of papers	39
5.1 Paper I	39
5.2 Paper II	40
5.3 Paper III	40
5.4 Paper IV	41
5.5 Paper V	41
6 Remarks	43
Bibliography	45

Chapter 1

Introduction

1.1 Background

The ocean is a major component of our planet, covering more than 70% of the surface of the Earth¹. It is a major heat reservoir, influencing weather patterns and moderating temperatures around the world (Rhein et al. 2013). Furthermore, the ocean absorbs almost all greenhouse gases created by humans (Takahashi et al. 2009) and it produces more than half of the oxygen on Earth (Field et al. 1998). It is also home to a wide variety of life forms, from the tiniest plankton to the largest whales. In particular, coral reefs, known as the "rain forests of the sea," are responsible for sustaining around 25% of all marine species, although they take up less than 1% of the ocean floor². The ocean provides a wide range of opportunities for the fishing industry, which supports the livelihoods of more than 600 million people around the world³. Additionally, maritime traffic is responsible for the transport of 80% of global trade, demonstrating its importance in international networking⁴.

Ocean fronts are the dynamic boundaries between different water masses. These

¹All About the Ocean by National Geographic

²Coral reef ecosystems by National Geographic

³Oceans fisheries and coastal economies by The World Bank

⁴Review of Maritime Transport 2022 by UNCTAD

fronts are essential for ocean ecosystems and climate, as they often have high gradients in salinity, temperature, density, and nutrient concentrations, leading to increased biodiversity and productivity (Chapman et al. 2020; Mańko et al. 2022). In addition, upwelling and downwelling of water masses can significantly influence the distribution and transport of nutrients, heat, and carbon, affecting global climate (Coogan et al. 2019). It is hence essential to understand ocean fronts to comprehend marine ecosystems and global climate, as well as to create conservation plans and regulate human effects on these ecologically significant areas. River plumes, which are a frequent occurrence in places where a river meets the sea, are an ideal focus for investigation.

Despite their importance, our knowledge of ocean fronts is limited due to the difficulty in sampling these dynamic and often remote regions. In the past, ship surveys, satellite observations, static buoys, and dynamic drifters have been employed in sampling efforts; however, their high operational costs, vulnerability to inclement weather, and need for human involvement are major drawbacks (Lin and Yang 2020). Recent advances in technology, such as autonomous underwater vehicles (AUVs) and high-fidelity numerical ocean models, offer promising solutions to improve ocean sampling (Das et al. 2015; Fossum et al. 2019; McCammon et al. 2021; Berget et al. 2023; Ge et al. 2023).

1.2 Project description

The Maritime Autonomous Sampling and Control (MASCOT) project, funded by the Norwegian Research Council (NRC), seeks to explore the characteristics of river plumes through a combination of multiple disciplines such as spatio-temporal statistics, robotics and oceanography. The project aims to combine traditional data sources such as satellite imagery and numerical models with modern approaches for robotic sampling (Slagstad and McClimans 2005; Lin and Yang 2020). Conventionally, AUVs utilize fixed, pre-programmed paths, limiting their adaptability in dynamic marine environments. This leads to a crucial need for ASS capable of real-time adjustments to fluctuating oceanographic conditions. To address these issues, an adaptive sampling system using AUVs is proposed. This method involves the use of dynamic sensing systems that can adjust

their sampling strategies based on observed data, allowing for more efficient data collection. This technique has been successfully applied in various oceanographic studies, including the examination of river plumes; see (Das et al. 2015; Fossum et al. 2019; McCammon et al. 2021; Berget et al. 2023; Anyosa et al. 2023; Ge et al. 2023).

1.3 Outline

In the subsequent chapters, we outline our approach to the adaptive sampling issue. Chapter 2 provides an overview of river plumes, their characteristics, and a review of existing systems for conducting autonomous sampling of such characteristics. In Chapter 3, we introduce our Adaptive Sampling System (ASS) developed to address these challenges. Chapter 4 describes our methods for testing and evaluating the ASS. Chapter 5 summarizes the papers produced during this Ph.D. research. Chapter 6 shows the final remarks.

Chapter 2

Problem Statement

2.1 River plumes

River plumes, which are freshwater bodies released from rivers into the sea, are essential for the transportation of sediments, nutrients, and pollutants, and thus have a major effect on coastal dynamics, marine ecosystems, and water quality (Devlin et al. 2015). The dynamics of river plumes is intricate and is governed by a variety of elements, such as river discharge, oceanic conditions, wind, and the Coriolis effect (Horner-Devine et al. 2015*a*). For example, the Coriolis effect shapes the dynamics of river plumes in large-scale coastal areas, and the wind direction and strength significantly affect plume behavior. River plumes also have a considerable impact on coastal and marine ecosystems. They carry nutrients and sediments from land-based sources, influencing primary productivity and sediment dynamics in coastal waters (Sagar et al. 2020).

Both Figure 2.1 and Figure 2.2 exhibit characteristic outbursts of river plumes, although they are situated in different geographical locations. Figure 2.1 illustrates a typical river plume outburst on the Douro River in Porto, Portugal, on 5 January 2020. The plume's scale is substantial, and suspended particulate matter contributes to its visible appearance. Moreover, the plume's dimensions and dynamics are influenced by factors such as water discharge, wind direction, and the Coriolis effect. During the




Figure 2.1: Sentinel-2 image of the Douro River plume. Captured on 5 January 2020. Courtesy of Copernicus Open Access Hub.

specified period, the water discharge¹ reached its peak, while the persistent east wind² enlarged the plume and induced a rightward drift as it deepened.

In contrast, Figure 2.2 portrays a typical river plume in the Nidelva River in Trondheim, Norway, on 30 October 2022. This plume exhibits a relatively smaller size compared to the one observed in the Douro River. Additionally, the water in Nidelva appears to be clearer, and a reduced presence of suspended particulate matter is observed in the surroundings. The annual average water discharge of the Douro River, as noted by Azevedo et al. (2010), is 505 cubic meters per second (m^3/s). In comparison, the Nidelva River's annual average water discharge, documented by Haraldstad et al. (2022), stands at 110 cubic meters per second (m^3/s).

In addition to the acquisition of satellite images, extensive fieldwork was conducted onboard a boat to pursue the river plume actively. Subsequently, our aerial survey utilizing a DJI Mini 3 Pro drone has provided valuable visualizations of the dynamic river

¹Redes de Monitorização by Sistema Nacional de Informação

²Weather history of Vila Nova da Telha, Porto, Portugal by Weather Underground




Figure 2.2: Sentinel-2 image of the Nidelva's river plume. Captured on 30 October 2022. Courtesy of Copernicus Open Access Hub.

plume from an elevated perspective. Figure 2.3 portrays the discernible boundaries between two distinct water masses, effectively capturing the striking characteristics of the river plume at two separate locations in close proximity to the Nidelva River in Trondheim, Norway, during the end of June, 2023.

In addition to satellite images, one can also use high-precision numerical ocean models such as SINMOD (Slagstad and McClimans 2005) and Delft3D (Mendes et al. 2016) to generate initial understanding of the desired field. An example of the numerical estimation of surface salinity at different locations from SINMOD and Delft3D is shown in Figure 2.4. Such outputs from the numerical ocean models often capture the large-scale dependencies of the ocean variables, but they can be biased. Hence, in-situ sampling is important for calibration and updating.


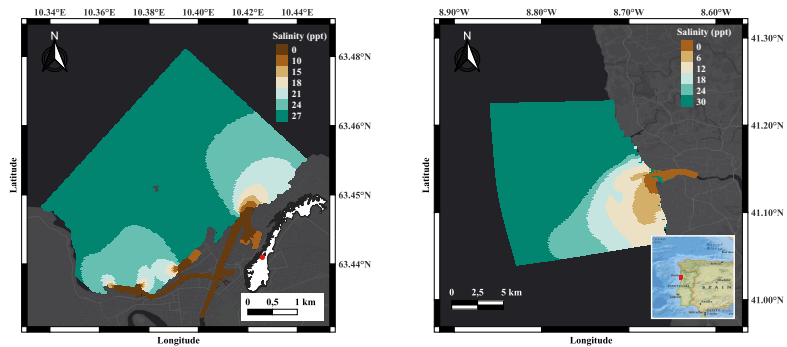


Figure 2.3: Aerial view captured by the DJI Mini 3 Pro drone, revealing the intricate boundaries between different water masses near the Nidelva River in Trondheim, Norway. The image showcases the vibrant dynamics of the river plume during the final days of June, providing valuable insights into its spatial distribution.



(a) SINMOD

(b) Delft3D

Figure 2.4: Comparison of surface salinity estimates from SINMOD and Delft3D in the Trondheimsfjord and Rio Douro regions.




Figure 2.5: LAUV-Thor undergoes remote operation check before its expedition.

2.2 Adaptive AUV sampling

The utility of AUVs extends to various domains of marine research, enabling effective contributions to oil plume detection with hydrocarbon sensors (Kukulya et al. (2018); Thorsnes et al. (2019)), thermocline monitoring via temperature sensors (Antunes and Cruz (2019)), upwelling fronts exploration with biological and chemical sensors (Jackson and Reneau (2014)), and internal waves analysis by measuring water density and salinity (Ratsimandresy et al. (2014); Toberman et al. (2017)). Furthermore, AUVs are indispensable in studying biological and chemical features to assess ecosystem health (Cowles and Donaghay (1998); Fossum et al. (2019)).

Recent developments in AUV maneuverability and computing power have enabled the integration of modern ASS, for autonomous decision-making, improved sensor technology, and advanced navigational autonomy. These AUV systems are designed to adjust their sampling strategies based on real-time environmental data, thus increasing the efficiency of data collection. Examples of such innovations include glider AUVs (Lagunas et al. (2018); Chave et al. (2018); Galarza Bogotá (2018)), multi-AUV systems (Ahmadzadeh et al. (2006); Fiorelli et al. (2006); Cui et al. (2015)), and the incorporation of statistical techniques into AUVs for real-time data processing (Fossum

(2019); Ge et al. (2023)).

The purpose of this study is to utilize an AUV to sample the spatial domain of interest over time in an adaptive way. An AUV (see Figure 2.5) has a variety of sensors, cameras, and navigation abilities that enable it to move autonomously and investigate various depths. Its capacity to modify its course and act in response to real-time data and environmental conditions makes adaptive sampling an increasingly attractive concept. In our case, we aim to build an ASS to recognize the boundaries between different water masses.

In order for the adaptive sampling to be effective, we must design systems that can assimilate in-situ data and then make decisions based on the updated knowledge of the uncertain spatio-temporal field. To do this, we are using Gaussian random field (GRF) models as the main surrogate model for data assimilation. A GRF has the capacity to model complex spatial relationships and quantify uncertainties in predictions. It can flexibly adapt to various spatial patterns without imposing rigid assumptions. By utilizing the Bayesian framework, a GRF can infer the underlying spatial structure by encoding prior beliefs about the data and updating the posterior with the observed measurements. This makes a GRF model an ideal choice for geo-spatial applications, where capturing spatial dependencies and accurately characterizing uncertainties are essential. Further, the limited memory and processing power needed make it possible for the AUV to carry out calculations while in use, in real-time. This should not be possible with a more complicated surrogate model building on complex differential equations or neural network models.

A key element of the system is the capacity to make decisions that can guide the AUV through the area, keeping in mind that the ultimate goal is to sample efficiently. This process of determining where the AUV should go is known as path planning. It is designed to provide optimized navigation in complex and often unpredictable underwater environments. Due to the difficult nature of underwater terrains and limited communication capabilities, it is essential to have intelligent path planning to guarantee the safety and success of the mission.

Chapter 3

Adaptive Sampling System

Leveraging the foundational research in the Sense-Plan-Act (SPA) framework, our system integrates key insights for enhanced decision-making and interaction with the environment. The SPA framework's effectiveness in various fields is well-established, with studies such as Oddi et al. (2020) demonstrating its extension in robotics through autonomous, open-ended learning, and Asmaa et al. (2019) highlighting its crucial role in Unmanned Aerial Vehicle (UAV) control architectures. Additionally, the application of the SPA paradigm in mobile health interventions is illustrated in Eapen et al. (2019), and its importance in autonomous navigation of mobile robots is explored by Nakhaeinia et al. (2011). The advancement of SPA through the multilayer architecture in UAV control, merging cognitive methods with traditional robotics, is presented in Emel'yanov et al. (2016).

Our system builds on these insights, adopting the SPA framework for intelligent, cyclical decision-making and environmental interaction. It starts by collecting environmental data through sensors, followed by planning based on this data to fulfill objectives. The system then executes these plans, enabling adaptive sampling and continuous adaptation, as depicted in Figure 3.1. This approach not only reflects the evolution of the SPA framework but also integrates the advancements from previous research into our system's unique application.




Figure 3.1: The Sense-Plan-Act (SPA) framework is responsible for organizing the cyclical process of collecting data from the environment, planning data-collection strategies, and carrying out accurate sampling activities.

3.1 Sense

In adaptive sampling, the 'Sense' phase is pivotal, encompassing both the collection and assimilation of environmental data. This phase involves deploying sensors and other tools to gather relevant data, which is crucial for the system to acquire a comprehensive understanding of its surroundings. The gathered data not only offers valuable insights but also aids in capturing the distribution of various environmental parameters, essential for accurate modeling and decision-making.

The process of data assimilation is equally important in this phase. It involves integrating the collected data points into the system's world model, ensuring that the model is continuously updated and reflective of the current state of the environment. This assimilation is critical for maintaining the accuracy and relevance of the model, allowing the system to adapt its responses based on the most current information.

For our system, the GRF acts as the proxy model for the environment. This choice is due to its effectiveness in handling spatial data, its ability to provide a probabilistic framework for understanding environmental variations, and the ease of model updating. The GRF excels in interpolating and predicting environmental states, making it an ideal

choice for adaptive systems that require a robust and dynamic understanding of their operational context.

The subsequent section will delve deeper into the specifics of how a GRF is utilized within our system, illustrating its role in enhancing the system’s sensing capabilities and its integration into the broader framework of adaptive sampling.

3.1.1 Gaussian random field

A GRF is a powerful and versatile non-parametric Bayesian modeling approach used for regression tasks. It is especially suitable for scenarios where the relationship between input and output data is not known beforehand and could be intricate or non-linear. A GRF not only offers point predictions but also estimates of the uncertainty related to those predictions, making it useful in the spatial analysis applications.

$$Posterior = \frac{Likelihood \times Prior}{NormalizationConstant}. \quad (1)$$

Equation (1) outlines the general expression for Bayesian inference. The prior distribution is a key element in the GRF as it reflects the initial assumptions about the trend, the variability and the smoothness and general behavior of the process being modeled. Selection of the prior distribution can have a significant effect on the predictions and uncertainty estimates in the GRF. By incorporating the observed data into the prior model, a GRF can make data-driven updates that refine the model and generate more accurate and meaningful predictions when more data is collected. Notably, we assume that the data that we observe are the variables of interest at the locations where they are taken. We also assume that the measurement noise is normally distributed, which implies that the posterior model is a GRF.

In our case, we use the GRF to model salinity in river plumes. We let $\xi_{\mathbf{u}}$ denote the salinity at location \mathbf{u} in a three-dimensional spatial domain $\mathcal{M} \subset \mathcal{R}^3$. Although we may sometimes focus on the surface layer, we keep it as a three-dimensional field for general reasons. In this presentation, we will focus on a spatial process, disregarding the temporal variation for simplicity. This GRF is used to estimate the mean, variability, and spatial dependence of the salinity field. To implement the model, the

spatial domain is often divided into n grid locations $\mathbf{u}_1, \dots, \mathbf{u}_n$, with the GRF at these locations represented by $\boldsymbol{\xi} = (\xi_{\mathbf{u}_1}, \dots, \xi_{\mathbf{u}_n})^T$. This GRF follows a normal distribution, specifically $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ is the mean salinity values and $\boldsymbol{\Sigma}$ is the covariance matrix. This matrix can be defined by a Matérn covariance function, characterized by elements $\Sigma(i, i') = \sigma^2(1 + \phi_1 h(i, i')) \exp(-\phi_1 h(i, i'))$, with σ^2 being the variance and ϕ_1 the correlation decay parameter (Cressie and Wikle 2015). The model can take into account anisotropy, which is essential for capturing the varying spatial relationships in the salinity field, particularly due to the differences in lateral and vertical stretches of river plumes (Horner-Devine et al. 2015b). This anisotropy is reflected in the weighted distance matrix between grid nodes, formulated as $h^2(i, i') = h_E^2(i, i') + h_N^2(i, i') + (\phi_1^2 / \phi_2^2) h_D^2(i, i')$, using separate correlation decay parameters ϕ_1 and ϕ_2 for the lateral and vertical dimensions, respectively.

The measurements at each stage $j = 1, \dots, J$ are modeled by a Gaussian likelihood model

$$\mathbf{y}_j | \boldsymbol{\xi} \sim N(\mathbf{F}_j \boldsymbol{\xi}, \mathbf{R}_j), \quad (2)$$

where \mathbf{F}_j is an $N_j \times n$ selection matrix containing an entry of 1 in each row and 0 otherwise. The 1 entry refers to the sampling indices. With the covariance matrix $\mathbf{R}_j = r^2 \mathbf{I}_{N_j}$, we assume that the data are conditionally independent, given the underlying salinity. Here, r indicates the measurement standard deviation of the AUV salinity observations. We denote the associated PDF by $p(\mathbf{y}_j | \boldsymbol{\xi})$.

Via Bayes' rule, data assimilation at stages $j = 1, \dots, J$, gives the sequential conditional PDF $p(\boldsymbol{\xi} | \mathcal{Y}_j) \propto p(\mathbf{y}_j | \boldsymbol{\xi}) p(\boldsymbol{\xi} | \mathcal{Y}_{j-1})$. At the initial stage, the data set \mathcal{Y}_0 is equal to Φ , and a prior GRF model is obtained. As time passes, the data set \mathcal{Y}_j increases, becoming a collection of \mathcal{Y}_{j-1} and \mathcal{Y}_j . Under the assumptions about a GRF prior model and a Gaussian measurement error model, this conditional PDF is also Gaussian with mean $\boldsymbol{\mu}_j$ and covariance matrix \mathbf{S}_j given by

$$\begin{aligned} \mathbf{G}_j &= \mathbf{S}_{j-1} \mathbf{F}_j^T (\mathbf{F}_j \mathbf{S}_{j-1} \mathbf{F}_j^T + \mathbf{R}_j)^{-1} \\ \boldsymbol{\mu}_j &= \boldsymbol{\mu}_{j-1} + \mathbf{G}_j (\mathbf{y}_j - \mathbf{F}_j \boldsymbol{\mu}_{j-1}) \\ \mathbf{S}_j &= \mathbf{S}_{j-1} - \mathbf{G}_j \mathbf{F}_j \mathbf{S}_{j-1}, \end{aligned} \quad (3)$$

where $\boldsymbol{\mu}_0 = \boldsymbol{\mu}$ and $\mathbf{S}_0 = \boldsymbol{\Sigma}$. The sequential updating resembles that of a spatio-temporal Kalman filter (Cressie and Wikle 2015).

A GRF offers a range of advantages, such as the capacity to measure the uncertainty of predictions and to perform kriging on the field with ease. However, the time complexity associated with the size of the field increases with the size of the field. Gaussian Markov random field (GMRF) provides a computationally efficient alternative to a GRF, particularly for large datasets, due to their sparse precision matrix $\mathbf{Q} = \boldsymbol{\Sigma}^{-1}$. This sparsity, which is a consequence of the Markov property, significantly accelerates computations by reducing the complexity associated with GRF's large and dense covariance matrices. Furthermore, a GMRF is particularly adept at modeling non-stationarity more naturally and flexibly. By adjusting the graph structure in the precision matrix, a GMRF can represent varying correlation degrees across different regions, offering a more nuanced approach to spatial variability that is especially useful in large-scale applications where uniformity cannot be assumed (Lindgren and Rue (2015); Fuglstad et al. (2015); Berild and Fuglstad (2023)).

3.1.2 Illustrative study

This study employs an illustrative case to elucidate the efficacy of using a GRF in spatial modeling. The prior scalar field is postulated to increment linearly along the east axis while maintaining uniformity along the north axis. The coefficients used in the Matérn covariance kernel are specified as $\sigma = 0.2$, $\phi_1 = \phi_2 = 7.5$, and $r = 0.1$. The excursion set below 0.5 and its corresponding probability are computed, as illustrated in Figure 3.2. A comprehensive examination of the methodologies for assessing the excursion set and probability is deferred to the following section. The excursion set delineates the threshold-based binary classification, assigning a value of one to regions falling below the threshold and zero to those exceeding it. The excursion probability represents the model's estimate of the veracity of the field's classification under current assumptions.

To establish simulated ground truth scenarios for subsequent sampling, three stochastic realizations of GRFs are generated, reflecting the inherent covariance structure and




Figure 3.2: Illustration of the prior mean field, the defined excursion set, and the corresponding excursion probability.

prior mean. These are depicted in Figure 3.3, showcasing the field's variability and the complexity of its boundaries.






Figure 3.3: Stochastic realizations representing hypothetical ground truth fields derived from the underlying covariance structure and prior mean.


Following the data acquisition along a predetermined trajectory, the prior field is updated to yield a conditional field. This field of interest embodies a refined estimate, integrating empirical observations to enhance the model's initial predictions. Figure 3.4 presents the posterior fields derived from the GRF realizations in Figure 3.3, utilizing data collected along the fixed path along the center line and juxtaposed against three distinct ground truth scenarios. These visualizations affirm the model's proficiency in assimilating new data to update its predictions substantively.



(a) Updated predictive field based on empirical data from ground truth scenario I.



(b) Updated predictive field based on empirical data from ground truth scenario II.



(c) Updated predictive field based on empirical data from ground truth scenario III.

Figure 3.4: Posterior fields demonstrating the mean, excursion set, and excursion probability after empirical data integration along a fix path from respective ground truth scenarios I, II, and III.

3.2 Plan

The planning stage of the ASS is essential for making decisions. It is necessary to gain a comprehensive understanding of the current situation in order to decide on the most suitable strategy to achieve the boundary mapping objectives, in our case, of characterizing fronts. This task can be divided into two components: the boundary classifier and the path planner. The boundary classifier attempts to classify the boundary area by utilizing the univariance in a statistical construction known as the excursion set. The path planner must then determine the most advantageous route to take, accounting for practical matters such as time constraints, as well as other factors such as balancing exploration and exploitation.

3.2.1 Boundary classification

We use the notion of an excursion set (ES) to characterize the river and ocean water masses (Fossum et al. 2019). The ES for salinity threshold t is defined by

$$\text{ES} = \{\mathbf{u} \in \mathcal{M} : \xi_{\mathbf{u}} \leq t\}. \quad (1)$$

Hence, salinity lower than this threshold will indicate river water. The associated excursion probability (EP) is

$$p_{\mathbf{u}} = P(\xi_{\mathbf{u}} \leq t), \quad \mathbf{u} \in \mathcal{M}. \quad (2)$$

When it is close to 1 or 0 at a given location, it is easy to classify the water mass to be river or ocean respectively. EP close to 0.5 reflects ambiguity in the characterization of water masses. The prior Bernoulli variance (BV) at location \mathbf{u} is $p_{\mathbf{u}}(1 - p_{\mathbf{u}})$ and the spatially integrated BV (IBV) is

$$\text{IBV} = \int p_{\mathbf{u}}(1 - p_{\mathbf{u}})d\mathbf{u}, \quad (3)$$

which is dominated by locations with probabilities near 0.5 and BV close to 0.25. In practice the integral will be approximated by a sum over the n grid nodes.

3.2.2 Planning metrics

In order to identify the most useful areas of the field, we need to focus our efforts on areas that are deemed to be the most informative or interesting. For optimization or sampling purposes, there exists a range of metrics including expected improvement (EI), variance reduction, and others. We have chosen certain metrics based on empirical evidence from previous studies and their demonstrated effectiveness for the task at hand. Specifically, we have adopted the following:

- **Expected integrated Bernoulli variance (EIBV):** This criterion facilitates the AUV in targeting locations that are likely to be ambiguous, thus channeling focused efforts around the boundary regions. The goal is to construct AUV sampling strategies that prioritize locations that are ambiguous, thus making the exploration more effective. At each stage, considering only the data at that stage, we define the EIBV by

$$\begin{aligned} \text{EIBV}(\mathbf{D}_j) &= \int E_{\mathbf{y}_j|\mathcal{Y}_{j-1},\mathbf{D}_j} [B_{\mathbf{u}}(\mathbf{y}_j)] d\mathbf{u}, \\ B_{\mathbf{u}}(\mathbf{y}_j) &= p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1})(1 - p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1})), \end{aligned} \quad (4)$$

where $B_{\mathbf{u}}(\mathbf{y}_j)$ is the conditional Bernoulli variance for outcome \mathbf{y}_j of data in design \mathbf{D}_j , and the conditional probability of an excursion is

$$p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1}) = P(\xi_{\mathbf{u}} \leq t | \mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1}). \quad (5)$$

With the GRF modeling assumptions, the EIBV in Equation 4 has a closed-form solution given via cumulative distribution functions of the bivariate Gaussian distribution (Fossum 2019).

- **Integrated variance reduction (IVR):** This approach seeks to spread the sampling efforts throughout the area, striving to reduce variance to the greatest extent possible. It embodies a more comprehensive strategy for understanding the fundamental structure of the field. The IVR for a specific design \mathbf{D}_j is defined as

$$\text{IVR}(\mathbf{D}_j) = \text{trace}(\mathbf{S}_{j-1} \mathbf{F}_{\mathbf{D}_j}^T (\mathbf{F}_{\mathbf{D}_j} \mathbf{S}_{j-1} \mathbf{F}_{\mathbf{D}_j}^T + \mathbf{R}_j)^{-1} \mathbf{F}_{\mathbf{D}_j} \mathbf{S}_{j-1}) \quad (6)$$

- **Operational constraints:** This component takes into account the complexities of the real world, including the presence of obstacles, the amount of time spent, and any other navigational difficulties that could obstruct the sampling procedure.

3.2.3 Path planning

Path planning is essential in robotics, facilitating efficient navigation for autonomous systems from one point to another. Dynamic Programming (DP) is a cornerstone in this area, breaking complex navigation tasks into simpler, recursive sub-problems. It relies on the Bellman equation, which posits that the optimal path in a route is an aggregation of optimal sub-paths (Bellman 1966). This principle allows DP to methodically evaluate all potential paths, selecting the most suitable one by using a cost function to gauge transitions between states. Moreover, DP's recursive nature optimizes computational resources by caching and reusing calculations for previously computed sub-paths.

However, the high computational load of traditional DP often makes it impractical for real-world scenarios, especially in robotics where high-dimensional state spaces are common. This challenge has led to the adoption of Approximate Dynamic Programming (ADP), which offers a more feasible approach for such applications (Powell 2007). While ADP retains the fundamental recursive framework of DP and its dependence on the Bellman equation, it introduces strategic approximations and heuristics. These modifications effectively manage the complexities and scalability issues inherent in standard DP. ADP strikes a balance between computational efficiency and the accuracy of solutions, focusing on the evaluation and selection of near-optimal paths. This approach not only ensures practicality in real-world implementations but also maintains the integrity of path planning processes, including both myopic and non-myopic strategies. For our situation, relying on computations on-board the AUV is the main limitation to the complexity of the ASS.

Myopic path planning

Myopic path planning employs a short-term approach that concentrates on immediate or local improvements instead of considering potential long-term effects. It is fast and




Figure 3.5: An illustration of myopic path planning is provided. The AUV plans its next waypoint to achieve its immediate objective, which may lead to overlooking more interesting areas.

responsive, making decisions quickly based on the current situation. Myopic planning is a type of ADP that only considers a single-stage expectation and does not take into account the potential implications of the current data for future stages of planning. This makes it suitable for scenarios that require quick reactions in a changing environment, however its greediness can lead to suboptimal solutions and getting stuck in local minima. The emphasis on short-term gains may overlook long-term benefits or opportunities, thus limiting its effectiveness in more strategic or complex circumstances, as demonstrated in Figure 3.5.

Non-myopic path planning

In contrast to myopic path planning, non-myopic path planning takes into account a broader perspective, considering potential future advantages along the route. This strategy goes beyond the immediate rewards and produces a more strategic path that evaluates the long-term benefits. It hence provides a much more complex ADP than the myopic approach. Depending on the particular approach and situation, it can achieve a close to optimal solution, but it can be a great challenge for the computational capacity. The long-term capability of non-myopic path planning usually involves higher




Figure 3.6: An example of non-myopic path planning is shown. The AUV plans its future path in order to reach its long-term goal, rather than settling for a locally optimal region which may be more convenient in the present.

computational costs and requires careful design of the constraints, but it is likely to generate a more robust and insightful solution for complex problems. An example of non-myopic path planning is depicted in Figure 3.6. The AUV formulates its future trajectory based on its long-term goal, rather than focusing on close-by convenient areas. One example of such long-horizon path planning algorithms is Rapidly-exploring Random Trees (RRT*), which is suitable for intricate, ever-changing environments in robotics and self-driving vehicles (Karaman and Frazzoli 2011).

3.3 Act

At this juncture, the AUV initiates the execution of its planned trajectory, a process intricately designed in the preceding strategic phase. Typically, this trajectory is constituted by a series of waypoints that the AUV must navigate through. The primary objectives during this phase is to traverse the desired path efficiently. Additionally, the

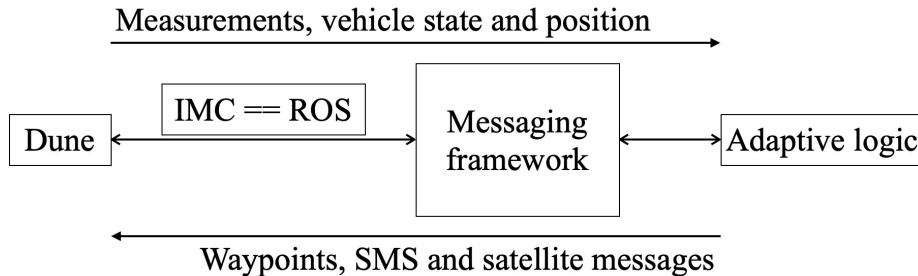


Figure 3.7: Illustration of the adaptive sampling framework developed at NTNU, highlighting the integration of ASS with hardware control algorithms via ROS messages (Mo-Bjørkelund et al. 2020).

AUV is equipped to rectify any navigational discrepancies that may arise during this operation.

To facilitate these objectives, the AUV utilizes an adaptive sampling framework in its embedded system, developed at NTNU. This framework (Mo-Bjørkelund et al. 2020), as depicted in Figure 3.7, serves a pivotal role in synchronizing the adaptive sampling mechanism with the AUV’s hardware control algorithms. The essence of this synchronization lies in the efficient exchange of information, which is realized through the Robot Operating System (ROS) messages. ROS, an open-source robotics middleware suite, offers a standardized communication protocol that allows diverse software components, including those responsible for data collection and hardware control, to interact seamlessly Quigley (2009).

This framework demonstrates how data collection and decision-making can be accomplished at the edge with an AUV. By utilizing ROS, the AUV is able to make real-time decisions based on the data it collects, decentralizing the decision-making process and increasing the responsiveness and adaptability of the AUV in dynamic underwater environments. This approach is statistically significant as it allows for dynamic adjustment of sampling strategies based on real-time environmental feedback, optimizing data collection for further analysis.

Chapter 4

Evaluation and Testing

We evaluate our ASS through two types of studies. We use a simulation study to uncover any potential hazards or undiscovered phenomena, and to gain insights from the model parameters. This allows us to fine-tune the parameters for the subsequent experiments. To further validate the system, we conduct field experiments in the fjord system in Trondheim, Norway and the Douro river in Porto, Portugal. The field trips in Trondheim require us to be physically present in the boat to conduct the experiment. On the other hand, the field trips in Porto are partially on-site and the rest of the mission are done remotely. We were assisted by the personnel in Porto when we started the mission.

4.1 Simulation

This study is designed to replicate the real-world situation as closely as possible, given all reasonable assumptions. To do this, we select an operational area that is likely to be the first place the river plume appears. We then obtain numerical results and satellite images to help us identify potential fronts. We build statistical prior models such as GRF and GMRF, which are trained on numerical results, including SINMOD and Delft3D (Figure 2.4). This agent is then run in a simulated environment, using simu-

lated ground truth. To reduce randomness, the proposed algorithms are compared to state-of-the-art algorithms for many replicates. We monitor a range of metrics, including root mean squared error (RMSE), integrated Bernoulli variance (IBV), variance reduction (VR), and potentially classification error (CE), to gain a deeper understanding of the properties of the system.

4.2 Experiments

The most stimulating part of the validation process is usually the field experiment. It is usually a multi-day mission that involves planning, executing, and post-processing. During the planning phase, two stages are typically involved in the preparation for the field experiment: software-in-the-loop (SIL) and hardware-in-the-loop (HIL). SIL tests the software used to ensure the accuracy of the algorithms and the overall functionality of the system. HIL integrates the system with the hardware units to test its interactions with real-world elements. The deployment to the ocean during the executing phase is the final stage of the experiment, where the performance of the system is observed under actual conditions, demonstrating its ability to handle the complexities of the real-world. Finally, data analysis is conducted to gain further insight into the successes and failures of the experiment.

The timeline in Figure 4.1 shows the field experiments I conducted during my PhD study. The challenges of conducting a real experiment underwater are numerous. Navigation can become inaccurate over time, and air bubbles can cause fluctuations in measured values, adding to measurement noise. Communication channels are not as reliable as on land, and the risk of collision with undetected objects such as fish or plastic bottles is high, potentially leading to the loss of the AUV. Additionally, marine traffic is unpredictable, and currents and waves can be strong, making it difficult to fix bugs on board a small boat while people are feeling seasick.

Despite these challenges, the gains from the experiment are substantial and multifaceted. Apart from the technical and scientific insights, we often find joy in the physical experiences involved. Fun moments arise from driving a boat across the vast and unpredictable ocean, feeling a mix of excitement and respect for nature's grandeur.




Figure 4.1: The timeline of all the field experiments conducted during my PhD period.

The opportunity to be outdoors, breathing fresh air, and being surrounded by the sea provides a refreshing break from the traditional office environment. The teamwork required for the experiment fosters actual face-to-face communication with colleagues, forging stronger relationships, and allowing for real-time problem-solving. Moreover, the memory of sailing around the ocean, facing unforeseen challenges, and overcoming them with collaboration and determination, becomes a cherished experience. The mix of technical mastery, adventure, and camaraderie forms a unique blend that sets such experimental missions apart and makes them invaluable in both professional growth and personal fulfillment.



Figure 4.2: Pictures taken during my doctoral studies demonstrate the successful completion of a field experiment.



Figure 4.3: The field experiment yielded some unexpected results, as evidenced by the image on the left which shows an AUV stranded close to the shore. The image on the right displays an AUV with its wings damaged.

Chapter 5

Summary of papers

5.1 Paper I

- **Title:** 3-D Adaptive AUV Sampling for Classification of Water Masses
- **Highlights:** This paper presents an adaptive sampling algorithm for AUVs to classify water masses in 3D oceanographic systems, particularly river plumes. It utilizes GRF models for salinity to differentiate between river and ocean waters. The study proposes a real-time myopic path planning strategy, which adjusts based on onboard AUV sensor data. The EIBV is used as a statistical design criterion for optimizing sampling locations, with the goal of reducing classification uncertainty of river plumes and fjord water masses. The efficacy of this methodology is tested through simulation studies and real-world field trials in Norway's Nidelva river plume, showing significant progress in adaptive AUV sampling techniques. The main contribution of this paper is to use an AUV for 3D adaptive sampling.

5.2 Paper II

- **Title:** RRT*-Enhanced Long-Horizon Path Planning for AUV Adaptive Sampling using a Cost Valley
- **Highlights:** This paper presents an ASS for AUVs that utilizes a long-horizon path planning approach with a cost valley approach for efficient oceanographic sampling, particularly of river plumes. The system combines a GRF model with a RRT* algorithm for onboard path planning, taking into account exploration, exploitation, and operational constraints. The system was tested through simulations and a field trial in Trondheim fjord, and was found to be able to effectively navigate dynamic marine environments, improving salinity field characterization. The main contribution of this paper is the use of an AUV to sample the river plume front of the Trondheim fjord region using the RRT*-enhanced long-horizon path planner and a spatio-temporal GRF.

5.3 Paper III

- **Title:** Efficient 3D real-time adaptive AUV sampling of a river plume front
- **Highlights:** This paper presents an adaptive AUV sampling technique for mapping ocean salinity. The method utilizes GMRF to learn from high-resolution ocean model data, and outperforms traditional statistical models in simulations. Field tests conducted near Trondheim, Norway, confirmed the efficacy of the approach, with one AUV dynamically adapting its path and outperforming another on a pre-planned route. The results suggest potential applications for ocean monitoring, although further refinement is needed for complex spatial modeling. The main contribution of this paper is to demonstrate the use of an AUV for the adaptive sampling of the river plume front in the Nidelva region, by employing a complex non-stationary GMRF model as the proxy model and a myopic path planner.

5.4 Paper IV

- **Title:** Using expected improvement of gradients for robotic exploration of ocean salinity fronts
- **Highlights:** This paper presents an algorithm for autonomous robotic exploration of ocean salinity fronts, which utilizes spatio-temporal GRFs to anticipate and investigate regions with considerable salinity variations. The algorithm, designed for an AUV, calculates expected improvements in directional derivatives to adjust sampling paths in real-time. The approach was tested through simulations and a field experiment in the Trondheim fjord, where the AUV successfully autonomously sampled salinity changes over two hours. This study provides a reliable method for dynamic sampling in marine environments, with potential applications in various domains requiring real-time adaptive monitoring. The main contribution of this paper is the use of an AUV for the adaptive sampling of the high-gradient area close to the river mouth region in the Nidelva river.

5.5 Paper V

- **Title:** Long-Horizon Informative Path Planning with Obstacles and Time Constraints
- **Highlights:** This paper presents a long-horizon path planning algorithm for AUVs that takes into account obstacles and time constraints. The algorithm uses a "cost valley" strategy to guide the AUV through a simulated river plume, maximizing information gain about salinity fronts while adhering to operational limits. RRT* is employed for path planning, accounting for various constraints such as distance budget, obstacles, and directional changes, with an emphasis on reducing the variance of the field and focusing on salinity excursions. Simulation results show that the algorithm outperforms traditional myopic strategies and pre-scripted lawnmower patterns, balancing exploration and the need to return to the base within a set time frame. The main contribution of this paper is to

demonstrate the feasibility of using RRT* as a path planner with a flexible cost valley for long-horizon path planning.

Chapter 6

Remarks

In this thesis, we have endeavored to address the complex challenge of adaptive sampling of river plume fronts, an area that sits at the confluence of statistical modeling, robotics, and oceanography. Our journey has been one of exploration and learning, characterized by both progress and the recognition of ongoing challenges. The development of an ASS, while a step forward, has also opened up avenues for future research and development.

Our approach, which combined GRF and GMRF with myopic and non-myopic path planning algorithms, aimed to refine the efficiency and accuracy of oceanographic sampling. Through simulations and field experiments, we have gained valuable insights into the workings of this system in diverse marine environments. These practical applications have provided a clearer understanding of the system's capabilities and have highlighted areas that require further exploration.

Looking to the future, the integration of multiple AUVs and advanced machine learning techniques presents a promising direction for oceanic exploration. Coordinating multiple AUVs could increase the efficiency and scope of data collection, providing a more comprehensive understanding of ocean dynamics. Simultaneously, applying machine learning to large environmental datasets could uncover new insights and improve predictive modeling. These advancements, combined with more sophisticated

modeling methods, could deepen our knowledge of complex marine ecosystems and inform future research in this essential field. In our study, we focused exclusively on changes in salinity. However, we anticipate that the methods and algorithms we used could be adapted for other bio-chemical elements, including plankton, oxygen, carbon dioxide, and more. However, this approach might introduce unique challenges related to modeling, scaling, and observation.

I am truly enthusiastic about the concept of using underwater robots for ocean exploration. It is like sending out a team of high-tech explorers into the immense, mysterious deep sea. Navigating the unpredictable seabed is a huge challenge for these robots. Once we solve that, I think the rest will follow more easily. At the moment, though, we are not quite there yet. It is a combination of anticipation and a bit of a reality check, but the potential is too exciting to ignore.

This thesis has highlighted the importance of ongoing innovation and interdisciplinary teamwork in addressing the difficulties posed by our marine environments. The potential for utilizing advanced technologies and fresh strategies to expand our knowledge of the oceans is immense. We are hopeful and accountable for further progress in this essential area of study.

Bibliography

- Ahmadzadeh, A., Jadbabaie, A., Kumar, V. and Pappas, G. J. (2006), Multi-uav cooperative surveillance with spatio-temporal specifications, *in* 'Proceedings of the 45th IEEE Conference on Decision and Control', IEEE, pp. 5293–5298.
- Antunes, H. M. and Cruz, N. A. (2019), Autonomous identification and tracking of thermoclines with a vertical profiler using extremum seeking control, *in* 'OCEANS 2019 MTS/IEEE SEATTLE', pp. 1–6.
- 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.
- Asmaa, I., Khalid, B. and Hicham, M. (2019), 'Uav control architecture', *International Journal of Advanced Computer Science and Applications* **10**(11).
- Azevedo, I. C., Bordalo, A. A. and Duarte, P. M. (2010), 'Influence of river discharge patterns on the hydrodynamics and potential contaminant dispersion in the douro estuary (portugal)', *Water research* **44**(10), 3133–3146.
- Bellman, R. (1966), 'Dynamic programming', *Science* **153**(3731), 34–37.
- Berget, G. E., Eidsvik, J., Alver, M. O. and Johansen, T. A. (2023), 'Dynamic stochastic modeling for adaptive sampling of environmental variables using an auv', *Autonomous Robots* pp. 1–20.

- Berild, M. O. and Fuglstad, G.-A. (2023), 'Spatially varying anisotropy for gaussian random fields in three-dimensional space', *Spatial Statistics* **55**, 100750.
- Chapman, C. C., Lea, M.-A., Meyer, A., Sallée, J.-B. and Hindell, M. (2020), 'Defining southern ocean fronts and their influence on biological and physical processes in a changing climate', *Nature Climate Change* **10**(3), 209–219.
- Chave, R., Buermans, J., Lemon, D., Taylor, J. C., Lembke, C., DeCollibus, C., Saba, G. K. and Reiss, C. S. (2018), Adapting multi-frequency echo-sounders for operation on autonomous vehicles, in 'OCEANS 2018 MTS/IEEE Charleston', IEEE, pp. 1–6.
- Coogan, J., Dzwonkowski, B. and Lehrter, J. (2019), 'Effects of coastal upwelling and downwelling on hydrographic variability and dissolved oxygen in mobile bay', *Journal of Geophysical Research: Oceans* **124**(2), 791–806.
- Cowles, T. and Donaghay, P. (1998), 'Thin layers: observations of small-scale patterns and processes in the upper ocean', *Oceanography* **11**(1), 2–2.
- Cressie, N. and Wikle, C. K. (2015), *Statistics for spatio-temporal data*, John Wiley & Sons.
- Cui, R., Li, Y. and Yan, W. (2015), 'Mutual information-based multi-auv path planning for scalar field sampling using multidimensional rrt', *IEEE Transactions on Systems, Man, and Cybernetics: Systems* **46**(7), 993–1004.
- Das, J., Py, F., Harvey, J. B., Ryan, J. P., Gellene, A., Graham, R., Caron, D. A., Rajan, K. and Sukhatme, G. S. (2015), 'Data-driven robotic sampling for marine ecosystem monitoring', *The International Journal of Robotics Research* **34**(12), 1435–1452.
- Devlin, M. J., Petus, C., Da Silva, E., Tracey, D., Wolff, N. H., Waterhouse, J. and Brodie, J. (2015), 'Water quality and river plume monitoring in the great barrier reef: an overview of methods based on ocean colour satellite data', *Remote Sensing* **7**(10), 12909–12941.

- Eapen, B. R., Archer, N., Sartipi, K. and Yuan, Y. (2019), Drishti: A sense-plan-act extension to open mhealth framework using flhir, *in* ‘2019 IEEE/ACM 1st International Workshop on Software Engineering for Healthcare (SEH)’, IEEE, pp. 49–52.
- Emel’yanov, S., Makarov, D., Panov, A. I. and Yakovlev, K. (2016), ‘Multilayer cognitive architecture for uav control’, *Cognitive Systems Research* **39**, 58–72.
- Field, C. B., Behrenfeld, M. J., Randerson, J. T. and Falkowski, P. (1998), ‘Primary production of the biosphere: integrating terrestrial and oceanic components’, *Science* **281**(5374), 237–240.
- Fiorelli, E., Leonard, N. E., Bhatta, P., Paley, D. A., Bachmayer, R. and Fratantoni, D. M. (2006), ‘Multi-auv control and adaptive sampling in monterey bay’, *IEEE journal of oceanic engineering* **31**(4), 935–948.
- Fossum, T. O. (2019), Adaptive sampling for marine robotics, PhD thesis, NTNU.
- 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.
- Fuglstad, G.-A., Simpson, D., Lindgren, F. and Rue, H. (2015), ‘Does non-stationary spatial data always require non-stationary random fields?’, *Spatial Statistics* **14**, 505–531.
- Galarza Bogotá, C. M. (2018), Diseño del sistema de navegación en inmersión del vehículo Guanay II para aplicaciones de detección y seguimiento de vertidos de hidrocarburos en zonas costeras, PhD thesis, Universitat Politècnica de Catalunya.
- Ge, Y., Eidsvik, J. and Mo-Bjørkelund, T. (2023), ‘3-d adaptive auv sampling for classification of water masses’, *IEEE Journal of Oceanic Engineering* **48**(3), 626–639.
- Haraldstad, T., Forseth, T., Olsen, E. M., Haugen, T. O. and Höglund, E. (2022), ‘Empirical support for sequential imprinting during downstream migration in atlantic salmon (*salmo salar*) smolts’, *Scientific Reports* **12**(1), 13736.

- Horner-Devine, A. R., Hetland, R. D. and MacDonald, D. G. (2015a), 'Mixing and transport in coastal river plumes', *Annual Review of Fluid Mechanics* **47**, 569–594.
- Horner-Devine, A. R., Hetland, R. D. and MacDonald, D. G. (2015b), 'Mixing and transport in coastal river plumes', *Annual Review of Fluid Mechanics* **47**(1), 569–594.
- Jackson, P. R. and Reneau, P. C. (2014), Integrated synoptic surveys of the hydrodynamics and water-quality distributions in two lake michigan rivermouth mixing zones using an autonomous underwater vehicle and a manned boat, Technical report, US Geological Survey.
- Karaman, S. and Frazzoli, E. (2011), 'Sampling-based algorithms for optimal motion planning', *The international journal of robotics research* **30**(7), 846–894.
- Kukulya, A. L., Bellingham, J., Stokey, R., Whelan, S., Reddy, C., Conmy, R. and Walsh, I. (2018), Autonomous chemical plume detection and mapping demonstration results with a cots auv and sensor package, in 'OCEANS 2018 MTS/IEEE Charleston', pp. 1–6.
- Lagunas, J., Marec, C., Leymarie, É., Penker'h, C., Rehm, E., Desaulniers, P., Brousseau, D., Larochelle, P., Roy, G., Fournier, G. et al. (2018), Sea-ice detection for autonomous underwater vehicles and oceanographic lagrangian platforms by continuous-wave laser polarimetry, in 'Ocean Sensing and Monitoring X', Vol. 10631, SPIE, pp. 250–264.
- Lin, M. and Yang, C. (2020), 'Ocean observation technologies: A review', *Chinese Journal of Mechanical Engineering* **33**(1), 1–18.
- Lindgren, F. and Rue, H. (2015), 'Bayesian spatial modelling with r-inla', *Journal of statistical software* **63**(19).
- Mańko, M. K., Merchel, M., Kwasniewski, S. and Weydmann-Zwolicka, A. (2022), 'Oceanic fronts shape biodiversity of gelatinous zooplankton in the european arctic', *Frontiers in Marine Science* **9**, 941025.

- McCammon, S., Marcon dos Santos, G., Frantz, M., Welch, T. P., Best, G., Shearman, R. K., Nash, J. D., Barth, J. A., Adams, J. A. and Hollinger, G. A. (2021), ‘Ocean front detection and tracking using a team of heterogeneous marine vehicles’, *Journal of Field Robotics* **38**(6), 854–881.
- Mendes, R., Sousa, M. C., DeCastro, M., Gómez-Gesteira, M. and Dias, J. M. (2016), ‘New insights into the western iberian buoyant plume: Interaction between the douro and minho river plumes under winter conditions’, *Progress in Oceanography* **141**, 30–43.
- 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–US Gulf Coast’, IEEE, pp. 1–5.
- Nakhaeinia, D., Tang, S. H., Noor, S. M. and Motlagh, O. (2011), ‘A review of control architectures for autonomous navigation of mobile robots’, *International Journal of the Physical Sciences* **6**(2), 169–174.
- Oddi, A., Rasconi, R., Santucci, V. G., Sartor, G., Cartoni, E., Mannella, F. and Baldassarre, G. (2020), Integrating open-ended learning in the sense-plan-act robot control paradigm, in ‘ECAI 2020’, IOS Press, pp. 2417–2424.
- Powell, W. B. (2007), *Approximate Dynamic Programming: Solving the curses of dimensionality*, Vol. 703, John Wiley & Sons.
- Quigley, M. (2009), Ros: an open-source robot operating system, in ‘ICRA 2009’.
- Ratsimandresy, A. W., Donnet, S., Goulet, P., Bachmayer, R. and Claus, B. (2014), Variation in the structure of the water column as captured by slocum glider ctd and by ctd from a research vessel and assessment of internal waves, in ‘2014 Oceans - St. John’s’, pp. 1–10.
- Rhein, M., Rintoul, S., Aoki, S., Campos, E., Chambers, D., Feely, R., Gulev, S., Johnson, G., Josey, S., Kostianoy, A. et al. (2013), ‘Observations: Ocean’, *IPCC AR5 Climate Change Report 2013* .

- Sagar, S., Falkner, I., Dekker, A., Huang, Z., Blondeau-Patissier, D., Phillips, C., Przeslawski, R. and Assessment, C. (2020), 'Earth observation for monitoring of australian marine parks and other off-shore marine protected areas', *Report to the National Environmental Science Program—Marine Biodiversity Hub* .
- 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.
- Takahashi, T., Sutherland, S. C., Wanninkhof, R., Sweeney, C., Feely, R. A., Chipman, D. W., Hales, B., Friederich, G., Chavez, F., Sabine, C. et al. (2009), 'Climatological mean and decadal change in surface ocean pco₂, and net sea–air co₂ flux over the global oceans', *Deep Sea Research Part II: Topical Studies in Oceanography* **56**(8-10), 554–577.
- Thorsnes, T., Chand, S., Brunstad, H., Lepland, A. and Lågstad, P. (2019), 'Strategy for detection and high-resolution characterization of authigenic carbonate cold seep habitats using ships and autonomous underwater vehicles on glacially influenced terrain', *Frontiers in Marine Science* **6**.
URL: <https://www.frontiersin.org/articles/10.3389/fmars.2019.00708>
- Toberman, M., Inall, M., Boyd, T., Dumont, E. and Griffiths, C. (2017), 'Nonlinear internal waves and plumes generated in response to sea-loch outflow, auv, and time-lapse photography observations', *Journal of Geophysical Research: Oceans* **122**(7), 5522–5544.

Paper I

3-D Adaptive AUV Sampling for Classification of Water Masses

Yaolin Ge, Jo Eidsvik, Tore Mo-Bjørkelund

Published to IEEE Journal of Oceanic Engineering on 26.04.2023

3-D Adaptive AUV Sampling for Classification of Water Masses

Yaolin Ge ¹, Jo Eidsvik ², and Tore Mo-Bjørkelund

Abstract—Autonomous underwater vehicles with onboard computing units foster innovative approaches for sampling oceanographic phenomena. Feedback of observations via the onboard model for planning algorithms enable adaptive sampling for such robotic units. In this work, we develop, implement, and test an adaptive sampling algorithm for efficient sampling of water masses in a 3-D frontal system. Focusing on a river plume, salinity variations are used to characterize the water masses. A threshold in salinity is assumed to distinguish the ocean and river waters, so that excursions below the threshold define river waters. The onboard model builds on a Gaussian random field representation of the salinity variations in (north, east, depth) coordinates. This model is initially trained from numerical ocean model data, and then updated with data gathered by the vehicle sensor. The Gaussian random field model further allows closed-form expressions of the expected spatially integrated Bernoulli variance of the salinity excursion set, which is used to reward sampling efforts. Combining these results with forward-looking planning algorithms, we suggest a workflow for 3-D adaptive sampling to map river plume systems. Simulation studies are used to compare the suggested approach with others. Results of field trials in the Nidelva river plume in Norway are presented and discussed.

Index Terms—Adaptive sampling, autonomous underwater vehicles (AUVs), excursion sets (ES), path planning, river plume.

I. INTRODUCTION

A RIVER plume is formed when the fresh water flowing out of the river encounters the saline water in the ocean [1]. When these two different water masses meet, they form a varying spatio-temporal boundary [2]. There have been increasing efforts using numerical models and data to investigate such phenomena in the past decades [3], [4], [5], [6], [7], [8].

Autonomous underwater vehicles (AUVs) with onboard sensors and computing resources provide rich opportunities for oceanographic sampling as they can calibrate numerical ocean model outputs with in situ data, and fill in the sampling resolution gaps at locations with large uncertainty [9], [10], [11], [12]. For

Manuscript received 10 February 2022; revised 2 November 2022; accepted 18 February 2023. Date of publication 26 April 2023; date of current version 14 July 2023. This work was supported by Norwegian Research Council (RCN) through the MASCOT project under Grant 305445. (Corresponding author: Yaolin Ge.)

Associate Editor: N. Cruz.

Yaolin Ge and Jo Eidsvik are with the Department of Mathematical Sciences, Norwegian University of Science and Technology, 7034 Trondheim, Norway (e-mail: yaolin.ge@ntnu.no; jo.eidsvik@ntnu.no).

Tore Mo-Bjørkelund is with the Department of Marine Technology, Norwegian University of Science and Technology, 7034 Trondheim, Norway (e-mail: tore.mo-bjorkelund@ntnu.no).

Digital Object Identifier 10.1109/JOE.2023.3252641

frontal regions such as river plumes, AUV sampling is helpful for classifying the different water masses more accurately. Previous AUV sampling efforts focus mainly on preprogrammed designs [13] or use event-triggered adaptation of designs [14], [15]. Recent efforts have shown added value of having model-based adaptive sampling plans [16]. Adaptive sampling strategies here refer to AUV planning schemes that enable the AUV plan to be updated based on the posterior knowledge from in situ sampling and the probabilistic model description. Ideas from statistical sampling design are highly useful in this field, because they can help guide the AUV to informative locations [16], [17].

The main contribution of this work is a 3-D full-scale adaptive AUV sampling strategy. With the AUVs limiting computing resources, a Gaussian random field (GRF) model serves as a statistical proxy models for the spatial salinity field in the 3-D domain (north, east, depth). This 3-D GRF model running onboard the AUV is sequentially refined using in situ observations. This refined probabilistic model is further a basis for evaluating AUV sampling designs. Starting with prior knowledge from a numerical ocean model, we use an AUV to adaptively explore the 3-D boundary between the water masses in the river plume. We suggest algorithms to speed up design computations and to enable efficient robotic maneuverability [18]. We use a statistical design criterion based on the uncertainty of the excursion set (ES) of low salinity which distinguishes the river from the ocean water. This ES is defined by spatial locations having salinity level below a user-defined threshold. Building on recently developed closed form expressions [16] for the expected integrated Bernoulli variance (EIBV) associated with the ES, we compare the EIBV associated with each candidate design location, and select the design which has the minimum EIBV. The EIBV is a useful criterion for improved classification of the river plume as it is large when probabilities of excursions are far from 0 and 1. One should select sampling designs that on expectation pull probabilities toward the 0 and 1 end-points to reduce the uncertainty of the ES.

Via simulation studies and in situ measurements from the Nidelva river plume in Trondheim, Norway, we study the properties of the EIBV sampling plans in the 3-D domain. For the real-world experiments we used a light AUV (LAUV) [19] with an on-board NVIDIA Jetson TX2 computing unit.

The rest of this article is organized as follows. In Section II, we provide the background and motivation for our work on adaptive AUV sampling to river plume water masses characterization. In Section III, we introduce the models and methods used in this article. In Section IV, we present our implementation used for

path planning. In Section V, we show a simulation study illustrating the properties of our 3-D adaptive sampling approaches. In Section VI, we show results from the Nidelva river plume experiments. Finally, Section VII concludes this article.

II. OCEAN SAMPLING

A. Data Sources

Numerical solutions of the complex differential equations governing spatio-temporal oceanographic variation with boundary conditions and forcing are essential in understanding the ocean variability. In our application we rely on a fjord-scale implementation of the SINMOD software [20]. Such ocean model data provide physical interpretability of the ocean variability, but they often need calibration or bias adjustments, and there have been growing interests in uncertainty quantification and data assimilation methods for various scales of this challenge, see, e.g., [21].

Traditional in situ measurements generating input or calibration data to numerical ocean models include stationary or floating buoys, gliders, moorings, and ships [22]. With the advent of smaller inexpensive sensor systems, one has capabilities of handling a variety of measurements for biological, chemical, and oceanographic purposes [22]. Ships data can be expensive, and buoys and gliders have limited flexibility in maneuverability given coverage constraints [23].

Satellite imagery has been a powerful and useful tool for analyzing ocean variables. Data from satellites can provide a large-scale coverage of the entire field of interest, and even output portraits of river plumes [6]. However, due to large latency and uncertainty (cloud coverage issues) of obtaining such images, the usage of satellite imagery is limited. Satellite data unfavorably cover only the surface of the ocean [24].

The development of underwater robotics have led to a large number of robot-assisted applications in oceanography. Thanks to the flexibility of the robots, there are growing numbers of autonomous sampling missions which are conducted by robots [9]. Benefits further include real-time sensing and high-resolution data gathering, with large opportunities to move in flexible paths in the ocean environment. In our case, an AUV is used as the target platform which is able to support 3-D adaptive sampling at high resolution.

B. Sequential AUV Sampling

We denote the salinity field by $\{\xi_{\mathbf{u}}; \mathbf{u} \in \mathcal{M} \subset \mathcal{R}^3\}$, where the location \mathbf{u} is (longitude, latitude, depth) and \mathcal{M} is the spatial domain of interest. Initially, we specify a probabilistic model for the salinity based on numerical ocean model data. This provides a realistic initial model for the 3-D salinity characteristics, one that it is much more physically inspired than a simple linear regression from available in situ AUV data [16]. We still use regression analysis to calibrate the 3-D ocean model data to the real-world ocean experiment by using a short preliminary AUV survey [25]. The objective of the survey is not to reveal the entire field, but rather provide some in situ measurements to adjust

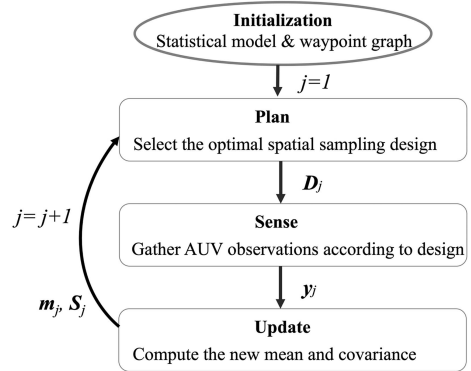


Fig. 1. Sequential loop where design D_j is chosen based on the updated model, y_j is the data collected in this design, and this is used to update the model (m_j, S_j) . This continues over stages $j = 1, 2, \dots, J$.

the ocean-model data and to form a reasonable prior model for the day of deployment. Therefore, the path for the preliminary survey can be as simple as a transect line with yo-yo movements in the vertical direction. As mentioned in the previous section, one can also use satellite data or even drone images in this initial model specification, if such data are available [26].

In situ salinity observations for the main part of the deployment are denoted by $\{y_j; j = 1, \dots, J\}$, for stages j of AUV measurements gathered over time. The vector y_j of measurements at stage j , holds N_j measurements made according to spatial sampling design D_j . The initial deployment location will then define D_1 . We denote by $\mathcal{Y}_j = \{(y_1, D_1), \dots, (y_j, D_j)\}$ the collection of data gathered with the selected designs up to stage j . Initially, this is an empty set; $\mathcal{Y}_0 = \emptyset$.

The sequential designs are selected adaptively based on what is evaluated to be the most informative AUV sampling locations. In this evaluation, the on-board model is conditional to all the data gathered until the current time. With new observations available, data assimilation methods are used to update the probabilistic representation for the salinity variables. This means that the model is “alive,” and changing at every stage, depending on the data. Adaptive sampling fits into the diagram loop in Fig. 1. In our setting the spatial design plan is optimized based on the current spatial statistical model. Then the AUV gathers new observations according to the chosen design, and the GRF model is updated. This continues over stages $j = 1, \dots, J$.

For prioritizing sampling efforts, one must impose an expected reward or value function associated with the different available sampling designs. At each stage, the expected rewards of all possible designs are evaluated. In our setting with river plumes, it makes sense to reward sampling locations that are expected to give data that improve the spatial characterization of the water masses [15], [16]. The setting is illustrated in Fig. 2, where we indicate the current location of the AUV, its path, and the sampling design opportunities at this stage. The information criterion (EIBV) is calculated for all feasible designs, shown as circular dots. Here, smaller dots with lighter colors are indicative




Fig. 2. Adaptive path example on a 3-D waypoint graph. There are 17 candidate locations in different layers, the blue dot shows the current AUV location whereas the green dot indicates the desired next waypoint selected based on the minimum EIBV criterion.

of larger expected uncertainty reduction. The adaptive sampling approach would act by moving to the location with lowest EIBV.

III. STATISTICAL MODELS AND METHODS FOR AUV SAMPLING

We next discuss our probabilistic modeling choices for the salinity field, and show how this enables efficient data assimilation as well as onboard design criteria. We then define ES and the EIBV as a design criterion, and finally present an adaptive sampling design algorithm for efficient 3-D characterization of the river plume.

A. On-Board Computing With GRFs

The prior model for river plume salinity $\{\xi_{\mathbf{u}}, \mathbf{u} \in \mathcal{M} \subset \mathcal{R}^3\}$, is defined via a GRF. A working assumption in our work is hence that the GRF provides a reasonable proxy model for the spatial salinity field in (latitude, longitude, depth). The initial model specification includes estimating the expected value of the field, its variability and spatial dependence. Note that the duration of the experiment will be short and the temporal variation in the river plume is ignored here. To check the Gaussian assumption, we made a quantile-quantile (QQ) plot from the SINMOD salinity data (see Fig. 3). Here, we have computed the mean and variance at each location in a gridded domain over replicates of time steps. The standardized residuals are used in the QQ plot. The QQ plot in Fig. 3 shows a crossplot of the theoretical Gaussian quantile of the residuals against the empirical quantile of residuals in the data set. The blue line that we achieve is quite close to the straight line (red). Of course, the physical model does not give a Gaussian model, and we notice a sharper distribution near 0, but nevertheless the discrepancy is rather small.

Critically, the GRF model enables onboard data assimilation and adaptive AUV sampling efforts, as we will describe next.




Fig. 3. Quantile-quantile plot of the residual based on SINMOD estimation. The residual is computed by subtracting the mean of the field and dividing the standard deviation.

For onboard implementation and computing, the spatial domain is discretized to a set of n grid locations; $\{\mathbf{u}_1, \dots, \mathbf{u}_n\}$. This grid is also used for the waypoint graph setting for the AUV sampling design. The prior or initial GRF model at these grid locations is denoted by

$$\xi = (\xi_{\mathbf{u}_1}, \dots, \xi_{\mathbf{u}_n})^T, \quad \xi \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (1)$$

with associated probability density function (PDF) $p(\xi)$. Here, length- n vector $\boldsymbol{\mu}$ represents the prior mean of the 3-D salinity variations, as will later be specified from ocean model data and a preliminary AUV transect run. The $n \times n$ covariance matrix $\boldsymbol{\Sigma}$ is defined via a Matérn covariance function with elements $\Sigma(i, i') = \sigma^2(1 + \phi_1 h(i, i')) \exp(-\phi_1 h(i, i'))$, where σ^2 is the variance and ϕ_1 a correlation decay parameter [27]. The distance between grid nodes \mathbf{u}_i and $\mathbf{u}_{i'}$ is defined for east, north, and depth Euclidean distances via

$$h^2(i, i') = h_E^2(i, i') + h_N^2(i, i') + (\phi_1^2/\phi_2^2)h_D^2(i, i')$$

with h being distance, and subscripts E, N, D indicating each of the three directions in vector $\mathbf{u}_v - \mathbf{u}_i$. Studies have shown that the lateral stretch of the river plume tends to be many magnitudes above the vertical stretch [1]. To model the correlation in different dimensions properly, we employ anisotropy between the lateral domain and the vertical domain. This means that the depth dimension is scaled differently (ϕ_2) using another correlation decay parameter than the one used in the lateral field (ϕ_1).

The measurements at each stage $j = 1, \dots, J$ are modeled by a Gaussian likelihood model

$$\mathbf{y}_j | \boldsymbol{\xi} \sim N(\mathbf{F}_j \boldsymbol{\xi}, \mathbf{R}_j) \quad (2)$$

where \mathbf{F}_j is an $N_j \times n$ selection matrix containing an entry of 1 in each row and 0 otherwise. The 1 entry refers to the sampling indices. With the covariance matrix $\mathbf{R}_j = r^2 \mathbf{I}_{N_j}$, we assume that the data are conditionally independent, given the underlying salinity. Here, r indicates the measurement standard deviation of the AUV salinity observations. We denote the associated PDF by $p(\mathbf{y}_j | \boldsymbol{\xi})$.

Via Bayes' rule, data assimilation at stages $j = 1, \dots, J$, gives the sequential conditional PDF $p(\boldsymbol{\xi} | \mathcal{Y}_j) \propto p(\mathbf{y}_j | \boldsymbol{\xi}) p(\boldsymbol{\xi} | \mathcal{Y}_{j-1})$. Under the assumptions about a GRF prior model and a Gaussian measurement error model, this conditional PDF is also Gaussian with mean \mathbf{m}_j and covariance matrix \mathbf{S}_j given by

$$\begin{aligned} \mathbf{G}_j &= \mathbf{S}_{j-1} \mathbf{F}_j^T (\mathbf{F}_j \mathbf{S}_{j-1} \mathbf{F}_j^T + \mathbf{R}_j)^{-1} \\ \mathbf{m}_j &= \mathbf{m}_{j-1} + \mathbf{G}_j (\mathbf{y}_j - \mathbf{F}_j \mathbf{m}_{j-1}) \\ \mathbf{S}_j &= \mathbf{S}_{j-1} - \mathbf{G}_j \mathbf{F}_j \mathbf{S}_{j-1} \end{aligned} \quad (3)$$

where $\mathbf{m}_0 = \boldsymbol{\mu}$ and $\mathbf{S}_0 = \boldsymbol{\Sigma}$. The sequential updating resembles that of a spatio-temporal Kalman filter [27]. In our case, we study the benefits of using a 3-D spatial model in the AUV sampling. Having a relatively short-term deployment, no explicit temporal dynamics are modeled.

B. Excursion Set and Expected Integrated Bernoulli Variance

We use the notion of an ES to characterize the river and ocean water masses [16]. The ES for salinity threshold t is defined by

$$\text{ES} = \{\mathbf{u} \in \mathcal{M} : \xi_{\mathbf{u}} \leq t\}. \quad (4)$$

Hence, salinity lower than this threshold will indicate river water. The associated excursion probability (EP) is

$$p_{\mathbf{u}} = P(\xi_{\mathbf{u}} \leq t), \quad \mathbf{u} \in \mathcal{M}. \quad (5)$$

When it is close to 1 or 0 at a given location, it is easy to classify the water mass to be river or ocean respectively. EP close to 0.5 reflects ambiguity in the characterization of water masses. The prior Bernoulli variance (BV) at location \mathbf{u} is $p_{\mathbf{u}}(1 - p_{\mathbf{u}})$ and the spatially integrated BV (IBV) is

$$\text{IBV} = \int p_{\mathbf{u}}(1 - p_{\mathbf{u}}) d\mathbf{u} \quad (6)$$

which is dominated by locations with probabilities near 0.5 and BV close 0.25. In practice, the integral will be approximated by a sum over the n grid nodes.

The goal is to construct AUV sampling strategies that prioritize locations that are ambiguous, thus, making the exploration more effective. At each stage, we define the EIBV by

$$\text{EIBV}(\mathbf{D}_j) = \int E_{\mathbf{y}_j | \mathcal{Y}_{j-1}; \mathbf{D}_j} [B_{\mathbf{u}}(\mathbf{y}_j)] d\mathbf{u}$$

$$B_{\mathbf{u}}(\mathbf{y}_j) = p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1})(1 - p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1})) \quad (7)$$

where $B_{\mathbf{u}}(\mathbf{y}_j)$ is the conditional Bernoulli variance for outcome \mathbf{y}_j of data in design \mathbf{D}_j , and the conditional probability of an excursion is

$$p_{\mathbf{u}}(\mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1}) = P(\xi_{\mathbf{u}} \leq t | \mathbf{y}_j, \mathbf{D}_j, \mathcal{Y}_{j-1}). \quad (8)$$

The notation in (7) indicates that the EIBV is an expectation with respect to the random data \mathbf{y}_j for design \mathbf{D}_j , conditional on the history of sampling results \mathcal{Y}_{j-1} .

The criterion for selecting design \mathbf{D}_j and then getting data \mathbf{y}_j at stage $j = 1, \dots, J$, is based on the minimum EIBV computed for all designs in a candidate waypoint set denoted \mathcal{D}_j . We have

$$\mathbf{D}_j = \underset{\mathbf{D}'_j \in \mathcal{D}_j}{\text{argmin}} \text{EIBV}(\mathbf{D}'_j). \quad (9)$$

Using expressions similar to that of [28], the EIBV in (7) can be evaluated in closed form. Denoting the variance reduction from data by $\mathbf{V}_j = \mathbf{G}_j \mathbf{F}_j \mathbf{S}_{j-1}$, see (3), the EIBV becomes

$$\begin{aligned} \text{EIBV}(\mathbf{D}'_j) &= \sum_{i=1}^n \text{EBV}_{\mathbf{u}_i}(\mathbf{D}'_j) \\ \text{EBV}_{\mathbf{u}_i}(\mathbf{D}'_j) &= \Phi_2 \left(\begin{bmatrix} t \\ -t \end{bmatrix}; \begin{bmatrix} m_{j-1}(i) \\ -m_{j-1}(i) \end{bmatrix}, \mathbf{W}_j(i, i) \right) \end{aligned} \quad (10)$$

where Φ_2 denotes the bivariate Gaussian cumulative distribution function, and with

$$\begin{aligned} \mathbf{W}_j(i, i) &= \begin{bmatrix} T(i, i) & -V_j(i, i) \\ -V_j(i, i) & T(i, i) \end{bmatrix} \\ T(i, i) &= S_j(i, i) + V_j(i, i). \end{aligned}$$

We next give some intuition for this EIBV criterion. Fig. 4 illustrates a Gaussian PDF (left) representing the current knowledge about salinity at some location. In this case it is standardized so that

$$Z_1 = \frac{\xi_{\mathbf{u}_i} - m_{j-1}(i)}{\sqrt{S_{j-1}(i, i)}}$$

for location \mathbf{u}_i . The scaled threshold $t - m_{j-1}(i)$ is shown as a vertical line. With variance $s_{j-1}^2(i) = S_{j-1}(i, i)$, the current BV $= p(1 - p)$, $p = \Phi(\xi_{\mathbf{u}_i}; m_{j-1}(i), s_{j-1}^2(i))$ is also displayed.

We can collect data and get more information. The expected BV (EBV) at this location is then available as a cumulative probability as indicated in Fig. 4 (right). The EBV depends on the mean value relative to the threshold. Assume that the mean is lower, meaning that the threshold $t - m_{j-1}(i)$ moves to the right in the left display. Then the BV decreases, and the EBV illustrated in the right display also decreases as the vertical line moves right and the horizontal line moves down. The EBV is further smallest when there is much negative correlation in the




Fig. 4. Left: Density curve represents the current knowledge at a selected location, while the vertical line indicates the threshold. The Bernoulli variance (BV) is indicated. Right: EBV calculation involves bivariate Gaussian cumulative probabilities, which is the volume below the contours in the bottom left region.

Algorithm 1: Informative Myopic Sampling Algorithm.

Initialization: $\mathbf{m}_0, \mathbf{S}_0, t, \mathcal{Y}_0 = \emptyset, \mathcal{D}_1$
 $j = 1$
while $j \leq N_{steps}$ **do**
 Plan: Evaluate $EIBV(\mathcal{D}'_j)$ for all $\mathcal{D}'_j \in \mathcal{D}_j \triangleright (7)$ and (10)
 $\mathcal{D}_j = \text{argmin}_{\mathcal{D}'_j \in \mathcal{D}_j} EIBV(\mathcal{D}'_j) \triangleright (9)$
 Go to design \mathcal{D}_j with the AUV, set design matrix \mathbf{F}_j , form set \mathcal{D}_{j+1} .
 Sense: Gather in situ AUV data \mathbf{y}_j according to design \mathcal{D}_j .
 $\mathcal{Y}_j = (\mathcal{Y}_{j-1}, \mathbf{y}_j)$.
 Update : $\mathbf{G}_j = \mathbf{S}_{j-1} \mathbf{F}_j^T (\mathbf{F}_j \mathbf{S}_{j-1} \mathbf{F}_j^T + \mathbf{R}_j)^{-1}$
 $\mathbf{m}_j = \mathbf{m}_{j-1} + \mathbf{G}_j (\mathbf{y}_j - \mathbf{F}_j \mathbf{m}_{j-1})$,
 $\mathbf{S}_j = \mathbf{S}_{j-1} - \mathbf{G}_j \mathbf{F}_j \mathbf{S}_{j-1} \triangleright (3)$
 $j = j + 1$
end while

density in Fig. 4 (right). From matrix $\mathbf{W}_j(i, i)$ in (11), we see that this occurs when the variance reduction $V_j(i, i)$ is large compared with $S_j(i, i) + V_j(i, i)$. The bivariate Φ_2 calculation in (10) is somewhat costly, and if the correlation term is small, one could approximate it with two univariate calculations to gain computational efficiency.

Previous research has demonstrated the possibility of using EIBV as the design criterion for AUV adaptive sampling in 2-D domains [16]. We next explain how we build on this to construct effective AUV operations in 3-D adaptive sampling plans.

IV. PATH PLANNING ALGORITHM

A. Adaptive Sampling

The GRF model updating in (3) and closed form EIBV calculation in (10) enable information-based adaptive AUV sampling. We summarize the approach in Algorithm 1.

Note that as outlined this defines a myopic or greedy approach to adaptive sampling. This is not necessarily optimal. The myopic evaluation is done by taking the expectation of data at this stage only, without anticipation of what future sampling efforts might bring. The optimal solution to the sequential sampling design problem would also account for the sampling efforts at future stages. However, from the mathematical and computational setting, it is not feasible to find the optimal design strategy because it involves combinatorial growth of possible paths requiring intermixed optimization and expected values. Instead, one often resorts to the outlined myopic strategy. More nuanced approaches exist for doing longer-horizon search, for instance variants of Markov decision processes (MDPs) or partially observed MDPs [29], rapidly exploring random trees [30] or those based on genetic algorithms [31]. Such approaches will typically perform better than the myopic heuristic in situations with forbidden regions or with high collision risks, but it is not easy to use these in large-scale computations onboard the AUV. Further, restricted Monte Carlo search or pruning of paths, these nonmyopic approaches will not necessarily improve performance compared with a myopic search on the regular waypoint graph case [16]. We will limit scope to the myopic calculations (Algorithm 1) in this work.

For the 3-D application we consider here, the sequential sampling is restricted to a path embedded on a predefined grid of waypoints. In practice, the EIBV is computed for a set of neighborhood waypoint locations, meaning that the candidate design \mathcal{D}'_j must be among those possible designs defining \mathcal{D}_j .

For small AUVs and large field, it might be possible to move the AUV wherever it needs to be. However, this might lead to an excess of manoeuvring time for the operation. To foster efficiency of the autonomous sampling process, a smooth-filtering method is applied to achieve AUV-friendly path planning (Algorithm 2). It first selects neighboring locations, and two vectors will be formed. Vector \vec{b}_1 is defined from the previous location to the current location, whereas vector \vec{b}_2 is from the current location to the potential candidate locations. Next, the inner products between these two vectors is calculated, and




Fig. 5. Smooth path planning. I: Arrive at the current location. II: Search all neighboring locations. III, IV: Compute inner products. V: Select qualified candidate locations. Blue thin crosses indicate the abandoned locations, whereas the green thick crosses indicate the filtered locations.

Algorithm 2: Smooth-Filtering Algorithm.

Require: D_{j-1}, D_{j-2}
 $D^* = \{u \in \mathcal{M} \text{ such that } |u - D_{j-1}| < \text{neighboring distance}\}$
 $\vec{b}_1 = D_{j-1} - D_{j-2}$
 $i = 1$
while $i \leq N_{D^*}$ **do**
 $\vec{b}_2 = D_i^* - D_{j-1}$
 if $\vec{b}_1 \cdot \vec{b}_2 < 0$ **then**
 Abandon D_i^* .
 end if
 $i = i + 1$
end while
 $D_j = D^*$

only candidate locations with positive inner products will be considered for EIBV evaluation.

A map view version of the smooth-filtering is depicted on a 2-D waypoint graph in Fig. 5. In 3-D, the principle is the same, except that it is expanded to include the vertical candidate locations as well. This path smooth-filtering algorithm is effective since it removes locations which might require a hydrobatic maneuver to go there [32]. The smooth-filtered trajectory further avoids time-consuming turning which would increase the traveling time and introduce location inaccuracy.

V. SIMULATION STUDY

To compare the performance between some existing algorithms and the 3-D myopic algorithm that we have developed here, a simulation study is conducted. We next describe the case, present the various methods, and discuss results.

A. Simulation Setup

We use data from the numerical ocean model SINMOD as a reference for specifying realistic trends and variabilities for the oceanographic fjord-river water masses. Fig. 6 shows the average surface salinity field predicted for the first week in May using SINMOD. Four outlets from the river are recognized. The salinity variation from the river mouth to the ocean changes dramatically from bins of [0, 3] to [28, 30] ppt. The boundary between the freshwater and the more saline fjord water is clearly depicted by the contours.




Fig. 6. Regional average surface salinity prediction in May 2021 from SINMOD. The blue rectangle indicates the designated simulation area (Section V), and the red dotted regions indicate the waypoint graph used in the field deployment (Section VI). The grid consists of 25×25 nodes in each lateral axis and 5 layers in depth. Courtesy of SINTEF Ocean and ESRI basemap.

To narrow down the focus on mapping the front of the river plume in 3-D, a smaller region of interest in the easternmost part is selected (see blue rectangle in Fig. 6). Five depth layers 0.5, 1.0, 1.5, 2.0, 2.5 m are used.

A 3-D GRF benchmark field is created based on the data extracted from SINMOD on the desired simulation region. The mean values are set from averaging SINMOD data. The coefficients used in the Matérn covariance kernel are specified as $\sigma = 0.71$, $\phi_1 = 0.008$, $\phi_2 = 2.25$, and $r = 0.2$.

Fig. 7 shows one realization from our GRF model with the specified mean and covariance model. This is regarded as the ground truth in the simulation. There is clearly river plume areas to the south-east and near the surface, and realistic variability in salinity extent with some mixing of water masses, indicating that the GRF model emulates the physical phenomenon rather well.

B. Simulation Approaches

We next describe two additional sampling strategies that are compared with our suggested 3-D adaptive sampling method. In

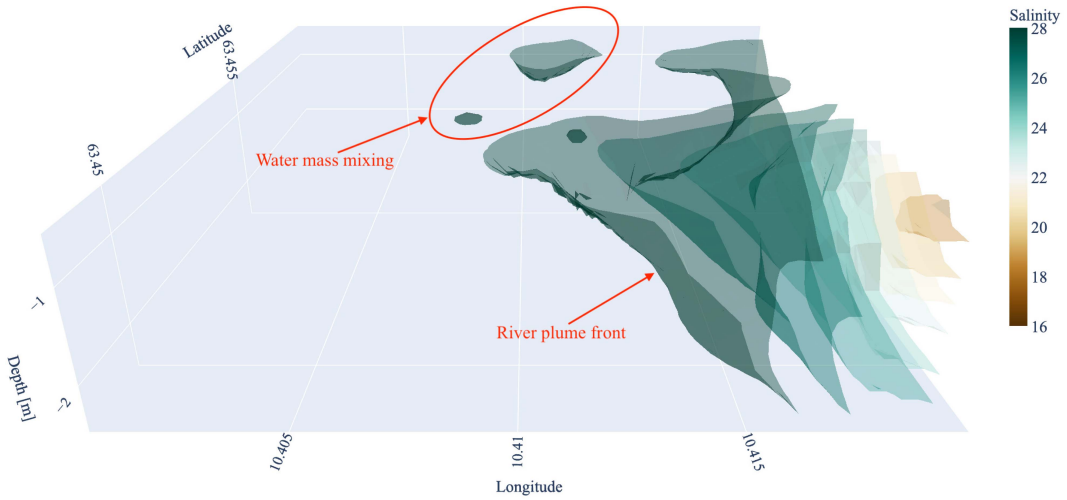


Fig. 7. One benchmark salinity field used in the simulation study. Some water blobs are shown on the north side of the region.




Fig. 8. Adaptive myopic 2-D algorithm illustration. The outmost envelope shows the estimated boundary after sampling 20 locations. Note that the yoyo pattern is shown as an illustration. It can be denser in the actual setting.

all three, the GRF proxy model provides an easy way to update the knowledge of the field by measuring the data at specified locations. The differences occur in how the data is included in the on-board computing and in what sampling strategy is used to explore the domain. When we compare results of the various approaches, they will be influenced by the sampling methodology used.

1) *Adaptive Myopic 2-D*: For the adaptive myopic 2-D, the AUV is only moving adaptively in the middle layer with the myopic strategy. It updates the entire field based on the data obtained from the middle layer at 1.5 m depth. In practice, the AUV needs to calibrate its navigational errors

by constantly popping up onto the surface and request accurate GPS locations and dive back to the place where it should continue. This is achieved by a yoyo pattern, as shown in Fig. 8.

2) *Nonadaptive Lawnmower*: For the nonadaptive lawnmower, Fig. 9 shows that the AUV will follow a predesigned 3-D lawnmower pattern. In the lateral direction, the surface-projected trajectory will be a typical lawnmower manoeuvre. To extend it into 3-D, a vertical yoyo maneuver is added in addition to the lateral lawnmower. This preprogrammed method requires no statistical computations at waypoints, and it uncovers the field with large coverage. But the approach is usually time-consuming

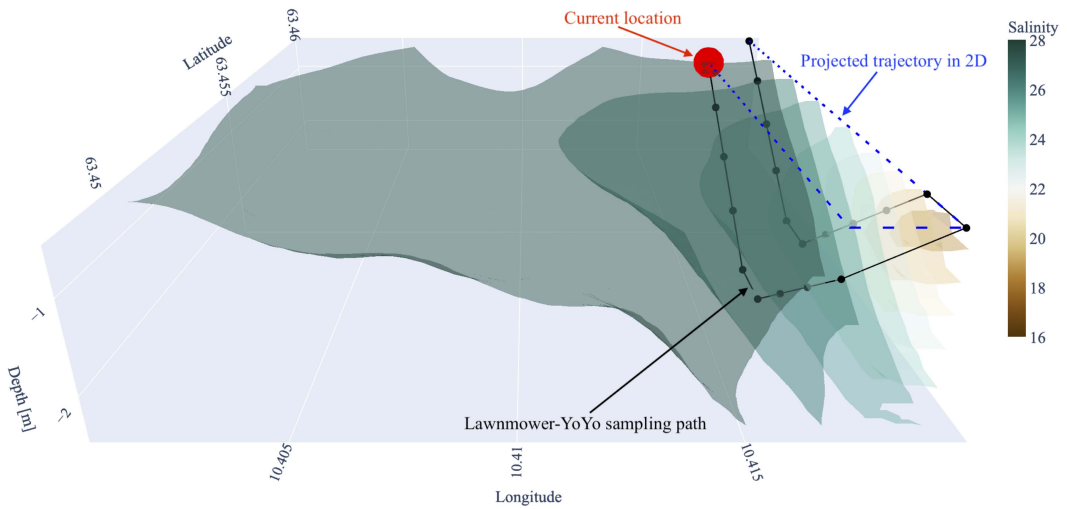


Fig. 9. Lawnmower-yoyo maneuver illustration. The estimated boundary after observing 20 sampling locations is shown as the outermost envelope.

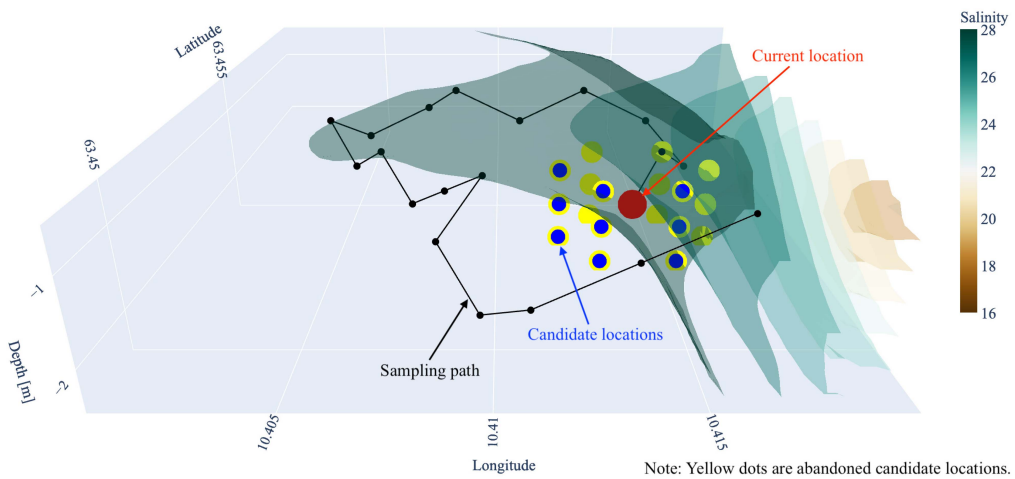


Fig. 10. Adaptive myopic 3-D sampling illustration. The outermost envelope shows the estimated river plume front after sampling 20 locations with the adaptive myopic 3-D path planning.

and inefficient in finding interesting features as it does not adapt to the data.

3) *Adaptive Myopic 3-D*: Our suggested adaptive myopic 3-D strategy extends the potential candidate sampling locations from one layer to include multiple layers. Therefore, it adapts to the field data with a much wider perspective. It is further both energy-efficient and time-efficient. One example of the adaptive 3-D myopic path planning is depicted in Fig. 10. One can see that at each stage, candidate locations will be generated in three dimensions. Only a few (shown as blue in Fig. 10) will be selected for the EIBV calculation due to the constraints of AUV maneuverability.

C. Simulation Results and Discussion

Figs. 8–10 show how each strategy behaves for one specific generated salinity field. To remove random effects, results of 100 replicate simulation results are averaged and shown in Fig. 11. At each time step of the runs, integrated Bernoulli variance (IBV), root-mean-squared error (RMSE), variance reduction and distance traveled are monitored for comparison of the three strategies.

The IBV indicator shows that the Lawnmower-yoyo pattern has the slowest reduction of the three strategies. However, it goes down quickly when the robot is in the area of interest, i.e., the boundary region or the front of the river plume, performing

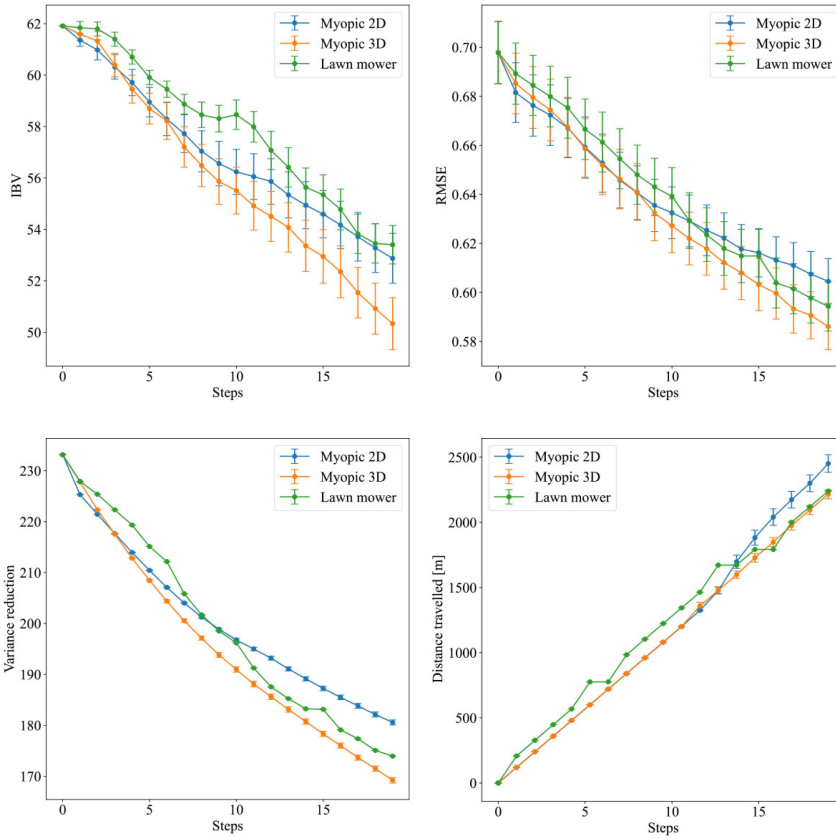


Fig. 11. Average results from 100 replicate simulations for 20 sampling locations. The standard error is depicted as vertical lines.

better than myopic 2-D after about 15 iterations (The same holds for RMSE and variance reduction.) This occurs because the lawn mower strategy can get lucky and the AUV runs into interesting parts of the domain, but it can also miss this entirely in the given time window. Even though the myopic 3-D strategy is guided by EIBV reduction, it also achieves large reduction in RMSE and variance, and more so than the other methods. It performs better than the 3-D lawnmower strategy because it explores new parts of the domains and in doing so avoids locations that are highly correlated to the ones already sampled.

With the same starting location and about the same traveled distance (see Fig. 11, lower right), the 3-D version of the myopic planning reveals the most information of the field within the three strategies. The flexibility in 3-D enables the AUV to both explore and exploit the environment effectively.

VI. AUV EXPERIMENTS IN THE NIDELVA PLUME

We next describe and show results of AUV experiments done in late Spring 2021 to map the Nidelva River plume, Trondheim, Norway. The adaptive AUV experiments were conducted on 6 July 2021. Before that, we gathered various complementary data. The phone footage on 27 May shows a visible river plume (see

Fig. 12). A satellite image on 2 June (see Fig. 13) shows how the river plume area is unfolded by pollen flushed away by the river in the spring season. That matches very well with the phone footage (see Fig. 12). Such data motivate AUV sampling for calibration, improved resolution, and 3-D characterization.

A. Experiment Setup

1) *Discretize the Grid*: Computational constraints and practical matters lead to a $25 \times 25 \times 5$ grid discretization within the $1 \text{ km} \times 1 \text{ km} \times 2 \text{ m}$ box region overlapping the river plume area as shown in Fig. 6 (red dots). We concentrate our effort on the near-surface regions (depth smaller than 2.5 m) because ocean model data and observations made during an initial AUV transect (see Fig. 14) show that the freshwater river plume tends to float close to the surface regions [1].

2) *Building the Prior*: To form a prior, we use SINMOD data as a core building block. First, we allocate mean values to each 3-D grid node, extracted from averages over many SINMOD runs. Second, we calibrate these mean values in a regression model using AUV data from a preliminary transect survey. A linear regression model $y_{u_k} = \beta_0 + \beta_1 y_{u_k}^{\text{SINMOD}}$ is fitted, where u_k indicate locations




Fig. 12. River plume zone captured by mobile phone on 27 May 2021. The camera perspective is shown as the white fan on the left corner which indicates the area where the plume occurs.




Fig. 13. Satellite image captured on 2 June 2021, showing the visible river plume thanks to the pollen flushed away by the river.

of transect line AUV data y_{u_k} and SINMOD data $y_{u_k}^{SINMOD}$. The fitted coefficients $\hat{\beta}_0, \hat{\beta}_1$ adjust the entire field, and $\hat{\beta}_0 + \hat{\beta}_1 y_{u_k}^{SINMOD}$ provides the prior mean in the onboard model used in the AUV deployment.

The coefficients for the Matérn kernel are approximated using empirical variograms of the AUV data collected from the initial survey. They are specified to $\sigma = 2, \phi_1 = 0.011, \phi_2 = 0.94,$ and $r = 0.55$. Careful assessment of these parameters is important




Fig. 14. Salinity versus depth plot from AUV in situ measurements and from SINMOD prediction. Both SINMOD and the in situ measurements show that most salinity variation happens close to the surface.



Fig. 15. LAUV Roald is taking a shower after the heavy duty.

when it comes to sharpening the performance of the adaptive sampling algorithm such that it recognizes the boundary more agilely. However, further tweaking of these parameters are out of the scope of this work.

3) *AUV Deployment*: LAUV Roald (see Fig. 15) from the Applied Underwater Robotics Laboratory at NTNU was employed in the Nidelva missions. All the essential scripts were integrated onboard on the backseat NVIDIA Jetson TX2 CPU. For hardware and software in the loop testing and the actual deployment we relied on the framework developed by [12]. The implementation of Algorithm 1 and 2 requires robot operating systems (ROS) [33] and a software bridge to the LAUV, running

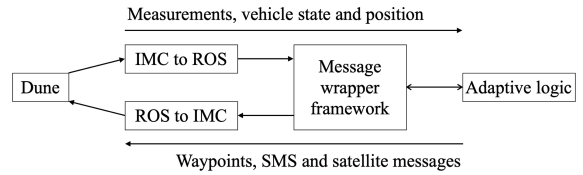


Fig. 16. Main software components in the communication between the adaptive code and the vehicle. DUNE [34] is running on the main CPU of the AUV while the IMC [35] messages are transmitted via TCP [37] to an auxiliary CPU, where ROS [33] and the adaptive code is run.

DUNE (DUNE :Unified Navigation Environment [34]) embedded and communicating over the intermodule communication (IMC) message protocol [35].

The software bridge between ROS and IMC was adapted from the Swedish Maritime Robotics Centers implementation of an ROS-IMC bridge [36]¹ to include messages going from ROS to the vehicle. In addition, a wrapper for the vehicle IMC messages was used, enabling easy interaction between the adaptive software and the vehicle. The communication bridge and framework between ROS and IMC use the same back-seat interface as [15], with IMC messages being transmitted over Transmission Control Protocol (TCP) [37] between the main CPU and the auxiliary CPU in the AUV. The adaptive code is run in the auxiliary CPU to preserve the integrity of the main CPU. For illustration, a flowchart containing the main software components is presented in Fig. 16.

B. Experiment Results and Discussion

Fig. 17 shows the posterior EPs after assimilating all the AUV measurements from the adaptive mission. When the EP is close to 1, it is classified as river water, while ocean water has probabilities close to 0. Some parts of the domain are still unexplored and have intermediate probabilities. In its adaptive sampling efforts to distinguish the water masses, the AUV travels between different layers and traverse the lateral domain. The sampling mainly takes place in the top three layers that mirror the buoyant river plume assumption, but it dips down to 2 and 2.5 m. The adaptive behavior guides the agent to be within the boundary region instead of putting too much effort on either side of the front. According to the updated field, there appears to be patches of river waters going down to 1 and 1.5 m, but most river water is near the surface.

In Fig. 18, we compare prior and posterior EPs for the top two layers. Clearly, the AUV reveals a bigger plume region than what is predicted by the SINMOD prior model. At 1.0 m there appears to be water mass separation. This kind of separation is likely very heterogeneous in space and time, and the displayed results only show predicted conditions at the day of the mission.

¹[Online]. Available: https://github.com/smarc-project/imc_ros_bridge




Fig. 17. Excursion probability for the posterior field. It describes how similar the water mass is to the river water. Values near 1 (blue) represents river water, while 0 (white) represents ocean water.




Fig. 18. Excursion probability comparison for the prior field (left) and the posterior field (right) at 0.5 m depth (top) and 1.0 m depth (bottom). The AUV trajectory is shown as the black line in the right column.

VII. CONCLUSION

The main contribution of this work is to apply Gaussian random field models for 3-D north-east-depth domains in the context of adaptive sampling with real-time computation and maneuverability routines on a robotic vehicle. The adaptive sampling routine presented here is tailored to frontal systems, and it relies on reduction of the EIBV. We conducted a simulation study comparing the suggested approach with more standard approaches. Results demonstrate the capability of the adaptive myopic 3-D sampling in a field deployment. The AUV managed to distinguish the different water masses in a river plume in a Norwegian fjord-river system.

River plumes are influenced by many factors such as winds, waves, and tides, and we could likely model statistical correlations more sensibly by using a nonstationary Gaussian random field prior [38]. Our method uses ocean model data to build a reasonable prior model of the salinity field in 3-D. However, when this type of information is lacking, the prior belief can also be constructed based on other data, possibly satellite imagery or buoy information. As AUV data are rather sparse, there is likely much to gain by using spatially covering physical modeling data and satellite data, as this allows a better initial model for sampling.

The time variation will play an important role if the AUV deployment lasts longer. This is naturally the case when the frontal region gets bigger and the distance traveled by the AUV increases. In long-term deployments it will also be important to capture such temporal effects [39]. The current myopic philosophy works well for a small river plume. As the plume gets bigger, or one has interest in capturing subregional plumes, there is likely some gain by using strategies that anticipate many stages [30], [31] or in using ocean physics for the 3-D navigation [40]. Other opportunities stem from using adaptive sampling in a cooperative fleet as discussed in [41].


ACKNOWLEDGMENT

The authors would like to thank AURLab NTNU for the support, collaboration, and easy access to testing equipment. The authors would also like to thank K. A. Skarpnes for his help during all the field-trials in 2021 and SINTEF Ocean for supplying SINMOD data.

REFERENCES

- [1] A. R. Horner-Devine, R. D. Hetland, and D. G. MacDonald, "Mixing and transport in coastal river plumes," *Annu. Rev. Fluid Mechanics*, vol. 47, no. 1, pp. 569–594, 2015.
- [2] S. Constantin, D. Doxaran, and S. Constantinescu, "Estimation of water turbidity and analysis of its spatio-temporal variability in the Danube River plume (Black Sea) using modis satellite data," *Continental Shelf Res.*, vol. 112, pp. 14–30, 2016.
- [3] A. A. Osadchiv and P. O. Zavialov, "Lagrangian model of a surface-advected river plume," *Continental Shelf Res.*, vol. 58, pp. 96–106, 2013.
- [4] S. Zheng, W. Guan, S. Cai, X. Wei, and D. Huang, "A model study of the effects of river discharges and interannual variation of winds on the plume front in winter in Pearl River Estuary," *Continental Shelf Res.*, vol. 73, pp. 31–40, 2014.
- [5] F. M. Falcieri, A. Benetazzo, M. Scavo, A. Russo, and S. Carniel, "Po river plume pattern variability investigated from model data," *Continental Shelf Res.*, vol. 87, pp. 84–95, 2014.
- [6] R. Mendes et al., "Observation of a turbid plume using modis imagery: The case of Douro estuary (Portugal)," *Remote Sens. Environ.*, vol. 154, pp. 127–138, 2014.
- [7] G. S. Saldias, J. L. Largier, R. Mendes, I. Pérez-Santos, C. A. Vargas, and M. Sobarzo, "Satellite-measured interannual variability of turbid river plumes off Central-Southern Chile: Spatial patterns and the influence of climate variability," *Prog. Oceanogr.*, vol. 146, pp. 212–222, 2016.
- [8] E. Park and E. M. Latrubesse, "Modeling suspended sediment distribution patterns of the Amazon River using Modis data," *Remote Sens. Environ.*, vol. 147, pp. 232–242, 2014.
- [9] J. Hwang, N. Bose, and S. Fan, "AUV adaptive sampling methods: A review," *Appl. Sci.*, vol. 9, no. 15, 2019, Art. no. 3145.
- [10] E. Fiorelli, N. E. Leonard, P. Bhatta, D. A. Paley, R. Bachmayer, and D. M. Fratantoni, "Multi-AUV control and adaptive sampling in Monterey Bay," *IEEE J. Ocean. Eng.*, vol. 31, no. 4, pp. 935–948, Oct. 2006.
- [11] T. O. Fossum et al., "Toward adaptive robotic sampling of phytoplankton in the coastal ocean," *Sci. Robot.*, vol. 4, no. 27, 2019, Art. no. aav3041.
- [12] T. Mo-Bjørkelund, T. O. Fossum, P. Norgren, and M. Ludvigsen, "Hexagonal grid graph as a basis for adaptive sampling of ocean gradients using AUVs," in *Proc. IEEE Glob. Oceans 2020: Singapore–U.S. Gulf Coast*, 2020, pp. 1–5.
- [13] P. Rogowski, E. Terrill, and J. Chen, "Observations of the frontal region of a buoyant river plume using an autonomous underwater vehicle," *J. Geophysical Res.: Oceans*, vol. 119, no. 11, pp. 7549–7567, 2014.
- [14] Y. Zhang, J. G. Bellingham, J. P. Ryan, B. Kieft, and M. J. Stanway, "Autonomous four-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle," *J. Field Robot.*, vol. 33, no. 1, pp. 67–81, 2016.
- [15] J. Pinto, R. Mendes, J. C. B. da Silva, J. M. Dias, and J. B. de Sousa, "Multiple autonomous vehicles applied to plume detection and tracking," in *Proc. IEEE OCEANS Conf. - MTS/IEEE Kobe Techno-Oceans*, 2018, pp. 1–6.
- [16] T. O. Fossum, C. Travalletti, J. Eidsvik, D. Ginsbourger, and K. Rajan, "Learning excursion sets of vector-valued Gaussian random fields for autonomous ocean sampling," *Ann. Appl. Statist.*, vol. 15, no. 2, pp. 597–618, 2021.
- [17] R. Cui, Y. Li, and W. Yan, "Mutual information-based multi-AUV path planning for scalar field sampling using multidimensional RRT*," *IEEE Trans. Syst., Man, Cybern. Syst.*, vol. 46, no. 7, pp. 993–1004, Jul. 2016.
- [18] P. Stankiewicz, Y. T. Tan, and M. Kobilarov, "Adaptive sampling with an autonomous underwater vehicle in static marine environments," *J. Field Robot.*, vol. 38, no. 4, pp. 572–597, 2021.
- [19] A. Sousa et al., "LAUV: The man-portable autonomous underwater vehicle," *IFAC Proc. Volumes*, vol. 45, no. 5, pp. 2–274, 2012.
- [20] D. Slagstad and T. A. McClimans, "Modeling the ecosystem dynamics of the Barents Sea including the marginal ice zone: I physical and chemical oceanography," *J. Mar. Syst.*, vol. 58, no. 12, pp. 1–18, 2005.
- [21] P. F. Lermusiaux, "Uncertainty estimation and prediction for interdisciplinary ocean dynamics," *J. Comput. Phys.*, vol. 217, no. 1, pp. 176–199, 2006.
- [22] M. Lin and C. Yang, "Ocean observation technologies: A review," *Chin. J. Mech. Eng.*, vol. 33, no. 1, pp. 1–18, 2020.
- [23] S. Martin, *An Introduction to Ocean Remote Sensing*. Cambridge, U.K.: Cambridge Univ. Press, 2014.
- [24] F. A. Al-Wassai and N. V. Kalyankar, "Major limitations of satellite images," 2013, *arXiv:1307.2434*.
- [25] S. Kemna, O. Kroemer, and G. S. Sukhatme, "Pilot surveys for adaptive informative sampling," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2018, pp. 6417–6424.
- [26] F. S. Longman, L. Mihaylova, and L. Yang, "A Gaussian process regression approach for fusion of remote sensing images for oil spill segmentation," in *Proc. 21st Int. Conf. Inf. Fusion*, 2018, pp. 62–69.
- [27] N. Cressie and C. K. Winkle, *Statistics for Spatio-Temporal Data*. Hoboken, NJ, USA: Wiley, 2015.
- [28] C. Chevalier, J. Bect, D. Ginsbourger, E. Vazquez, V. Picheny, and Y. Richet, "Fast parallel Kriging-based stepwise uncertainty reduction with application to the identification of an excursion set," *Technometrics*, vol. 56, no. 4, pp. 455–465, 2014.
- [29] D. Silver and J. Veness, "Monte-Carlo planning in large POMDPs," in *Proc. 23rd Int. Conf. Neural Inf. Process. Syst.*, 2010, pp. 2164–2174.
- [30] C. Xiong, H. Zhou, D. Lu, Z. Zeng, L. Lian, and C. Yu, "Rapidly-exploring adaptive sampling tree*: A sample-based path-planning algorithm for unmanned marine vehicles information gathering in variable ocean environments," *Sensors*, vol. 20, no. 9, 2020, Art. no. 2515.

- [31] M. Bresciani et al., "Path planning for underwater information gathering based on genetic algorithms and data stochastic models," *J. Mar. Sci. Eng.*, vol. 9, no. 11, 2021, Art. no. 1183.
- [32] S. Bhat, "Hydrobatatics: Efficient and agile underwater robots," Ph.D. dissertation, Dept. Vehicle Maritime Eng., KTH Royal Inst. Technol., Stockholm, Sweden, 2020.
- [33] M. Quigley, "ROS: An open-source robot operating system," in *Proc. ICRA Workshop Open Source Softw.*, 2009, vol. 3, no. 3.2, p. 5.
- [34] J. Pinto, P. S. Dias, R. Martins, J. Fortuna, E. Marques, and J. Sousa, "The LSTS toolchain for networked vehicle systems," in *Proc. IEEE/MTS OCEANS Conf.*, Bergen, Norway, 2013, pp. 1–9.
- [35] R. Martins et al., "IMC: A communication protocol for networked vehicles and sensors," in *Proc. IEEE Oceans 2009 - Europe*, 2009.
- [36] S. Bhat et al., "A cyber-physical system for hydrobat AUVs: System integration and field demonstration," in *Proc. IEEE/OES Auton. Underwater Veh. Symp.*, 2020, pp. 1–8.
- [37] V. Cerf and R. Kahn, "A protocol for packet network intercommunication," *IEEE Trans. Commun.*, vol. 22, no. 5, pp. 637–648, May 1974.
- [38] G.-A. Fuglstad, F. Lindgren, D. Simpson, and H. Rue, "Exploring a new class of non-stationary spatial Gaussian random fields with varying local anisotropy," *Statistica Sinica*, vol. 25, pp. 115–133, 2015.
- [39] K. H. Foss, G. E. Berget, and J. Eidsvik, "Using an autonomous underwater vehicle with onboard stochastic advection-diffusion models to map excursion sets of environmental variables," *Environmetrics*, vol. 33, 2022, Art. no. e27102.
- [40] C. S. Kulkarni and P. F. Lermusiaux, "Three-dimensional time-optimal path planning in the ocean," *Ocean Model.*, vol. 152, 2020, Art. no. 101644.
- [41] M. J. Kuhlman, D. Jones, D. A. Sofge, G. A. Hollinger, and S. K. Gupta, "Collaborating underwater vehicles conducting large-scale geospatial tasks," *IEEE J. Ocean. Eng.*, vol. 46, no. 3, pp. 785–807, Jul. 2021.



Yaolin Ge received the dual M.Sc. degrees in maritime engineering in 2020 from the KTH Royal Institute of Technology, Stockholm, Sweden, and Norwegian University of Science and Technology (NTNU), Trondheim, Norway, where he is currently working toward the Ph.D. degree in statistics with the Department of Mathematical Sciences.



Jo Eidsvik received the M.Sc. degree in applied mathematics from the University of Oslo, Oslo, Norway, in 1997, and the Ph.D. degree in statistics from the Norwegian University of Science and Technology (NTNU), Trondheim, Norway, in 2003.

He is currently a Professor of Statistics with NTNU. He has industry experience with the Norwegian Defence Research Establishment, Kjeller, Norway and with Equinor, Stavanger, Norway. His research interests include computational methods and sampling design approaches related to spatio-temporal statistics. He works mainly with applications from the earth sciences.



Tore Mo-Bjørkelund received the M.Sc. degree in marine cybernetics in 2017 from the Norwegian University of Science and Technology (NTNU), Trondheim, Norway, where he is currently working toward the Ph.D. degree in autonomous underwater systems.

Paper II

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

Yaolin Ge, Jo Eidsvik, André Julius Hovd Olaisen

Submitted to Autonomous Robots

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

Yaolin Ge¹, Jo Eidsvik¹, André Julius Hovd Olaisen¹

¹Department of Mathematical Sciences, Norwegian University of Science
and Technology, Trondheim, Norway.

Contributing authors: yaolin.ge@ntnu.no; jo.eidsvik@ntnu.no;
andre.j.h.olaisen@ntnu.no;

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.




Fig. 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

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 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.

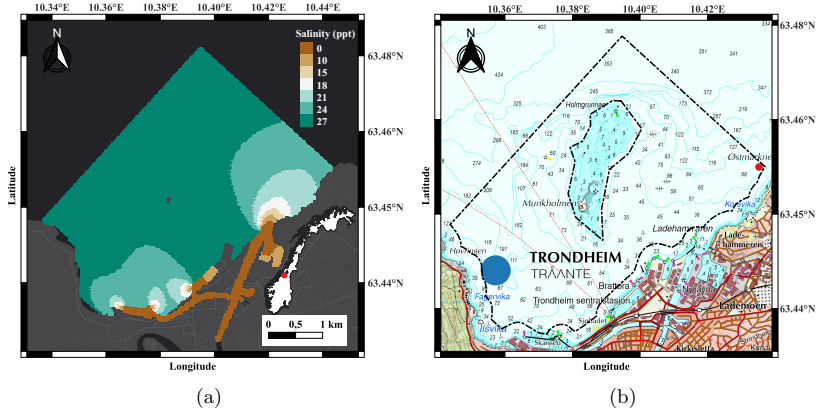


Fig. 2: (a) Numerical ocean model data for surface salinity on May 11th, 2022. (b) The selected operational area, the outer polygon draws the border for the operation whereas the inner polygon shows an unsafe region. The small red dot shows where to start, and the big blue dot shows where to end.

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).

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 σ^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 ϕ and with Euclidean distance $h(i, i')$ between sites \mathbf{u}_i and $\mathbf{u}_{i'}$, $i, i' = 1, \dots, n$. We specify parameters σ and ϕ 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 ρ . 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 ρ 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

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

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.




Fig. 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.

4.1 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.

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

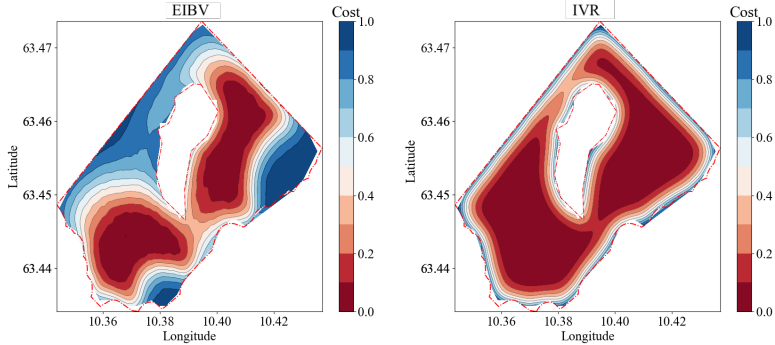


Fig. 4: Informative cost fields at the initial stage. The left display shows the EIBV cost field, which assigns low cost to areas where the river plume front might exist at this initial time. The right display shows the IVR cost field, which assigns low cost almost everywhere as it has not yet sampled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

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 ζ 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 Φ_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 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 ζ because locations close too 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

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**
 $\mathbf{f} = \mathbf{0}^{n \times 1}, \mathbf{f}(i) = 1$
 $\mathbf{R} = \mathbf{S}_{t|t-1} \mathbf{f} (\mathbf{f}^T \mathbf{S}_{t|t-1} \mathbf{f} + r^2)^{-1} \mathbf{f}^T \mathbf{S}_{t|t-1}$
 $\mathbf{IVR}(i) = \sum_{i'=1}^n \text{diag}(\mathbf{R})$
 $\mathbf{EIBV}(i) = \sum_{i'=1}^n \Phi_2 \left(\begin{bmatrix} \zeta \\ -\zeta \end{bmatrix}; \begin{bmatrix} m_{t|t-1}(i') \\ -m_{t|t-1}(i') \end{bmatrix}, \mathbf{W}(i', i') \right),$
 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 $\mathbf{C}_{\text{EIBV}}, \mathbf{C}_{\text{IVR}}$

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

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

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}_{EIBV} and \mathbf{C}_{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,

RRTStar and Cost valley illustration




Fig. 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.

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

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

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 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.

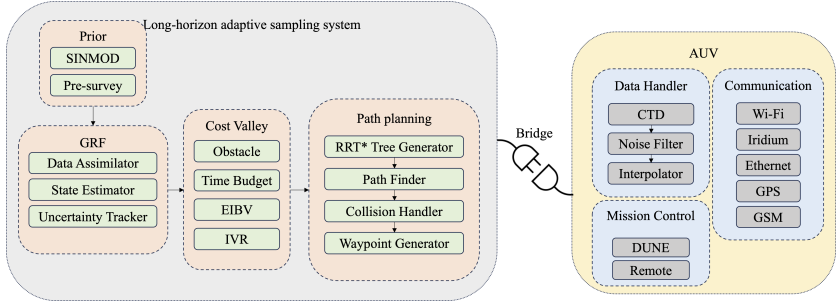


Fig. 6: System architecture of our proposed adaptive sampling system using RRT* for path planning and a cost valley for guidance.

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: $w_1 = 1, w_2 = 0$
- 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




Fig. 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.

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.

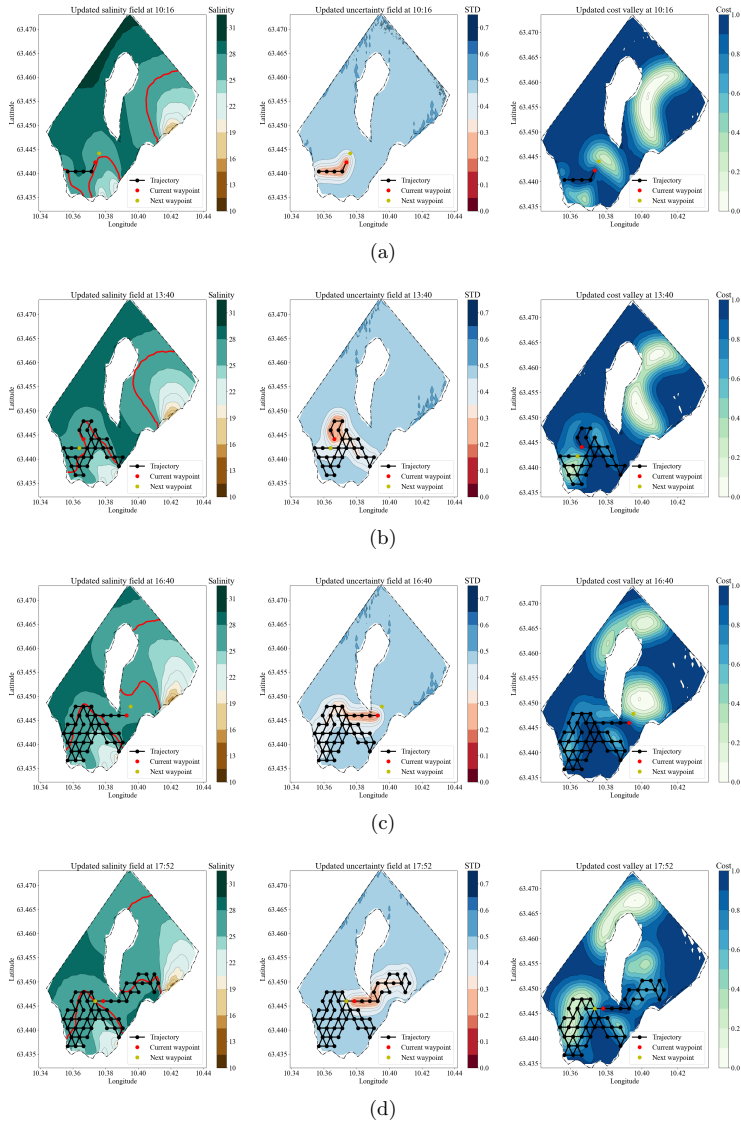


Fig. 8: The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using the myopic strategy over a series of time iterations. In this particular example, most of the effort was focused on the southwest side before the agent moved eastward.

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 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.

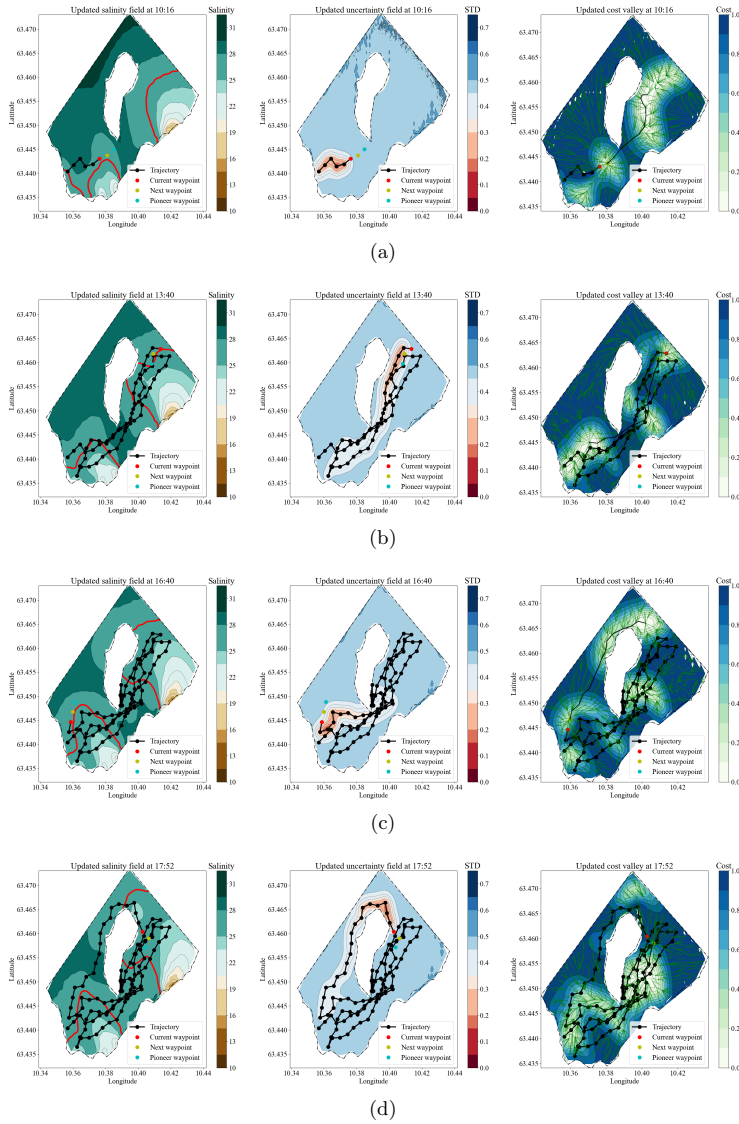


Fig. 9: The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using our proposed long-horizon RRT* strategy over a series of time iterations. The RRT* trees are displayed in the cost valley column as well to better illustrate the planning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

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, rapidly dispersing across the field. Consequently, distinguishing between the two planners becomes challenging in this context. They both effectively



(a)



(b)

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

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

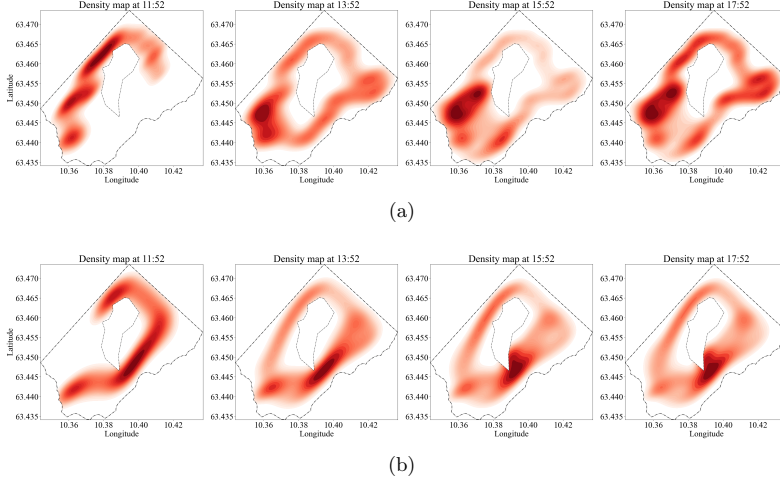


Fig. 11: The traffic flow density map for the myopic agent (a) and the long-horizon RRT* agent (b) using the IBV dominant cost valley.

obstacle are quite balanced: it is about as likely to take a route south of the obstacle as it is to go north.

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

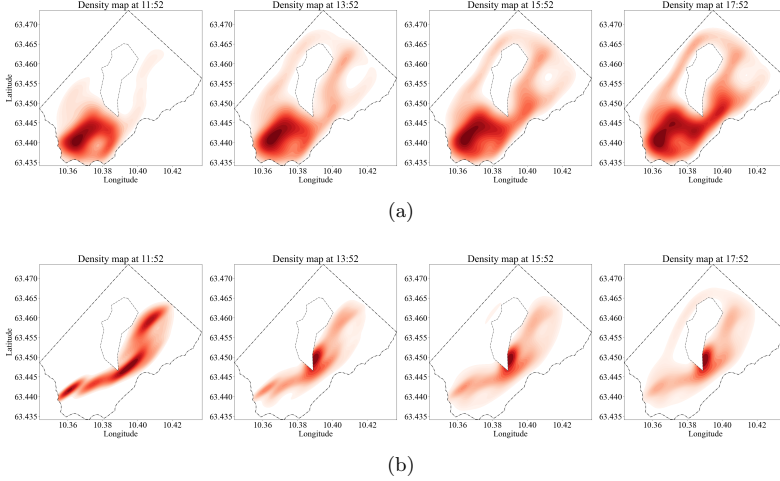


Fig. 12: The traffic flow density map for the myopic agent (a) and the long-horizon RRT* agent (b) using the equally weighted cost valley.

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 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

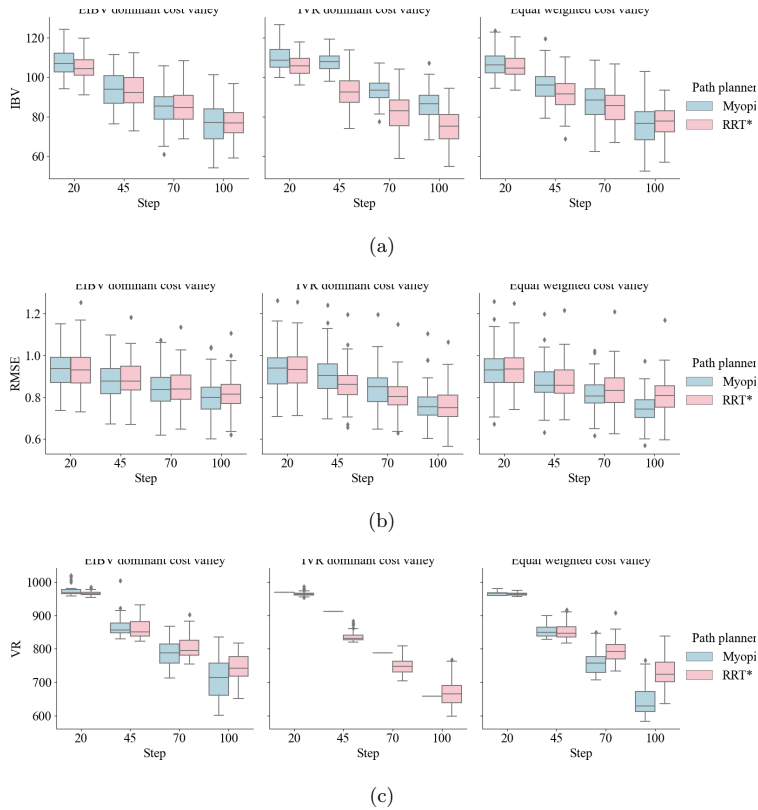


Fig. 13: The IBV (top), RMSE (middle) and VR (bottom) metrics from the 100 replicate simulation study. The solid line is the mean while the light-colored band indicates 90 % uncertainty over the 100 replicates.

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.

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




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

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.




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

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




Fig. 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.

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 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




Fig. 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.




Fig. 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).

(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.




Fig. 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).

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 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.

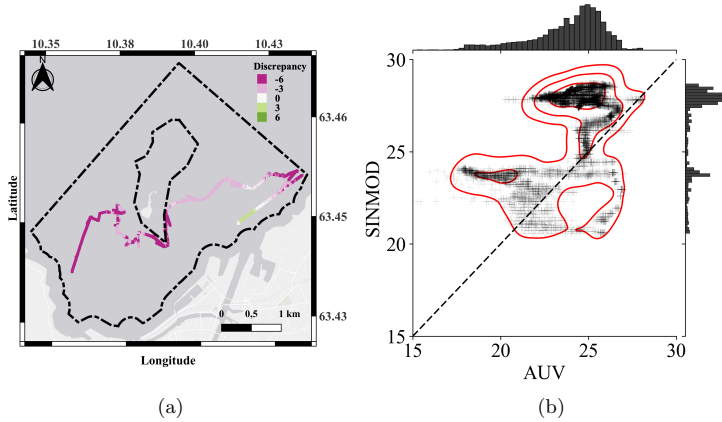


Fig. 20: (a) The difference between the AUV measurements and the SINMOD prediction, referred to as the residual. The in-situ data reveal a considerable river plume in the western part, characterized by low salinity values, which contradicts the SINMOD prediction. (b) The cross plot between the SINMOD 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.

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.

Declarations

Ethical statements

This study, centered on efficient ocean exploration by AUVs, adheres to ethical guidelines, focusing solely on software aspects without involving human or animal subjects. Field experiments, including AUV deployments in the Trondheim fjord, were environmentally considerate, ensuring no adverse impact on the marine ecosystem. The research strictly involved oceanographic data and algorithmic simulations, upholding data privacy and scientific integrity.

Author Contributions

Ge was responsible for the conceptualization and methodology of the study and played a key role in drafting the original manuscript. Eidsvik contributed to data interpretation, significantly revised the manuscript, oversaw the project administration, and provided critical resources. Olaisen assisted in drafting and

critically revising the manuscript. Both Ge and Olaisen took part in the field experiment to collect in-situ data.

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.

Funding information

This research was funded by the Norwegian Research Council (RCN) through the MASCOT project 305445.

Data Availability

The datasets generated and analyzed during this study are available upon request. We are committed to promoting collaborative research and transparency in scientific endeavors. Researchers interested in accessing the data for academic, research, or non-commercial purposes are encouraged to contact the corresponding author. Requests will be processed in accordance with our open-source policy and the ethos of community-driven research, ensuring that the data is shared responsibly and ethically.

References

Anyosa, S., J. Eidsvik, and O. Pizarro. 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., T. Shan, F. Chen, L. Liu, and B. Englot. 2021. Information-driven path planning. *Current Robotics Reports* 2(2): 177–188. <https://doi.org/10.1007/s43154-021-00045-6> .
- Belkin, I.M., P.C. Cornillon, and K. Sherman. 2009. Fronts in large marine ecosystems. *Progress in Oceanography* 81(1-4): 223–236 .
- Berget, G.E., T.O. Fossum, T.A. Johansen, J. Eidsvik, and K. Rajan. 2018. Adaptive sampling of ocean processes using an auv with a gaussian proxy model. *IFAC-PapersOnLine* 51(29): 238–243. <https://doi.org/https://doi.org/10.1016/j.ifacol.2018.09.509> .
- Binney, J., A. Krause, and G.S. Sukhatme. 2013. Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research* 32(8): 873–888 .
- Cerf, V. and R. Kahn. 1974. A protocol for packet network intercommunication. *IEEE Transactions on Communications* 22(5): 637–648. <https://doi.org/10.1109/TCOM.1974.1092259> .
- Cressie, N. and C.K. Wikle. 2015. *Statistics for spatio-temporal data*. John Wiley & Sons.
- Das, J., K. Rajany, S. Frolovy, F. Pyy, J. Ryany, D.A. Caronz, and G.S. Sukhatme. 2010. Towards marine bloom trajectory prediction for auv mission planning. In *2010 IEEE International Conference on Robotics and Automation*, pp. 4784–4790. IEEE.
- Eidsvik, J., T. Mukerji, and D. Bhattacharjya. 2015. *Value of information in the earth sciences: Integrating spatial modeling and decision analysis*. Cambridge University Press.
- Enevoldsen, T.T. and R. Galeazzi. 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 H. Bleuler 2006. Rough-terrain robot motion planning based on obstacle-ness. In *2006 9th International Conference on Control, Automation, Robotics and Vision*, pp. 1–6.
- Fedorov, K.N. 1986. *The physical nature and structure of oceanic fronts*, Volume 333. Springer.
- Fossum, T.O., J. Eidsvik, I. Ellingsen, M.O. Alver, G.M. Fragoso, G. Johnsen, R. Mendes, M. Ludvigsen, and K. Rajan. 2018. Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics* 35(7): 1101–1121. <https://doi.org/https://doi.org/10.1002/rob.21805>. <https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21805> .
- Fossum, T.O., G.M. Fragoso, E.J. Davies, J.E. Ullgren, R. Mendes, G. Johnsen, I. Ellingsen, J. Eidsvik, M. Ludvigsen, and K. Rajan. 2019. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics* 4(27): eaav3041. <https://doi.org/10.1126/scirobotics.aav3041>. <https://www.science.org/doi/pdf/10.1126/scirobotics.aav3041> .
- Fossum, T.O., C. Travelletti, J. Eidsvik, D. Ginsbourger, and K. Rajan. 2021. Learning excursion sets of vector-valued Gaussian random fields for autonomous ocean sampling. *The Annals of Applied Statistics* 15(2): 597 – 618. <https://doi.org/10.1214/21-AOAS1451> .
- Ge, Y., J. Eidsvik, and T. Mo-Bjørkelund. 2023. 3d adaptive auv sampling for classification of water masses. *IEEE Journal of Ocean Engineering* 48: 626–639 .
- Ge, Y., A.J.H. Olaisen, J. Eidsvik, R.P. Jain, and T.A. Johansen. 2022. Long-horizon informative path planning with obstacles and time constraints. *IFAC-PapersOnLine* 55(31): 124–129 .
- Hollinger, G.A. and G.S. Sukhatme. 2014. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research* 33(9): 1271–1287 .

- Horner-Devine, A.R., R.D. Hetland, and D.G. MacDonald. 2015. Mixing and transport in coastal river plumes. *Annual Review of Fluid Mechanics* 47: 569–594 .
- Hwang, J., N. Bose, and S. Fan. 2019. Auv adaptive sampling methods: A review. *Applied Sciences* 9(15): 3145 .
- Jaillet, L., J. Cortés, and T. Siméon. 2010. Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics* 26(4): 635–646. <https://doi.org/10.1109/TRO.2010.2049527> .
- Jakuba, M. and D.R. Yoerger 2008. Autonomous search for hydrothermal vent fields with occupancy grid maps. In *Proc. of ACRA*, Volume 8, pp. 2008. Citeseer.
- Karaman, S. and E. Frazzoli. 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., L. Liu, H.K. Heidarsson, and G.S. Sukhatme. 2018. Data-driven learning and planning for environmental sampling. *Journal of Field Robotics* 35(5): 643–661 .
- Mo-Bjørkelund, T., T.O. Fossum, P. Norgren, and M. Ludvigsen 2020. Hexagonal grid graph as a basis for adaptive sampling of ocean gradients using auvs. In *Global Oceans 2020: Singapore – U.S. Gulf Coast*, pp. 1–5.
- Okopal, G. 2019. Multi-objective autonomy for auv adaptive sampling missions. In *OCEANS 2019 MTS/IEEE SEATTLE*, pp. 1–6. IEEE.
- Pinto, J., P.S. Dias, R. Martins, J. Fortuna, E. Marques, and J. Sousa 2013. The lsts toolchain for networked vehicle systems. In *2013 MTS/IEEE OCEANS - Bergen*, pp. 1–9.
- Preston, V., G. Flaspohler, A.P. Michel, J.W. Fisher III, and N. Roy. 2022. Robotic planning under uncertainty in spatiotemporal environments in expeditionary science.
- Quigley, M. 2009. Ros: an open-source robot operating system. In *ICRA 2009*.

- Richards, J.A. and J.A. Richards. 2022. *Remote sensing digital image analysis*, Volume 5. Springer.
- Rigby, P., O. Pizarro, and S.B. Williams. 2010. Toward adaptive benthic habitat mapping using gaussian process classification. *Journal of Field Robotics* 27(6): 741–758 .
- Sharples, J., J.J. Middelburg, K. Fennel, and T.D. Jickells. 2017. What proportion of riverine nutrients reaches the open ocean? *Global Biogeochemical Cycles* 31(1): 39–58 .
- Slagstad, D. and T.A. McClimans. 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., Y.T. Tan, and M. Kobilarov. 2021. Adaptive sampling with an autonomous underwater vehicle in static marine environments. *Journal of Field Robotics* 38(4): 572–597 .
- Suh, J., J. Gong, and S. Oh. 2017. Fast sampling-based cost-aware path planning with nonmyopic extensions using cross entropy. *IEEE Transactions on Robotics* 33(6): 1313–1326. <https://doi.org/10.1109/TRO.2017.2738664> .
- Xiao, C. and J. Wachs. 2022. Nonmyopic informative path planning based on global kriging variance minimization. *IEEE Robotics and Automation Letters* 7(2): 1768–1775. <https://doi.org/10.1109/LRA.2022.3141458> .
- Zacchini, L., M. Franchi, and A. Ridolfi. 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., J.G. Bellingham, J.P. Ryan, B. Kieft, and M.J. Stanway. 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., B. Kieft, B.W. Hobson, J.P. Ryan, B. Barone, C.M. Preston, B. Roman, B.Y. Raanan, R. Marin III, T.C. O'Reilly, C.A. Rueda, D. Pargett, K.M. Yamahara, S. Poulos, A. Romano, G. Foreman, H. Ramm, S.T. Wilson, E.F. DeLong, D.M. Karl, J.M. Birch, J.G. Bellingham, and C.A. Scholin. 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. <https://doi.org/10.1109/JOE.2019.2920217> .
- Zuhlke, M., N. Fomferra, C. Brockmann, M. Peters, L. Veci, J. Malik, and P. Regner 2015. Snap (sentinel application platform) and the esa sentinel 3 toolbox. In *Sentinel-3 for Science Workshop*, Volume 734, pp. 21.

Paper III

Efficient 3D real-time adaptive AUV sampling of a river plume front

Martin Outzen Berild, Yaolin Ge, Jo Eidsvik, Geir-Arne Fuglstad, and Ingrid Ellingsen

Submitted to Frontiers in Marine Science

Efficient 3D real-time adaptive AUV sampling of a river plume front

Martin Outzen Berild¹, Yaolin Ge¹, Jo Eidsvik¹,
Geir-Arne Fuglstad¹, Ingrid Helene Ellingsen²

¹Department of Mathematical Sciences, Norwegian University of Science
and Technology, Trondheim, Norway.

²Fisheries and New Biomarine Industry, SINTEF Ocean AS, Trondheim,
Norway.

Contributing authors: martin.o.berild@ntnu.no; yaolin.ge@ntnu.no;
jo.eidsvik@ntnu.no; geir-arne.fuglstad@ntnu.no; ingrid.ellingsen@sintef.no;

The coastal environment faces multiple challenges due to climate change and human activities. Sustainable marine resources management necessitates knowledge, and the development of efficient ocean sampling approaches is increasingly important for understanding the ocean processes. Currents, winds, and freshwater runoff make ocean variables such as salinity very heterogeneous, and standard statistical models can be unreasonable for describing such complex environments. We employ a class of Gaussian Markov random fields that learns complex spatial dependencies and variability from numerical ocean model data. The suggested model further benefits from fast computations using sparse matrices, and this facilitates real-time model updating and adaptive sampling routines on an autonomous underwater vehicle. To justify our approach we compare its performance in a simulation experiment with a similar

approach using a more standard statistical model. We show that our suggested modeling framework outperforms the current state of the art for modeling such spatial fields. Then the approach is tested in a field experiment using two autonomous underwater vehicles for characterizing the three dimensional fresh-/saltwater front in the sea outside Trondheim, Norway. One vehicle is running an adaptive path planning algorithm while the other runs a pre-programmed path. The objective of adaptive sampling is to reduce the variance of the excursion set to classify freshwater and more saline fjord water masses. Results show that the adaptive strategy conducts effective sampling of the frontal region of the river plume.

Keywords: Adaptive sampling, Ocean modeling, Autonomous underwater vehicle, Gaussian random field, Stochastic partial differential equations, Surrogate model

1 Introduction

Human activities and pollution are heavily impacting the world's oceans (Halpern et al., 2008). Anthropogenic climate change and local intrusion from industries can lead to fundamentally altered ocean ecosystems, challenging species distributions, loss in biodiversity, incidence of disease, and more (Hoegh-Guldberg and Bruno, 2010; Doney et al., 2012). The changes in ecosystem structure further influence important services such as carbon sequestration, oxygen production and nutrient food chains. In order to achieve a more sustainable utilization of marine resources and services we need to enhance our insight. Developing smart technologies for efficient monitoring of the ocean can provide information that enables us to identify adverse effects and guide development of countermeasures, and it can hence be vital in saving or maintaining local ecosystems. Commonly used ocean observation technologies are buoys, drifters, satellites, unmanned surface vehicles, Argo floats, underwater gliders, cabled seafloor observatories, autonomous underwater vehicles (AUVs), hadal landers, or some coupled system of these technologies (see e.g. (Lin and Yang, 2020) for a overview). Ocean

monitoring systems are advancing from simple and static single sensors systems to dynamic and multi sensor systems that can cover a large spectrum of temporal and spatial scales. With the drive in artificial intelligence and robotic systems, there is also a development towards intelligent sampling systems where observations of various kinds are gathered and processed where and when it is considered valuable.

With the improved affordability and functionality of AUVs, the research literature has seen many advances lately; Zhang et al. (2012, 2013) used deterministic algorithms to map coastal temperature upwellings; Das et al. (2015) demonstrated AUV mission planning for informative plankton sampling; Fossum et al. (2018) monitored large temperature gradients by adaptively choosing surveys paths that substantially reduce the uncertainty in the statistical temperature model; Fossum et al. (2019) conducted a 3D AUV survey for chlorophyll-a mapping; Mo-Bjørkelund et al. (2020a) employed hexagonal grids for equilateral survey paths to adaptively explore large temperature gradients; Foss et al. (2022) used a 2D spatio-temporal model onboard an AUV to supervise mining waste seafill; Fonseca et al. (2023) compared satellite imagery and adaptive AUV sampling results for predicting algal blooms. These examples from recent research activity have advanced the field of ocean monitoring with AUVs by going from planar (sea-surface) fields to volumetric fields, in the combination of various data sources, or by presenting a novel algorithm for adaptive exploration.

Considering the vastness of our ocean, it is extremely difficult to obtain sufficient data to cover the full range of scale and resolution desired. Instead, one must rely on a combination of different data sources and sophisticated modeling tools. To fill in the gaps effectively one can further proactively plan targeted and high precision sampling campaigns that will improve predictions and support decision-making. At its core, these tasks relate to statistical methods that can combine various data sources for prediction and for evaluating data sampling designs to optimize further data-gathering efforts.

In this work, we combine the fields of oceanography, statistics, and robotics to effectively monitor freshwater frontal regions of river outlets in three-dimensions (north, east, depth) using AUVs. Specifically, we conduct sampling in the Nidelva river running into the fjord outside Trondheim, Norway (see Figure 1). The freshwater coming from




Fig. 1: Map of the operational area in the fjord outside Trondheim, Norway. The location of Trondheim is indicated by the red circle on the map of Scandinavia in the top left corner. We have 3D numerical ocean model data at a high lateral and depth resolution in these coastal waters. The blue square just north of the Nidelva river outlet indicates the boundaries of the autonomous underwater vehicle mission in a map view. The operation domain extends from the sea surface to 5 meter depth.

the river is mixing slowly with the more salinity fjord waters which can cause a sharp gradient between the different water masses.

At our availability, we have output from a complex numerical ocean model Slagstad and McClimans (2005), henceforth referred to as SINMOD. Along with many other physical oceanography variables, SINMOD outputs salinity at every grid node in a dense spatial (3D) and temporal grid. Even though this model carries much physical insight, the salinity output can be systematically biased, and we will calibrate and update the salinity by deploying an AUV. In this way, the SINMOD data is used to form a prior model for the salinity trends and variations at the time of the AUV

deployment. We fit a Gaussian process prior as a surrogate model to the numerical ocean model SINMOD. This surrogate model has the advantage that it can be updated onboard the AUV, and it can hence assimilate in-situ data efficiently. Moreover, this surrogate model enables fast evaluation of various AUV sampling designs in real-time while it is maneuvering in the water. Fossum et al. (2021) used similar methods to fit a surrogate model from SINMOD, but only in 2D space. Ge et al. (2023) used a 3D Gaussian surrogate prior model, but only for a small-size grid and assuming a much simpler spatial dependency structure.

This paper brings together many elements, and the novelty lies in a more realistic description of spatial correlations with a complex model learned from SINMOD data (Berild and Fuglstad, 2023). The approach is made computationally feasible through modern techniques using a 3D Gaussian random fields with a Markov property, and this enables adaptive AUV sampling based on the new surrogate model in a large-size 3D waypoint graph used during AUV deployments. Additionally, we

- fit the more realistic 3D statistical model to 3D numerical ocean data, and develop a fast algorithm for updating this model onboard an AUV during field deployment,
- develop methods for adaptive path-planning in the context of 3D space with the more realistic model onboard the AUV,
- show through a simulation study, based on SINMOD, that the more realistic statistical model allows an AUV to sample and map the ocean domain better than with a standard statistical model,
- run two AUVs simultaneously in the ocean and show that the combination of an intelligent adaptive survey design and the more realistic model outperforms a standard pre-scripted AUV sampling plan.

In Section 2, we describe the numerical ocean model and its statistical surrogate model. In Section 3, we present the data assimilation part and our approach for adaptive AUV sampling designs. In Section 4, we study properties of the suggested methods in

a simulation study. In Section 5, we show results of deployments with one adaptive AUV mission and one pre-programmed mission. In Section 6, we conclude and point to future work.

2 Prior model for salinity

Consider a three-dimensional ocean domain $\mathcal{D} \subseteq \mathbb{R}^3$, where $x(\mathbf{s})$ represents the salinity field at a specific location $\mathbf{s} = (\text{latitude}, \text{longitude}, \text{depth})^T \in \mathcal{D}$. The salinity in this ocean domain exhibits both spatial and temporal variations. However, we focus on short-term AUV deployments and simplify our analysis by excluding temporal effects.

2.1 Numerical Ocean model

An approximation of the salinity field is achieved using the complex numerical ocean model SINMOD, developed by SINTEF ocean (Slagstad and McClimans, 2005). SINMOD is a three-dimensional model based on the primitive equations, solved using finite difference methods on a regular grid with horizontal cell sizes of $20\text{km} \times 20\text{km}$, which are nested in several steps down to $32\text{m} \times 32\text{m}$ for the bay outside Trondheim. The model employs varying vertical resolution, allowing for higher resolution near the dynamic surface and more uniform resolution in deeper waters. Atmospheric forces (obtained from forecasts available at <https://www.met.no>), freshwater outflows (data from HBV model (Beldring et al., 2003) provided by the Norwegian Water Resources and Energy Directorate (NVE)), and tides (<https://www.tpxo.net/>) drive the model. SINMOD offers numerical simulations of multiple ocean variables, including temperature and currents as well as salinity. It is a multi-purpose tool that has been used for instance in the prediction of Arctic ocean primary production by leveraging physical-biologically coupling (Slagstad et al., 2015; Vernet et al., 2021), in quantifying the effects of the aquaculture structures for large-scale cages by specifying and incorporating drag parameters in SINMOD (Broch et al., 2020), and coupled with the particle dispersion

of waste from fish farming (Broch et al., 2017), oil production (Nepstad et al., 2022) or mine tailings (Berget et al., 2018; Nepstad et al., 2020; Berget et al., 2023). For a more comprehensive explanation of the SINMOD methodology, readers are directed to Slagstad and McClimans (2005).

In the current paper, we are only using the salinity outputs from SINMOD. Figure 2 shows an example of SINMOD salinity data and an excursion set (salinity ≤ 25.4 g/kg) separating water masses into freshwater/saltwater in the fjord outside Trondheim. We notice that the river plume has lower salinity than the surrounding brackish water. There are very low salinity levels (about 5 – 10 g/kg) in the river outlet, while the salinity increases further out in the fjord (about 31 g/kg). Salinity is measured in grams salinity per kilogram water (g/kg) which is dimensionless and equal to ‰ and sometimes referred to as the practical salinity unit (PSU).

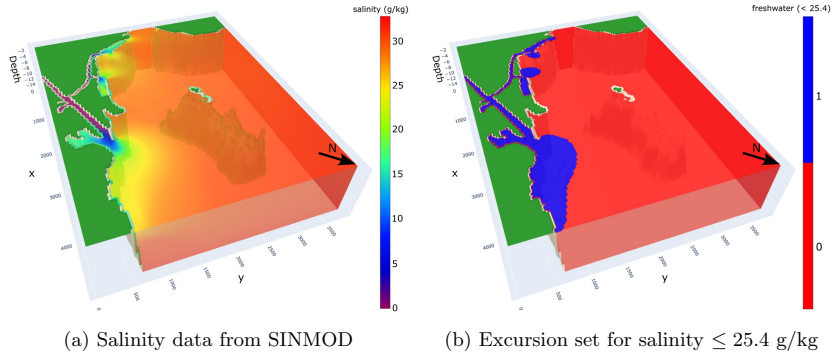


Fig. 2: Simulation from the numerical ocean model SINMOD for September 8th, 2022.

2.2 Surrogate model with spatially varying anisotropy

In-situ salinity observations made with an AUV are assumed to be more accurate than the forecast provided by SINMOD. However, an AUV measurement only characterizes the salinity at the specific location where the measurement was taken, whereas a model

like SINMOD or similar is required to extrapolate variables in space and time. For onboard computing, SINMOD is however too computationally intensive, and it is challenging to assimilate AUV observations in real-time with a full-fledged numerical ocean model. Instead, a surrogate model can be trained from the numerical model. It forms an approximate representation of the underlying physical model and is highly applicable for different tasks that require fast updating. Statistical models with spatial effects have shown very suitable for such a task (Gramacy, 2020), and we employ a particular statistical surrogate model for the numerical ocean model salinity data here.

We use the spatial statistical model presented in Berild and Fuglstad (2023), where the 3D salinity field is modeled as a Gaussian Markov random field (GMRF) that allows sparse matrix computations and realistic modeling via spatial variability in the directional dependencies and the variance components. This model is an extension of Lindgren et al. (2011) and Fuglstad et al. (2015).

Assuming that the 3D discretization of the domain $\mathcal{D} \subseteq \mathbb{R}^3$ consists of $n_1 \times n_2 \times n_3$ grid cells, the salinity field, $x(\mathbf{s})$, is represented by a vector of concatenated field values of size $n = n_1 n_2 n_3$. In the application, we have $n_1 = 50$, $n_2 = 45$, $n_3 = 6$ with 32×32 m² lateral resolution and 1 m depth resolution. The vector \mathbf{x} of salinity values is modeled by a Gaussian distribution, i.e.

$$\mathbf{x} \sim \mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad \boldsymbol{\Sigma} = \mathbf{Q}^{-1}. \quad (1)$$

Here, the \sim symbol means 'distributed according to', and $\mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ refers to the n -variate Gaussian (or normal) distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, where its inverse, namely the precision matrix, is denoted \mathbf{Q} .

There is much flexibility in choosing the mean vector and covariance matrix in Equation (1), and the Gaussian distribution can hence form quite realistic surrogate models. The mean vector $\boldsymbol{\mu}$ of the salinity field captures the spatial trends of the field,

which in our case entails fresher water near the river gradually getting more saline going out in the fjord. To form a realistic covariance structure, the idea of Berild and Fuglstad (2023) is to form a random process for $\mathbf{u} = \mathbf{x} - \boldsymbol{\mu}$ via differential operators and Gaussian noise forming a stochastic partial differential equation (SPDE) as

$$(\kappa^2(\mathbf{s}) - \nabla \cdot \mathbf{H}(\mathbf{s})\nabla)u(\mathbf{s}) = \mathcal{W}(\mathbf{s}). \quad (2)$$

Here, \mathbf{s} is a location in the domain of interest $\mathcal{D} \subseteq \mathbb{R}^3$, $u(\mathbf{s})$ is the spatially varying deviation from the trend, and $\mathcal{W}(\mathbf{s})$ is a Gaussian white noise process (with zero mean and statistically independent values), while $\kappa(\mathbf{s}) = \kappa(\mathbf{s}; \boldsymbol{\theta}) > 0$ and $\mathbf{H}(\mathbf{s}) = \mathbf{H}(\mathbf{s}; \boldsymbol{\theta}) > 0$ is differentiable are model components controlled by parameters $\boldsymbol{\theta}$ that regulate the variability and dependency within the process.

Equation (2) is solved locally for the zero-mean random field $u(\mathbf{s})$ using numerical integration and differentiation on a discretization of the domain of interest \mathcal{D} . The solution is $\mathbf{u} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{Q}^{-1})$, where the precision matrix $\mathbf{Q} = \mathbf{Q}(\boldsymbol{\theta})$ inherits the sparsity of the differential operators in Equation (2) and it describes the Markov structure in the GMRF model. This structure is very important for our purposes because it enables fast matrix factorization and matrix-vector computations. Hence, the GMRF formulation means that we can update the model onboard the AUV. It is also used in the sampling design evaluations. Without this sparsity, the Gaussian surrogate model could not scale up the magnitude of the ocean mass in 3D (Berild and Fuglstad, 2023).

A detailed description of the model is provided in the Supplemental material.

2.3 Parameter estimation for salinity field

In order to estimate the parameters and components of the statistical GMRF model for salinity, we utilize numerical ocean model data from SINMOD as the training dataset. This data is denoted as $y(\mathbf{s}_i, t_j)$, where $\mathbf{s}_i \in \mathcal{D}$ represents the location of cell

$i \in [1, \dots, n]$ at time t_j for SINMOD realization $j = 1, \dots, T$. The surrogate data model is then

$$y(\mathbf{s}_i) \sim \mathcal{N}(x(\mathbf{s}_i), \sigma_{\mathbb{S}}^2), \quad i = 1, \dots, n, \quad (3)$$

where $\sigma_{\mathbb{S}}^2$ is an unstructured noise variance of the SINMOD dataset.

We estimate a location-dependent mean $\mu(\mathbf{s}_i)$ of the GMRF using the empirical average across all replicates t_j as:

$$\hat{\mu}(\mathbf{s}_i) = \frac{1}{T} \sum_{j=1}^T y(\mathbf{s}_i, t_j), \quad \forall \mathbf{s}_i \in \mathcal{D}. \quad (4)$$

We compute an estimate $\hat{\boldsymbol{\theta}}$ of the covariance parameters of the GMRF by maximizing the likelihood function $\mathcal{L}(\boldsymbol{\theta})$, given residual data from an autoregressive model fitted to the SINMOD data (See Supplemental material and (Berild and Fuglstad, 2023)). The covariance parameters for the models used in Sections 4 and 5 are fit on 144 timesteps SINMOD data of the whole field. Optimizing the likelihood of this rather sophisticated covariance model is not straightforward, but it gets less difficult with more data and at this also improves the accuracy of estimates. Berild and Fuglstad (2023) suggest that at least 10 timesteps of the whole field should be used to find reasonable parameters values for such a flexible model.

Figure 3 shows the prior mean (Equation (4)), the prior variance of the n-variate Gaussian distribution \boldsymbol{x} and the corresponding spatial correlation of the marked location. The mean salinity clearly increases going north in the fjord, away from the river outlet. The salinity variance is larger near the river. For the correlation, we notice non-circular contours indicative of anisotropy. Here, the correlation appears to be stronger in the directions where salinity is expected to be similar to that of the reference location.

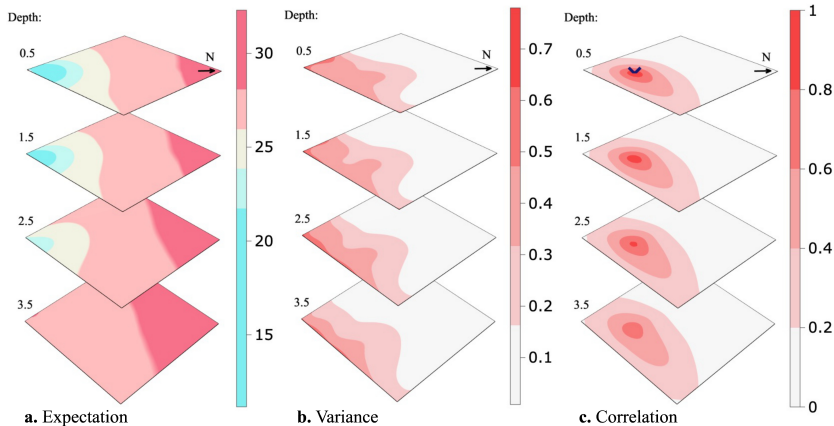


Fig. 3: Prior expectation (a), the variance of the process model (b), and the spatial correlation of the location highlighted (c). The N-arrow is the cardinal north.

3 Adaptive AUV sampling

We now delve into our approach for adaptive AUV exploration. One part of this involves continuous updates of the GMRF surrogate model through onboard data assimilation of the in-situ AUV salinity data. Another part is the strategic planning of the next AUV sampling locations.

3.1 Conditioning to AUV data

Assume that the AUV gathers in-situ data at m locations or design points $\mathbf{d} = \{\mathbf{d}_1, \dots, \mathbf{d}_m\}$, where $\mathbf{d}_j \in \mathcal{D}$. In practice these locations form an AUV design (a trajectory). Data $y(\mathbf{d}_j)$, $j = 1, \dots, m$, are noisy measurements of the salinity $x(\mathbf{d}_j)$ at the location \mathbf{d}_j where they are made. We organize the data in a length- m vector \mathbf{y} , and we allocate these observations to the correct grid locations by using a size $m \times n$ selection matrix \mathbf{A} . This matrix has a single 1 entry in each row, and otherwise only 0 entries. With this structure, it selects the m indices in the length- n vector \mathbf{x} of

discretized salinity field variables in Equation (1). The measurement model is then

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \boldsymbol{\epsilon}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}_m(0, \sigma_{\text{auv}}^2 \mathbf{I}_m). \quad (5)$$

Here, the variance σ_{auv}^2 of the independent additive noise terms aggregates the AUV positioning error and measurement noise. This variance parameter is specified from existing AUV data.

The conditional model for salinity \mathbf{x} , given measurements \mathbf{y} , is Gaussian distributed with updated precision matrix

$$\mathbf{Q}_C = \mathbf{Q} + \mathbf{A}^T \mathbf{A} / \sigma_{\text{auv}}^2, \quad (6)$$

and conditional mean

$$\boldsymbol{\mu}_C = \boldsymbol{\mu} + \mathbf{Q}_C^{-1} \mathbf{A}^T (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}) / \sigma_{\text{auv}}^2. \quad (7)$$

With the sparse precision matrices, the updating in Equation (6) and (7) can be computed very fast.

Given a series of observations collected with the AUV along a straight line from the river plume and straight north we calculate the conditional precision matrix and mean using Equation(6)-(7) of the model estimated in Section 2.3. Using the precision we calculate the inverse diagonal (the conditional variance of the field), and from this the correlation about a location in space. We demonstrate the effect of data conditioning using a visualization of the conditional expectation, conditional variance, and conditional correlation given a series of updates are shown in Figures 4, 5, and 6. Figure 4 indicates that the river water is going further north than anticipated in the prior mean. In Figure 5 we see that the variance is reduced where the AUV has visited,




Fig. 4: Conditional expectation given AUV measurements along a fixed transect path at 0.5 meter depth.

and as a consequence the correlation range shown in Figure 6 gets lower. Dense data sampling tends to reduce the spatial correlation.

3.2 Excursion sets and plume mapping criterion

One goal of the AUV sampling is to improve the characterization of the plume front defined in our case as the zone separating fresh river waters and more saline fjord waters. Following Fossum et al. (2021) and Ge et al. (2023), we use the uncertainty in the random set of excursions below a salinity threshold ℓ to distinguish river and fjord water. The excursion set is defined by

$$ES = \{\mathbf{s} \in \mathcal{D} : x(\mathbf{s}) < \ell\}. \quad (8)$$

The associated excursion probability (EP) and the Bernoulli variance (BV) is

$$EP(\mathbf{s}) = P(x(\mathbf{s}) < \ell), \quad BV(\mathbf{s}) = EP(\mathbf{s})[1 - EP(\mathbf{s})] \quad \mathbf{s} \in \mathcal{D}. \quad (9)$$




Fig. 5: Conditional variance in the process model given AUV measurements along a fixed transect path at 0.5 meter depth.

The BV is near 0 at locations where the EP is near 0 or 1, while it is at its maximum value 0.25 at locations which have EP equal to 0.5.

When AUV salinity data \mathbf{y} are available, we get a conditional GMRF, and conditional EPs and BVs. Effective AUV sampling designs get salinity data that can pull these EPs closer to 0 or 1 and in doing so one reduces the uncertainty of the river plume front. Conditional on salinity data $\mathbf{y} = \mathbf{y}_d$ according to design $\mathbf{d} \in \mathcal{D}$. The conditional EPs and BVs are

$$P(x(\mathbf{s}) < \ell | \mathbf{y}_d), \quad P(x(\mathbf{s}) < \ell | \mathbf{y}_d)[1 - P(x(\mathbf{s}) < \ell | \mathbf{y}_d)]. \quad (10)$$

Design plans must be made before the data \mathbf{y}_d is revealed, and we take the expectation over the data when calculating the most effective design. Focusing on improved spatial mapping of the river plume front, it is natural to integrate the objective criterion over all locations in the domain. The expected integrated Bernoulli variance (EIBV) of a




Fig. 6: Conditional correlation of the marked point given AUV measurements along a fixed transect path at 0.5 meter depth.

design \mathbf{d} is then defined by

$$\text{EIBV}_{\mathbf{d}}(\boldsymbol{\mu}, \mathbf{Q}) = \int_{\mathcal{D}} E_{\mathbf{y}_{\mathbf{d}}} \{P(x(\mathbf{s}) < \ell|\mathbf{y}_{\mathbf{d}})[1 - P(x(\mathbf{s}) < \ell|\mathbf{y}_{\mathbf{d}})]\} d\mathbf{s}. \quad (11)$$

For the GMRF surrogate model specified by mean $\boldsymbol{\mu}$ and precision \mathbf{Q} , the EIBV for a design \mathbf{d} has a closed form involving sums of bivariate cumulative distribution functions Φ_2 for the Gaussian distribution. In this expression, the design is here involved via 1-entries structure of the selection matrix $\mathbf{A} = \mathbf{A}_{\mathbf{d}}$. The closed-form solution facilitates very fast computations of multiple sampling designs. The complete derivations of the closed forms are in the Supplemental material. See also Fossum et al. (2021) and Ge et al. (2023). In our approach with the sparse GMRF model, we use Monte Carlo sampling from the conditional model to approximate the variance reduction components that are required in the EIBV (see Supplemental material).

3.3 Adaptive AUV sampling algorithm

The AUV cannot navigate to all possible design locations. Rather, its continued path is constrained by the current location and the possible maneuvers it can perform. We let $\mathcal{P} \subset \mathcal{D}$ denote the possible designs the AUV can choose from, defined by directions (straight, left, right, up, down) from the current AUV location. The chosen design is the one that minimizes the EIBV in Equation (11). This means that

$$\mathbf{d}^* = \operatorname{argmax}_{\mathbf{d} \in \mathcal{P}} \text{EIBV}_{\mathbf{d}}(\boldsymbol{\mu}, \mathbf{Q}). \quad (12)$$


During the AUV operation this kind of design choice is done at many time points, and with an updated model that is conditional on all the data gathered up to this point. In this way we utilize the benefits of robotic intelligence to navigate the uncertain ocean plume zone.

We outline a myopic adaptive sampling algorithm in the 3D domain. This is a sequential selection of waypoints or grid nodes where the AUV sample data and update the model. The myopic approach represents a heuristic optimization strategy for the AUV operation that does not anticipate potential data or navigation choices beyond the current time. It makes the optimal choice based on the expected values at the current time alone.


Figure 7a shows the idea of adaptive sampling in a sketch with a cycle of tasks where one leads to the next. Here, the AUV senses the salinity, updates its onboard model, and plans where to navigate to, and then it continues on the next cycle.

Hence, at the planning stage, the computer onboard the AUV solves Equation (12) to navigate in promising 3D directions Figure 7b. To compensate for the time it takes to do the computation, and to make the system near real-time, asynchronous parallel computing is applied to compensate for the excessive computing time onboard.

Algorithm 1 shows the main steps of this adaptive AUV sampling approach. In



(a) Caption for the first subfigure



(b) Caption for the second subfigure

Fig. 7: Illustration of adaptive sampling mechanism. Visualization of the adaptive sampling design (a). The AUV evaluates the potential high-value locations to determine the next visiting waypoint (b). Red colors represent more interesting next waypoint.

Algorithm 1 Myopic EIBV minimizing sampling with a GMRF surrogate model.

Initialization: Prior model $\mu_{C,0} = \mu$, $\mathbf{Q}_{C,0} = \mathbf{Q}$. Set start location \mathbf{d}_0 . Set $t = 1$.

while True **do**

Plan:

 Choose design that reduces EIBV the most

$$\mathbf{d}_t = \arg \min_{\mathbf{d} \in \mathcal{P}_t} \text{EIBV}_{\mathbf{d}}(\mu_{C,t-1}, \mathbf{Q}_{C,t-1})$$

 Form selection matrix $\mathbf{A}_t = \mathbf{A}(\mathbf{d}_t)$.

Act/sense:

 Move according to design \mathbf{d}_t collecting measurements \mathbf{y}_t

Model Update:

 With the collected measurements, update the GMRF

$$\begin{aligned} \mathbf{Q}_{C,t} &= \mathbf{Q}_{C,t-1} + \mathbf{A}_t^T \mathbf{A}_t / \sigma_{\text{auv}}^2 \\ \mu_{C,t} &= \mu_{C,t-1} + \mathbf{Q}_{C,t}^{-1} \mathbf{A}_t^T (\mathbf{y}_t - \mathbf{A}_t \mu_{C,t-1}) / \sigma_{\text{auv}}^2 \end{aligned}$$

 Set $t = t + 1$.

end while

this algorithm we use t to indicate subsequent stages of AUV sampling. At stage t , the updated mean in the onboard surrogate model is denoted $\mu_{C,t}$ and the updated precision is $\mathbf{Q}_{C,t}$. The selection matrix $\mathbf{A}_t = \mathbf{A}(\mathbf{d}_t)$ is formed based on the most promising design \mathbf{d}_t at each stage. This design \mathbf{d}_t is chosen among several possible designs $\mathcal{P}_t \subset \mathcal{D}$ that vary depending on where the AUV is at the current stage and the operational navigation opportunities it has according to the grid. In our implementation,

the AUV can continue from its current location to go straight ahead, or turn left, right, up, down. It cannot return back to its previous grid location (Figure 7b). There are natural exceptions at the grid boundary.

4 Simulation study

In this section, we conduct a simulated experiment to evaluate the performance of our approach for monitoring the three-dimensional freshwater plume of the Nidelva river in Trondheim, Norway. The operational area is outlined in Figure 1. Specifically, we will compare the effectiveness of the suggested complex GRF model and a more standard model. The complex model is discretized with a resolution of 32m x 32m square cells in the horizontal plane and the standard model with a hexagonal grid with a lateral neighbor distance of 120m. Both models have 1m depth increments ranging from 0.5 m to 5.5 m, resulting in a total of $n = 50 \times 45 \times 6$ spatial location for the complex model and 1098 for the standard model. This is inline with the capabilities of the AUVs' onboard computer.

Initially, both models are estimated on the SINMOD data within the operational area in order to form a prior field. The standard model is specified using a standard variogram analysis, resulting in a Matérn covariance with lateral correlation range of 550 m, vertical range of 2 m, a prior marginal variance of 1, and a nugget effect of 0.4 (see Section 2.4 of Cressie (1993) for a description of this spatial data analysis method). The parameters of the complex model are estimated through the approach described in Section 2.3, and detailed in the Supplemental material and Berild and Fuglstad (2023). Both models use the empirical average across all timesteps (replicates) of the SINMOD data, Equation (4), as its prior expectation.

In order to obtain performance statistics, we ran $L = 100$ simulated field experiments where the AUV is equipped with either one of the models estimated above and tasked with monitoring the salinity field according to Algorithm 1. The AUV is in this

simulation environment exploring a SINMOD dataset from the 9th of November 2022 with an assumed additional Gaussian noise term with standard deviation 0.12. This noise represents positional error and measurement error in a real experimental setting. Moreover, the AUV is set to travel at 1 m/s and each simulated field experiment is run for $T = 25$ sequential steps of Algorithm 1, i.e. visiting 25 spatial locations, where the starting location is kept the same for each run.

Within the l th simulated experiment and after visiting t th location, the following three metrics are calculated: integrated Bernoulli variance (IBV), root mean squared error (RMSE), and classification error (CE). Let x_l be the ground truth (SINMOD data) in the l th experiment. Then, we calculate the metrics as

$$\text{IBV}_{l,t} = \sum_{i=1}^n \text{EP}_{l,t}(\mathbf{s}_i)[1 - \text{EP}_{l,t}(\mathbf{s}_i)], \quad (13)$$

$$\text{RMSE}_{l,t} = \sqrt{\frac{1}{n} \sum_{i=1}^n [x_l(\mathbf{s}_i) - \mu_{C,l,t}(\mathbf{s}_i)]^2}, \quad (14)$$

$$\text{CE}_{l,t} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(\text{E}_l(\mathbf{s}_i) \neq \hat{\text{E}}_{l,t}(\mathbf{s}_i)), \quad \text{E}_l(\mathbf{s}_i) = \mathbb{I}(x_l(\mathbf{s}_i) \leq \ell) \quad \hat{\text{E}}_{l,t}(\mathbf{s}_i) = \mathbb{I}(\mu_{C,l,t}(\mathbf{s}_i) \leq \ell), \quad (15)$$

where \mathbb{I} is the indicator function, $t \in [0, T]$ where $T = 25$ indicates the sequential step, and $l = 1, \dots, L$ with $L = 100$ replicate field experiment. Summary statistics of these metrics from the L replicated experiments are shown in Figure 8. The solid lines are the average across all L replicates at time t for each metric, e.g.

$$\widehat{\text{IBV}}_t = \frac{1}{L} \sum_{l=1}^L \text{IBV}_{l,t}, \quad (16)$$




Fig. 8: Variation in integrated Bernoulli variance (**a**), root mean square error (**b**), and classification error (**c**) over the 100 replicate runs with the standard model (blue) and the complex model (orange). The solid lines show the averages and the vertical errorbars show the empirical standard deviations of the respective metrics.

and similarly the errorbars show the empirical standard deviation at time t across all L replicates as

$$\mathbb{SD}(\text{IBV}_t) = \sqrt{\frac{1}{L-1} \sum_{l=1}^L (\text{IBV}_{l,t} - \widehat{\text{IBV}}_t)^2}. \quad (17)$$

Each display has one of the metrics on the second axis and time stages on the first axis. For the IBV and CE criteria, the percentages reduction compared to the starting value are shown since the models are constructed differently and therefore will also differ prior to the mission as can be seen in the middle RMSE display.

The IBV reduction (Figure 8a) indicates the ability of the AUV to capture the river plume boundary. A lower IBV means that the AUV is better at sampling the frontal salinity region separating river and fjord water masses. In this spatial example the IBV has a tendency of going down, even though it could increase at some stages (because data pull probabilities closer to 0.5). The complex GMRF model clearly achieves lower IBV than the simpler model. After some stages, the curve for the complex GMRF model declines rapidly, indicating that the AUV is efficient at exploring the boundary. This means that incorporating a more realistic covariance structure helps the AUV choose the best designs and it tends to move in the right direction.

The RMSE plot (Figure 8b) reflects the similarity between the ground truth and the updated field. The ground truth is here the same as what the AUV is sampling, i.e.

the SINMOD dataset from the 9th of November 2022, but without the added noise term. A lower RMSE means that the AUV is gathering data that helps in predicting the salinity field. Again, the complex model is performing much better than the simpler one. For CE (Figure 8c), a lower value means that the updated model is good at classifying the excursion set associated with the ground truth. The complex model has CE results that are declining faster than the simpler model. The complex model performs better than the standard model due to its versatile capability and flexibility. However, we do realize that training such models often requires expert knowledge and it can be a laborious process to fine-tune the parameters for such a complex model.

In all displays of Figure 8, we observe larger metric variability for the complex GMRF model. The underlying reason for this is the Monte Carlo variance in the EIBV calculation for the GMRF model (see Section 3.2). With the relatively small sampling size, the Monte Carlo error is still not negligible and, influenced by this estimate, the directional sampling decision made by the AUV exhibit more small-scale variability than that of the standard model which has a closed form variance expression. Over many replicates, the variability in metrics then gets larger for the GMRF model, especially for the IBV which relates directly to the AUV sampling decision criterion.

5 Results of Nidelva Mission

The field experiment was executed in the Nidelva river plume outside Trondheim, Norway, on the 8th of September, 2022. The duration of this field deployment spanned 1.5 hours. Figure 1 shows the operational area.

5.1 Experimental set-up

For this experiment two AUVs are deployed. This is intended to not only increase the amount of data collected, but also enable us to compare the performance of our embedded system under similar conditions. One of the AUVs was programmed with

the adaptive sampling algorithm, while the other was running with a pre-programmed path plan onboard.

LAUV (Light Autonomous Underwater Vehicle) Harald and LAUV Roald (Figure 9) from the Applied Underwater Robotics Laboratory at NTNU were employed for this mission. LAUV Roald was programmed to carry out the adaptive experiment and LAUV Harald was programmed to conduct the pre-designed plan. To measure the salinity in the water LAUVs Harald and Roald uses CTD sensors, or conductivity, temperature, and depth sensor. Harald uses a SeaBird SBE 49 FastCAT and Roald a AML OEM SV Xchange. Despite being from different manufacturers, the specifications from the suppliers indicate that they should have the same level of precision and accuracy. All the essential scripts were integrated onboard on the backseat NVIDIA Jetson TX2 CPU. For hardware and software in the loop testing and the actual deployment, we relied on the framework developed by Mo-Bjørkelund et al. (2020b). The onboard implementation of Algorithm requires Robot Operating Systems (ROS) Quigley (2009) and a software bridge to the LAUV, running DUNE (DUNE: Unified Navigation Environment Pinto et al. (2013)) embedded and communicating over the Inter Module Communication (IMC) message protocol (LSTS, 2022).

The software bridge between ROS and IMC was adapted from the Swedish Maritime Robotics Centers implementation of a ROS-IMC bridge Bhat et al. (2020)(https://github.com/smarc-project/imc_ros_bridge) to include messages going from ROS to the vehicle. In addition, a wrapper for the vehicle IMC messages was used, facilitating interaction between the adaptive software and the vehicle. The communication bridge and framework between ROS and IMC use the same back-seat interface as Pinto et al. (2018), with IMC messages being transmitted over Transmission Control Protocol (TCP) (Cerf and Kahn, 1974) between the main CPU and the auxiliary CPU in the AUV. The adaptive code was running on the auxiliary CPU in order to preserve the



Fig. 9: LAUV Roald is on expedition with an adaptive sampling algorithm onboard. The AUV is about 2 m long and runs at about 1 m/s. It is doing a 3D sampling mission at depths ranging from 0.5 to 5.5 m.

integrity of the main CPU. For illustration, a flowchart containing the main software components is presented in Figure 10.

Before conducting the principal deployment, we gained understanding of the sea conditions via a preliminary survey. We first launched the pre-survey adaptive mission with a reasonable threshold based on our belief field and then updated the threshold to be 25.4 g/kg after observing the updated salinity field from the pre-survey run.

5.2 Field operation

The AUVs started moving from their starting location around 12:50am. We received the "Mission Complete" message from the AUVs around 14:15pm which marks the end of the operation.

In Figure 11, the results of the AUV conducting adaptive sampling are displayed. Here, we plot the AUV path (black) and the updated posterior mean salinity field over time steps. The AUV began near the river mouth and gradually moved towards the frontal region, occasionally diving to the deeper layers. In total the AUV traveled approx

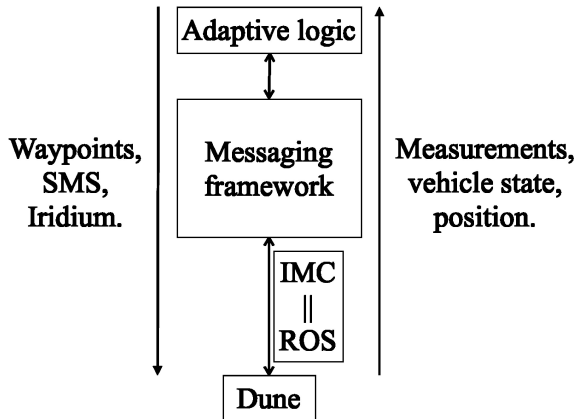


Fig. 10: The diagram of the software component in the adaptive sampling system. The main CPU of the AUV is running DUNE (Pinto et al., 2013), while IMC (LSTS, 2022) messages are sent through TCP (Cerf and Kahn, 1974) to a secondary CPU, where the adaptive code and ROS (Quigley, 2009) are executed.

9.734 km, with a coverage of 6.9% of the field at 0.5m, 6.3% at 1.5m, 1.6% at 2.5m, and 0.1% at 3.5m. Both the AUVs was set to travel at 1.5 m/s, but the speed varies widely because of conditions in the ocean. During the mission the plume expanded due to the tidal effect, so the AUV attempted to follow the front more closely. Interestingly, the AUV did not dive deeper than 2.5m. This can be attributed to the fact that the water becomes more homogeneous and saline when it is too deep, and the river plume tends to stay close to the upper layers. Also this can be an effect of the model learning from observation closer to the surface. Note that the path in Figure 11 appear somewhat disjointed because the viewed measurements are assigned to the nearest grid cell to their actual location and the grid cell center of each measurements are shown here.

In Figure 12, the salinity prediction results of the AUV's pre-planned sampling are displayed. The path was designed to maximize the sampling coverage and consequently reduce the variance of the field. The AUV was programmed to move along the path with a consistent YoYo pattern. This pattern involves the AUV moving between 0.5m and 5.5m repeatedly. The pre-programmed path approach ensures a more systematic




Fig. 11: Salinity prediction during the adaptive sampling mission September 8th 2022. The AUV (black) began close to the river mouth and gradually moved towards the frontal region and dived to deeper layers occasionally.




Fig. 12: Salinity prediction during the fixed path mission September 8th 2022. The AUV (black) aims to cover the spatial domain.

and exhaustive coverage of the volume, providing a broader perspective but lacks the pinpoint accuracy on such a large and the rapidly changing ocean volume. The path the AUV traveled along was approx 9.346km with a coverage of 7.9% at 0.5m, 8.3% at 1.5m, 9.1% at 2.5m, 8.6% at 3.5m, 7.9% at 4.5m, and 7.2% at 5.5m.

Given the unpredictable nature of the location of the freshwater front, it is virtually impossible to pre-plan precise sampling paths. The shifts and movements of the plume demand a real-time responsive approach like adaptive sampling. This is also evident

when comparing Figure 11 and 12. On the other hand, if a broad overview of the ocean volume is the goal, then a pre-planned design likely is useful to ensure a systematic coverage of the region, leaving minimal gaps in the data collection. Also, note that diving 1 m is significantly more time efficient than moving 32 m in the horizontal plane for the AUV, this can be viewed by the coverage in each layer by two missions. The fixed path mission has good coverage within each layer whilst the adaptive mission mostly considered the top two layers. Because of this it could be interesting in future work to consider adaptive sampling in only the horizontal plane and to always move in a YoYo pattern.

5.3 SINMOD and AUV data comparison

Even though salinity is only one state variable in SINMOD, it is useful and interesting to compare the AUV salinity measurements with the predictions made by SINMOD, as it will give information on the overall performance of the hydrodynamic model. Using all the data collecting by both AUVs, we compare the location-specific observations made by the AUVs with the associated SINMOD predictions for this day.

As illustrated in Figure 13, there is a clear inclination of SINMOD to overpredict salinity values. This trend is evident as a majority of the AUV measurements are situated below the zero error line (dotted line). In shallow water regions, both SINMOD and the AUV measurements exhibited high salinity variability which is reasonable considering the freshwater influx from the river and local disturbances. For deeper waters the salinity is higher for both models and more concentrated, with a bias about 3.5 g/kg between the SINMOD predictions and AUV measurements. The highest measured salinity value by the AUV was 28.0 g/kg, while the highest value from SINMOD was 31.5 g/kg. Further confirming that the numerical model overestimates salinity both for the water in the river plume and in the brackish layer in the fjord.




Fig. 13: Comparison of salinity values estimated from the numerical ocean model (x-axis) and salinity measurements collected with the AUV on the 8th of September 2022 (y-axis) where the color shows the corresponding depths. The dotted line shows the zero error line.

While this discrepancy between SINMOD and the actual salinity field is evident, this will not impact the learning of the covariance structure, which captures the spatial correlation and variability within the data, and is independent of any systematic bias. That said, the observed overestimation in SINMOD does set a prior expectation in our model that is skewed slightly high and will initial impact the adaptive sampling algorithm.

6 Conclusions

We have presented an approach for effective 3D (north, east, depth) sampling of the salinity in a river plume front with a realistic flexible spatial covariance model running onboard the AUV in real-time. Results of a deployment in the Trondheim fjord show that prior inputs from the SINMOD numerical ocean model are effectively calibrated with the in-situ AUV measurements. In a mission focusing on mapping the frontal

region, the AUV adapts naturally to the updated situation and zig-zags near the plume front to improve its spatial characteristics. Moreover, it is evident that the adaptive approach holds a distinct advantage over the pre-planned method when it comes to accurately monitoring dynamic zones like the river plume front.

Although our focus in this study is centered on separating ocean masses of low (freshwater plume) and high (brackish water) salinity concentrations, we believe that this approach transfers to other applications in physical or biological oceanography, such as polar melting water, high chlorophyll concentrations, or pollution detection.

The statistical model, while promising, does necessitate further refinement to fully realize its potential in this context. Firstly, refining its parameterization slightly to simplify the likelihood surface can potentially improve the optimization process significantly. Furthermore, estimating the covariance structure to innovations constructed from the SINMOD data, as described in Berild and Fuglstad (2023), is not guaranteed to be very accurate in removing the temporal effect in the data. Thus, making it challenging to ascertain if the final structure is only capturing the spatial effect.

The prior models used in this work included 3D space with no temporal variation. A natural extension is to include temporal variation in the prior, which could be done in a Gaussian framework assuming known advection and diffusion (Foss et al., 2022). But more research is required to develop realistic space-time models for frontal regions, such as that associated with river plumes, while maintaining the computational efficiency required to conduct expansive field surveying as considered in this study. Lastly, our exploration was confined to a near-sighted myopic sampling scheme. Future avenues might explore more sophisticated strategies (Bai et al., 2021), using longer sampling horizons where one can look ahead and anticipate the information gained by traversing longer distances with the AUV while also accounting for operational constraints.

Conflict of Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

MB and YG implemented the algorithm, ran the simulation study, and conducted the AUV field deployment with all the software and hardware testing as well as the adaptive and pre-planned AUV missions. JE and GAF supervised the methodological development. IE prepared SINMOD data and provided oceanographic motivation for the study. All contributed in the article writing but MB drove this process forward.

Funding

This work is funded by the Norwegian Research Council through the MASCOT project 305445.

Acknowledgments

We acknowledge the aid and support of AURLab (<https://www.ntnu.edu/aur-lab>) in lending us the AUVs and help in running multiple field mission.

Supplemental Material

This article contains Supplemental material on the statistical details of spatial modeling and sampling methodology.

Data Availability Statement

The datasets for this study can be found in the repository [Data for "Efficient 3D real-time adaptive AUV sampling of a river plume front"] [<https://doi.org/10.17605/OSF.IO/RHDKM>].

References

- Bai, S., T. Shan, F. Chen, L. Liu, and B. Englot. 2021. Information-driven path planning. *Current Robotics Reports* 2(2): 177–188. <https://doi.org/10.1007/s43154-021-00045-6> .
- Beldring, S., K. Engeland, L.A. Roald, N.R. Sælthun, and A. Voksø. 2003, June. Estimation of parameters in a distributed precipitation-runoff model for Norway. *Hydrology and Earth System Sciences* 7(3): 304–316. <https://doi.org/10.5194/hess-7-304-2003> .
- Berget, G.E., J. Eidsvik, M.O. Alver, and T.A. Johansen. 2023. Dynamic stochastic modeling for adaptive sampling of environmental variables using an auv. *Autonomous Robots* 47: 483–502 .
- Berget, G.E., T.O. Fossum, T.A. Johansen, J. Eidsvik, and K. Rajan. 2018. Adaptive sampling of ocean processes using an auv with a gaussian proxy model. *IFAC-PapersOnLine* 51(29): 238–243 .
- Berild, M.O. and G.A. Fuglstad. 2023. Spatially varying anisotropy for gaussian random fields in three-dimensional space. *Spatial Statistics* 55: 100750 .
- Bhat, S., I. Torroba, Ö. Özkahraman, N. Bore, C.I. Sprague, Y. Xie, I. Stenius, J. Severholt, C. Ljung, J. Folkesson, et al. 2020. A cyber-physical system for hydrobatic auvs: system integration and field demonstration. In *2020 IEEE/OES Autonomous Underwater Vehicles Symposium (AUV)*, pp. 1–8. IEEE.
- Broch, O.J., R.L. Daae, I.H. Ellingsen, R. Nepstad, E.Å. Bendiksen, J.L. Reed, and G. Senneset. 2017. Spatiotemporal dispersal and deposition of fish farm wastes: a

- model study from central norway. *Frontiers in Marine Science* 4: 199 .
- Broch, O.J., P. Klebert, F.A. Michelsen, and M.O. Alver. 2020. Multiscale modelling of cage effects on the transport of effluents from open aquaculture systems. *Plos one* 15(3): e0228502 .
- Cerf, V. and R. Kahn. 1974. A protocol for packet network intercommunication. *IEEE Transactions on Communications* 22(5): 637–648. <https://doi.org/10.1109/TCOM.1974.1092259> .
- Chevalier, C., J. Bect, D. Ginsbourger, E. Vazquez, V. Picheny, and Y. Richet. 2014. Fast parallel kriging-based stepwise uncertainty reduction with application to the identification of an excursion set. *Technometrics* 56(4): 455–465 .
- Cressie, N. 1993. *Statistics for spatial data*. John Wiley & Sons.
- Das, J., F. Py, J.B. Harvey, J.P. Ryan, A. Gellene, R. Graham, D.A. Caron, K. Rajan, and G.S. Sukhatme. 2015. Data-driven robotic sampling for marine ecosystem monitoring. *The International Journal of Robotics Research* 34(12): 1435–1452 .
- Doney, S.C., M. Ruckelshaus, J. Emmett Duffy, J.P. Barry, F. Chan, C.A. English, H.M. Galindo, J.M. Grebmeier, A.B. Hollowed, N. Knowlton, J. Polovina, N.N. Rabalais, W.J. Sydeman, and L.D. Talley. 2012. Climate Change Impacts on Marine Ecosystems. *Annual Review of Marine Science* 4(1): 11–37. <https://doi.org/10.1146/annurev-marine-041911-111611> .
- Fonseca, J., S. Bhat, M. Lock, I. Stenius, and K.H. Johansson. 2023. Adaptive sampling of algal blooms using autonomous underwater vehicle and satellite imagery: Experimental validation in the baltic sea. *arXiv preprint arXiv:2305.00774* .
- Foss, K.H., G.E. Berget, and J. Eidsvik. 2022. Using an autonomous underwater vehicle with onboard stochastic advection-diffusion models to map excursion sets of environmental variables. *Environmetrics* 33: e2702 .

- Fossum, T.O., J. Eidsvik, I. Ellingsen, M.O. Alver, G.M. Fragoso, G. Johnsen, R. Mendes, M. Ludvigsen, and K. Rajan. 2018. Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics* 35(7): 1101–1121. <https://doi.org/10.1002/rob.21805> .
- Fossum, T.O., G.M. Fragoso, E.J. Davies, J.E. Ullgren, R. Mendes, G. Johnsen, I. Ellingsen, J. Eidsvik, M. Ludvigsen, and K. Rajan. 2019, February. Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics* 4(27): eaav3041. <https://doi.org/10.1126/scirobotics.aav3041> .
- Fossum, T.O., C. Travelletti, J. Eidsvik, D. Ginsbourger, and K. Rajan. 2021. Learning excursion sets of vector-valued Gaussian random fields for autonomous ocean sampling. *The Annals of Applied Statistics* 15(2): 597 – 618. <https://doi.org/10.1214/21-AOAS1451> .
- Fuglstad, G.A., F. Lindgren, D. Simpson, and H. Rue. 2015. Exploring a new class of non-stationary spatial gaussian random fields with varying local anisotropy. *Statistica Sinica* 25: 115–133 .
- Fuglstad, G.A., D. Simpson, F. Lindgren, and H. Rue. 2015. Does non-stationary spatial data always require non-stationary random fields? *Spatial Statistics* 14: 505–531 .
- Ge, Y., J. Eidsvik, and T. Mo-Bjørkelund. 2023. 3d adaptive auv sampling for classification of water masses. *IEEE Journal of Ocean Engineering* 48: 626–639 .
- Gramacy, R.B. 2020. *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. CRC press.
- Halpern, B.S., S. Walbridge, K.A. Selkoe, C.V. Kappel, F. Micheli, C. D’Agrosa, J.F. Bruno, K.S. Casey, C. Ebert, H.E. Fox, R. Fujita, D. Heinemann, H.S. Lenihan, E.M.P. Madin, M.T. Perry, E.R. Selig, M. Spalding, R. Steneck, and R. Watson. 2008, February. A Global Map of Human Impact on Marine Ecosystems. *Science* 319(5865): 948–952. <https://doi.org/10.1126/science.1149345> .

- Hoegh-Guldberg, O. and J.F. Bruno. 2010, June. The Impact of Climate Change on the World's Marine Ecosystems. *Science* 328(5985): 1523–1528. <https://doi.org/10.1126/science.1189930> .
- Lin, M. and C. Yang. 2020, April. Ocean Observation Technologies: A Review. *Chinese Journal of Mechanical Engineering* 33(1): 32. <https://doi.org/10.1186/s10033-020-00449-z> .
- Lindgren, F., H. Rue, and J. Lindström. 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73(4): 423–498 .
- LSTS. 2022. Inter module communication protocol. <https://lsts.pt/docs/imc/master>. Accessed: 2022-11-01.
- Mo-Bjørkelund, T., T.O. Fossum, P. Norgren, and M. Ludvigsen 2020a, October. Hexagonal Grid Graph as a Basis for Adaptive Sampling of Ocean Gradients using AUVs. In *Global Oceans 2020: Singapore – U.S. Gulf Coast*, pp. 1–5. ISSN: 0197-7385.
- Mo-Bjørkelund, T., T.O. Fossum, P. Norgren, and M. Ludvigsen 2020b. Hexagonal grid graph as a basis for adaptive sampling of ocean gradients using auvs. In *Global Oceans 2020: Singapore – U.S. Gulf Coast*, pp. 1–5.
- Nepstad, R., M. Liste, M.O. Alver, T. Nordam, E. Davies, and T. Glette. 2020. High-resolution numerical modelling of a marine mine tailings discharge in western norway. *Regional Studies in Marine Science* 39: 101404 .
- Nepstad, R., T. Nordam, I.H. Ellingsen, L. Eisenhauer, E. Litzler, and K. Kotzakoulakis. 2022. Impact of flow field resolution on produced water transport in lagrangian and eulerian models. *Marine Pollution Bulletin* 182: 113928 .
- Pinto, J., P.S. Dias, R. Martins, J. Fortuna, E. Marques, and J. Sousa 2013. The lsts toolchain for networked vehicle systems. In *2013 MTS/IEEE OCEANS - Bergen*,

pp. 1–9.

- Pinto, J., R. Mendes, J.C.B. da Silva, J.M. Dias, and J.B. de Sousa 2018, May. Multiple autonomous vehicles applied to plume detection and tracking. In *2018 OCEANS - MTS/IEEE Kobe Techno-Oceans (OTO)*, pp. 1–6.
- Quigley, M. 2009. Ros: an open-source robot operating system.
- Slagstad, D. and T.A. McClimans. 2005, October. Modeling the ecosystem dynamics of the Barents sea including the marginal ice zone: I. Physical and chemical oceanography. *Journal of Marine Systems* 58(1): 1–18. <https://doi.org/10.1016/j.jmarsys.2005.05.005> .
- Slagstad, D., P.F. Wassmann, and I. Ellingsen. 2015. Physical constrains and productivity in the future arctic ocean. *Frontiers in Marine Science* 2: 85 .
- Vernet, M., I. Ellingsen, C. Marchese, S. Bélanger, M. Cape, D. Slagstad, and P.A. Matrai. 2021. Spatial variability in rates of net primary production (npp) and onset of the spring bloom in greenland shelf waters. *Progress in Oceanography* 198: 102655 .
- Wackernagel, H. 2003. *Multivariate geostatistics: an introduction with applications*. Springer Science & Business Media.
- Whittle, P. 1954. On Stationary Processes in the Plane. *Biometrika* 41(3/4): 434–449. <https://doi.org/10.2307/2332724> .
- Zammit-Mangion, A. and J. Rougier. 2018. A sparse linear algebra algorithm for fast computation of prediction variances with gaussian markov random fields. *Computational Statistics & Data Analysis* 123: 116–130 .
- Zhang, Y., J.G. Bellingham, J.P. Ryan, B. Kieft, and M.J. Stanway 2013, September. Two-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle. In *2013 OCEANS - San Diego*, pp. 1–4. ISSN: 0197-7385.

Zhang, Y., J.P. Ryan, J.G. Bellingham, J.B.J. Harvey, and R.S. McEwen. 2012.

Autonomous detection and sampling of water types and fronts in a coastal upwelling system by an autonomous underwater vehicle. *Limnology and Oceanography: Methods* 10(11): 934–951. <https://doi.org/10.4319/lom.2012.10.934> .

Supplementary Material

Describing GMRFs through SPDEs

Whittle (1954) showed that the solution, $u(\mathbf{s})$, of the stochastic partial differential equation (SPDE)

$$(\kappa(\mathbf{s}))^2 - \nabla \cdot \mathbf{H}(\mathbf{s})\nabla u(\mathbf{s}) = \mathcal{W}(\mathbf{s}), \quad \mathbf{s} \in \mathcal{D} \subseteq \mathbb{R}^3, \quad (18)$$

is a Gaussian Markov random field (GMRF) with a Matérn covariance function. The approach was popularized by Lindgren et al. (2011); then, extended to non-stationary and anisotropic 2D fields in Fuglstad et al. (2015,?), and recently to 3D fields by Berild and Fuglstad (2023).

In Equation (18), $\mathcal{W}(\mathbf{s})$ is Gaussian white noise, while κ is a parameter controlling both variance and range of the GMRF. The component \mathbf{H} is also used to regulate the variance and the range, but more importantly, it is controlling the anisotropy of the Laplacian, $\nabla \cdot \mathbf{H}\nabla$, and thus the anisotropy of the resulting field. With this anisotropy, the model can account for varying properties depending on directions. Also, note that $\kappa(\mathbf{s})$ and $\mathbf{H}(\mathbf{s})$ in Equation (18) are allowed to vary through space.

In the following, we will describe the parametrization of the anisotropy (Section 6), the non-stationarity (Section 6), properties of the GMRF described through the SPDE (Section 6), and how we infer parameters from data (Section 6).

Parametrizing Anisotropy

The spatially and directionally varying covariances are described by controlling the eigenvalues and eigenvectors of the matrix \mathbf{H} in the anisotropic Laplacian. Berild and Fuglstad (2023) proposed the following parsimonious and interpretable parameterization

$$\mathbf{H}(\mathbf{s}) = \gamma(\mathbf{s})\mathbf{I}_3 + \mathbf{v}(\mathbf{s})\mathbf{v}(\mathbf{s})^\top + \boldsymbol{\omega}(\mathbf{s})\boldsymbol{\omega}(\mathbf{s})^\top, \quad (19)$$

where $\gamma(\mathbf{s}) > 0$, $\mathbf{v} = (v_x, v_y, v_z)^\top \in \mathbb{R}^3$ and $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^\top \in \mathbb{R}^3$ whereby $\mathbf{v} \perp \boldsymbol{\omega}$. Thus, the eigenvalues are $\lambda_1 = \gamma$, $\lambda_2 = \gamma + \|\mathbf{v}\|^2$, and $\lambda_3 = \gamma + \|\boldsymbol{\omega}\|^2$ with eigenvectors $\mathbf{v}_1 = \mathbf{v} \times \boldsymbol{\omega}$, $\mathbf{v}_2 = \mathbf{v}$, and $\mathbf{v}_3 = \boldsymbol{\omega}$.

The vector \mathbf{v} is simply parametrized with its Cartesian components, v_x , v_y , and v_z . Further, $\boldsymbol{\omega}$ is parametrized by two scalars, ρ_1 and ρ_2 , controlling the linear combination of two orthogonal vector, $\boldsymbol{\omega}_1 = (-v_y, v_x, 0)^\top$ and $\boldsymbol{\omega}_2 = \mathbf{v} \times \boldsymbol{\omega}_1$, in the plane with \mathbf{v} as normal vector such that

$$\boldsymbol{\omega} = \rho_1 \frac{\boldsymbol{\omega}_1}{\|\boldsymbol{\omega}_1\|} + \rho_2 \frac{\boldsymbol{\omega}_2}{\|\boldsymbol{\omega}_2\|}. \quad (20)$$

Parametrizing non-stationarity

The non-stationarity is obtained by allowing the parameters κ , γ , v_x , v_y , v_z , ρ_1 and ρ_2 to vary throughout space. This is achieved by describing these parameters as spline functions:

$$g(\mathbf{s}) = (\mathbf{s})^\top \boldsymbol{\alpha}_g. \quad (21)$$

Here, $\boldsymbol{\alpha}_g \in \mathbb{R}^p$ is a vector of weights or the *new* parameters for the spline function $g(\cdot)$ or the *old* parameter, and $(\mathbf{s}) = (\mathbf{f}_1(\mathbf{s}), \dots, \mathbf{f}_p(\mathbf{s}))^\top$ is a p -dimensional vector of B-spline basis functions evaluated at location \mathbf{s} . These basis splines are constructed as a tensor product of three clamped 1D second-order B-splines in each dimension as

$$f_{ijk}(\mathbf{s}) = B_{x,i}(x) \cdot B_{y,j}(y) \cdot B_{z,k}(z), \quad \mathbf{s} = (x, y, z)^\top \in \mathcal{D}, \quad (22)$$

where $B_{x,i}$ is the i -th component in the x -direction and similarly for the other directions. Thereby, the number of parameters for each spline function $g(\cdot)$ is $p = 27$ which in total is 189 parameters or B-spline weights for all spline functions in our model. We collect them as

$$\boldsymbol{\theta} = (\boldsymbol{\alpha}_{\log(\kappa^2)}, \boldsymbol{\alpha}_{\log \gamma}, \boldsymbol{\alpha}_{v_x}, \boldsymbol{\alpha}_{v_y}, \boldsymbol{\alpha}_{v_z}, \boldsymbol{\alpha}_{\rho_1}, \boldsymbol{\alpha}_{\rho_2}).$$

General properties

As mentioned in Section 6 the parameters controlling the covariance structure in the spatial effect are learned from the residuals of an autoregressive model of order one fit to the SINMOD dataset \mathbf{y} . This gives the underlying process

$$\mathbf{x}_R | \boldsymbol{\theta} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{Q}^{-1}(\boldsymbol{\theta})), \quad (23)$$

with zero mean and the inverse covariance matrix $\mathbf{Q}(\boldsymbol{\theta})$. In this case $\mathbf{x}_R = \mathbf{u}$, where \mathbf{u} is the random variable specified through the SPDE.

The marginal variance for the solution \mathbf{u} of the SPDE, can be derived through a series of calculations. The transfer function for Equation (18) is $g(w) = (\kappa^2 + \mathbf{w}^T \mathbf{H} \mathbf{w})^{-1}$ and given the spectral density of Gaussian white noise in \mathbb{R}^3 is $(2\pi)^{-3}$ the spectral density of the solution of the SPDE is

$$f_S(\mathbf{w}) = (2\pi)^{-3} (\kappa^2 + \mathbf{w}^T \mathbf{H} \mathbf{w})^{-1}. \quad (24)$$

Lastly, the variance of the solution is found by integrating its spectral density over the whole domain \mathbb{R}^3 , and thus, the variance of the process we are trying to describe is

$$\text{Var}(x(\mathbf{s})) = \frac{1}{8\pi\kappa(\mathbf{s})\sqrt{\det(\mathbf{H}(\mathbf{s}))}}. \quad (25)$$

Moreover, in a stationary case where $\kappa(\mathbf{s}) = \kappa$ and $\mathbf{H}(\mathbf{s}) = \mathbf{H}$ for all \mathbf{s} the covariance of the process between two locations \mathbf{s}_1 and \mathbf{s}_2 in \mathbb{R}^3 can be written as

$$\text{Cov}(x(\mathbf{s}_1), x(\mathbf{s}_2)) = \frac{1}{8\pi\kappa\sqrt{\det(\mathbf{H})}} \exp\left(-\kappa\|\mathbf{H}^{-1/2}(\mathbf{s}_1 - \mathbf{s}_2)\|\right), \quad (26)$$

or namely an exponential covariance function.

Parameter inference

Following the notation from the process defined in Equation (23) the data model for the innovations are

$$\mathbf{y}_R | \mathbf{x}_R, \sigma_S^2 \sim \mathcal{N}_m(\mathbf{A}\mathbf{x}_R, \sigma_S^2 \mathbf{I}_n). \quad (27)$$

Here, matrix \mathbf{A} is a m by n matrix linking the locations of the observation \mathbf{y}_R to our discretization in the process \mathbf{x}_R , and σ_S^2 is the independent noise in the innovations.

To find the optimal parameters such that our models best describe the innovations we will maximize the likelihood function of the parameters. Specifically, the parameters have a multivariate Gaussian distribution, and following common practice we will optimize the logarithmic transformation of this likelihood, thereby the log-likelihood:

$$\begin{aligned} \ell(\boldsymbol{\theta}, \sigma_S^2 | \mathbf{y}) &= \text{Const} + \log \pi(\boldsymbol{\theta}, \sigma_S^2) + \frac{1}{2} \log \det(\mathbf{Q}) - \frac{m}{2} \log(\sigma_S^2) \\ &\quad - \frac{1}{2} \log \det(\mathbf{Q}_C) - \frac{1}{2} \boldsymbol{\mu}_C^T \mathbf{Q}_C \boldsymbol{\mu}_C - \frac{1}{2\sigma_S^2} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_C)^T (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}_C). \end{aligned} \quad (28)$$

The reader is referred to the supplementary material of Berild and Fuglstad (2023) for a full derivation of the log-likelihood. In Equation (28), \mathbf{Q}_C is the conditional precision matrix, i.e. the precision matrix given that the model has seen \mathbf{y} ,

$$\mathbf{Q}_C = \mathbf{Q} + \mathbf{A}^T \mathbf{A} / \sigma_S^2. \quad (29)$$

Similarly, the conditional mean is

$$\boldsymbol{\mu}_C = \mathbf{Q}_C^{-1} \mathbf{A}^T \mathbf{A} \mathbf{y} / \sigma_S^2. \quad (30)$$

Note that this is similar to the equations used in the model updating, but with $\boldsymbol{\mu} = 0$ as assumed by our process of these innovations.

The parameter space is quite challenging to explore so we use an analytical expression for the gradient to determine the search directions. In order to speed up the gradient calculations a stochastic version is calculated, and therefore, ultimately the optimization strategy is a stochastic gradient descent algorithm. Furthermore, we have employed a root mean square propagation (RMSprop) in the optimization to improve the stability and convergence.

EIBV design criteria

For any design \mathbf{d} at stage t giving data $\mathbf{y} = \mathbf{y}_d$, the EIBV reduction in can be rephrased as

$$\begin{aligned} \text{EIBV}(\boldsymbol{\mu}_{C,t-1}, \mathbf{Q}_{C,t-1}) &= \int E_{\mathbf{y}|\mathcal{Y}_{t-1}} \{p_{\mathbf{s}}(\mathbf{y}, \mathcal{Y}_{t-1}) [1 - p_{\mathbf{s}}(\mathbf{y}, \mathcal{Y}_{t-1})]\} d\mathbf{s}, \\ p_{\mathbf{s}}(\mathbf{y}, \mathcal{Y}_{t-1}) &= P(\mathbf{x}(\mathbf{s}) \leq \ell | \mathbf{y}, \mathcal{Y}_{t-1}), \end{aligned} \quad (31)$$

where \mathcal{Y}_{t-1} denotes all the data gathered at stages before t . The probability $p_{\mathbf{s}}(\mathbf{y}, \mathcal{Y}_{t-1})$ is a Gaussian cumulative distribution function (CDF) with linear conditioning to \mathbf{y} in the mean and with a variance that does not depend on the outcome of the data.

The conditional mean at stage $t - 1$ is $\boldsymbol{\mu}_{C,t-1}$ with entries $\mu_{C,t-1}(\mathbf{s}_i)$, $i = 1, \dots, n$. The conditional covariance matrix is $\boldsymbol{\Sigma}_{C,t-1} = \mathbf{Q}_{C,t-1}^{-1}$ with diagonal entries $\sigma_{C,t-1}^2(\mathbf{s}_i)$, and after the updating we have $\boldsymbol{\Sigma}_{C,t} = \mathbf{Q}_{C,t}^{-1}$ with diagonal entries $\sigma_{C,t}^2(\mathbf{s}_i)$. Based on results of Chevalier et al. (2014) and Fossum et al. (2021), the EIBV in Equation (31)

can be evaluated in closed form as a bivariate Gaussian CDF depending on these parameters.

$$\begin{aligned} \text{EIBV}(\boldsymbol{\mu}_{C,t-1}, \mathbf{Q}_{C,t-1}) &= \sum_{i=1}^n \text{EBV}(\mu_{C,t-1}(\mathbf{s}_i), \sigma_{C,t-1}^2(\mathbf{s}_i), \sigma_{C,t}^2(\mathbf{s}_i)) \\ \text{EBV}(\eta, v^2, w^2) &= \Phi_2 \left(\begin{bmatrix} \ell \\ -\ell \end{bmatrix}; \begin{bmatrix} \eta \\ -\eta \end{bmatrix}, \begin{bmatrix} v^2 & (w^2 - v^2) \\ (w^2 - v^2) & v^2 \end{bmatrix} \right), \end{aligned} \quad (32)$$

where Φ_2 denotes the bivariate Gaussian cumulative distribution function.

We note that this closed form EIBV calculation in equation (32) relies on the variance terms and their reduction in the updating step ($w^2 = \sigma_{C,t}^2(\mathbf{s}_i)$ compared with $v^2 = \sigma_{C,t-1}^2(\mathbf{s}_i)$). In the formulation with the precision matrix, these terms are not immediately available. Matrix recursions exist for computing the marginal variance terms from the precision matrix, see e.g. Zammit-Mangion and Rougier (2018), but the solution can be computationally challenging in 3D because of rather large fill-in of non-zeros in the sparse matrix structure during the recursion. We instead approximate the required variance terms by Monte Carlo sampling from the GMRF model.

Conditional samples are here generated by a trick known as conditioning by Kriging equation (see e.g. Wackernagel (2003)). This relies on the following steps; first an unconditional sample $\mathbf{x}_{C,t-1}^b$ of the field is generated. In our case this comes from the Gaussian distribution with mean $\boldsymbol{\mu}_{C,t-1}$ and precision matrix $\mathbf{Q}_{C,t-1}$. Next, a synthetic data sample $\mathbf{y}^b = \mathbf{A}\mathbf{x}_{C,t-1}^b + \boldsymbol{\epsilon}^b$, $\boldsymbol{\epsilon}^b \sim N(0, \sigma_{\text{auv}}^2 \mathbf{I}_m)$ is generated according to the specified design. Finally, a conditional sample is formed by solving the equation for the conditional mean (also known as the Kriging equation), given the synthetic data:

$$\mathbf{x}_{C,t}^b = \mathbf{x}_{C,t-1}^b + \mathbf{Q}_{C,t-1}^{-1} \mathbf{A}^T (\mathbf{y}^b - \mathbf{A}\mathbf{x}_{C,t-1}^b) / \sigma_{\text{auv}}^2. \quad (33)$$

This procedure is repeated for $b = 1, \dots, B$ independent Monte Carlo samples. We used $B = 100$ in our implementation. Equation (33) requires matrix-vector solves with the sparse precision matrix which can be done very fast.

When the optimal design is selected, the AUV acts to move in the direction of the selected design. It senses salinity data \mathbf{y} , and then Equation (33) is used with these in-situ observations in place of \mathbf{y}^b to get the conditional samples for time stage t . The resulting Monte Carlo sample $\mathbf{x}_{C,t}^b$, $b = 1 \dots, B$ forms the basis for the EIBV evaluation at the next time step, when $t \rightarrow t + 1$.

Paper IV

Using expected improvement of gradients for robotic exploration of ocean salinity fronts

André Julius Hovd Olaisen, Yaolin Ge and Jo Eidsvik

Submitted to Environmetrics

Using expected improvement of gradients for robotic exploration of ocean salinity fronts

André Julius Hovd Olaisen, Yaolin Ge and Jo Eidsvik
Department of Mathematical Sciences, NTNU,
7491 Trondheim, Norway
*Corresponding author: andre.j.h.olaisen@ntnu.no

Abstract

We study opportunities for dynamic sampling designs in spatio-temporal random field. Considering a situation with a robotic agent, we develop an algorithm that enables autonomous exploration of spatial domains with large gradients. The modeling assumptions rely on a spatio-temporal Gaussian random field, which means that the directional derivatives of the field are Gaussian distributed. Using computational tricks at the dimension of relatively sparse data, the robot updates its on-board Gaussian random field model in real-time. Moreover, it computes the expected improvement in directional derivatives along a set of possible paths in a spider-leg search space to choose intelligent exploration designs over time stages. We study statistical properties of this suggested approach in a simulation study, where we compare the design criterion with several other viable design selection criteria. The new algorithm is embedded on an autonomous underwater vehicle which is deployed for characterizing a river plume frontal system in a Norwegian fjord. Using expected improvement for the salinity field derivatives, the vehicle successfully sampled the river front for more than two hours without intervention.

Keywords: Expected improvement; Gaussian Random Field; Spatial design; Robotics; Oceanography

1. Introduction

Inspired by new sensor technology and small-size computing units, there is currently a drive to develop intelligent monitoring systems. This development is pushed by engineers and multi-disciplinary visions on how to put innovative solutions into practical use. Recent examples include internet-of-things for smart sensor networks monitoring air pollution (Dhingra et al., 2019), embedded systems and AI for agriculture (Shadrin et al., 2019), robotic systems for understanding environmental processes (Dunbabin and Marques, 2012) and cyber-physical systems that can re-configure themselves for ecological monitoring (Schranz et al., 2021).

The capabilities of such embedded systems can clearly be improved by leveraging knowledge from spatio-temporal statistics and design of experiments, see e.g. Mateu and Müller (2012), Wang et al. (2012, 2020) or Brus (2022). In doing so, one can develop more principled approaches for what, where and when to gather additional data samples, and integrate this new information in a consistent statistical modeling framework. Even so, solutions to these situations tend to be case-specific, and often of a heuristic type as the search space is too large to find the optimal solution. Impactful examples of spatial statistics and design for embedded systems include Krause et al. (2008) who studied the NP-hard problem of sensor placement using Gaussian random fields (GRFs) with the goal of finding designs that optimize the mutual information and Manohar et al. (2018) who suggested using machine learning methods to facilitate the search for constructive design patterns.

As robotic units and sensor systems often have limited computing, storage and communication capabilities, one must often simplify the modeling to guide the optimization challenge. Ideas from designing computer experiments and building surrogate models (Gramacy, 2020; Fuhg et al., 2021) are hence also highly relevant in this context of optimal spatio-temporal design and fast integration of data.

In this article we focus on an application of underwater robotics. An autonomous underwater vehicle (AUV) with onboard computing units uses a GRF surrogate model to plan where and when to explore various parts of an ocean domain. The AUV is hence a dynamic agent that can navigate to extract informative samples in an uncertain dynamic environment. In particular, the goal here is to find large derivatives of the field which are indicative of important frontal zones between different water masses. The spatio-temporal variable that we target here is ocean salinity, and via the real-world deployment, we show results of an AUV deployment characterizing salinity changes in a river plume front in a Norwegian fjord.

The main contributions of this paper are

- A spatio-temporal sampling approach with a dynamic agent searching for large derivatives in the field.
- A fast scalable algorithm for updating and planning based on GRFs and sparse observation points in the vicinity of the current position.
- A field deployment with an AUV adaptively sampling a river plume front for 2 hours and 10 minutes.

In Section 2, we describe the motivation for our work and define the necessary notation. In Section 3, we set up the required building blocks from theory on spatio-temporal GRFs and the properties of their derivatives. In Section 4, we present the method and algorithm for adaptive sampling of large directional derivatives. In Section 5, we demonstrate properties of the suggested algorithm in a simulation study. In Section 6, we show results of the AUV deployment in the Trondheim fjord in Norway. In Section 7, we provide conclusions and point to future work.

2. Background and notation

Fronts are important in meteorology and oceanography as they tend to be key drivers of the physical dynamic behavior, see e.g. Fedorov (1986) and Catto and Pfahl (2013). Frontal zones in the ocean are further known to be biological hot-spots that shape parts of the marine ecosystem (Belkin et al., 2009). In this paper, the spatio-temporal variable of interest is ocean salinity and its derivatives which capture the frontal zone near river plumes.

Ocean fronts can be detected from satellite data, see e.g. Hopkins et al. (2010), but this is only on the surface and not available on a cloudy day. Numerical ocean models, see e.g. Lermusiaux (2006), can mimic fronts at various scales, but even though they are incredibly useful at predicting ocean variables, they tend to be biased in space and time. AUVs have become an important tool for oceanographic in-situ sampling, and they are commonly used to detect frontal zones. These vehicles can navigate autonomously underwater and can hold a range of sensors such as a standard tool that provides salinity measurements. Many AUVs also have an onboard computer that enables for instance data assimilation in a model, and using this to adapt its trajectory and move in more interesting directions. This is important because communication is limited under water, and the full benefits of an AUV are gained only when it acts on its own as an intelligent agent.

To further motivate the detection of fronts in the ocean, we highlight a few examples. Figure 1 shows three different deployments where researchers aimed to map the frontal zone between water masses. Zhang et al. (2019) used temperature data from satellites along with in-situ AUV measurements to follow the zone of mixing cold and warm water masses in the Monterey Bay, California. Fossum et al. (2021) used an AUV to understand the frontal zone near the ice shelf in Arctic waters. Fonseca et al. (2023) compared satellite imagery and AUV samples to map the front of chlorophyll

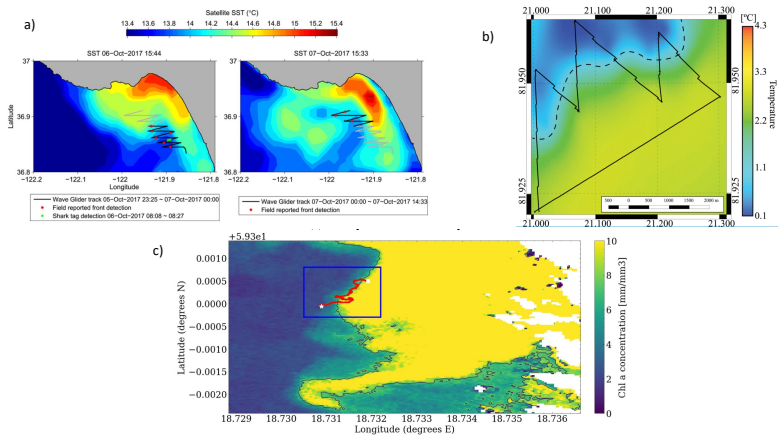


Figure 1: Examples of AUV exploration of fronts. a) Zhang et al. (2019) characterizing the Monterey Bay front between water masses using temperature information processed from satellites and that gathered by an AUV. b) Fossum et al. (2021) conducting frontal AUV sampling in water masses in the Arctic. c) Fonseca et al. (2023) showing an AUV path zig-zagging the chlorophyll front in the Baltic sea as extracted from satellite data.

in the Baltic sea. These studies attempt to find the gradient or derivative in the ocean variable of interest. In doing so, the AUV reacts to data, but none of them use spatio-temporal statistical models or approaches from spatial design, which would likely have improved the mapping performance.

We next define the notation used in our statistical model and sampling design approach. Let $\mathbf{s} = (s_e, s_n, s_d, t)$ be a point in space and time. Here, s_e , s_n and s_d represent east, north and depth coordinates, respectively, while $t > 0$ is a temporal index. A spatial operational domain is defined so that $(s_e, s_n, s_d) \in \mathcal{D} \subset \mathcal{R}^3$. The spatio-temporal variable of interest is denoted $x(\mathbf{s}) \in \mathcal{R}$. In our application this is ocean salinity.

In this paper, we are primarily interested in detecting large derivatives or changes in this spatio-temporal variable in the lateral plane close to the sea surface. In our application this would indicate ocean front zones. We define the directional difference

from location \mathbf{s} to \mathbf{s}' by

$$g(\mathbf{s}, \mathbf{s}') = \frac{x(\mathbf{s}') - x(\mathbf{s})}{d(\mathbf{s}, \mathbf{s}')}, \quad (1)$$

where $d(\mathbf{s}, \mathbf{s}')$ is the Euclidean distance between the two locations. Letting this distance go to 0, we obtain the field derivative at \mathbf{s} in the direction towards \mathbf{s}' . In practice, we instead consider the distance $d(\mathbf{s}, \mathbf{s}')$ as a tuning parameter that can be specified in the context of the application and the operational constraints.

An observation made at space-time location \mathbf{s} is denoted $y(\mathbf{s})$. Because of sensor noise and positioning error, this observation does not carry perfect information about the salinity. Observations from a set of sampling points $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_N\}$ are denoted by $\mathbf{y}(\mathcal{S}) = (y(\mathbf{s}_1), y(\mathbf{s}_2), \dots, y(\mathbf{s}_N))$. AUV data are gathered sequentially. At stage k , the AUV gathers a batch of data size N_k , and we denote batch sampling locations by $\mathcal{S}_k = \{\mathbf{s}_1^{(k)}, \dots, \mathbf{s}_{N_k}^{(k)}\}$ with associated data $\mathbf{y}_k = \{y(\mathbf{s}_1^{(k)}), \dots, y(\mathbf{s}_{N_k}^{(k)})\}$. This means that we at stage k have measured at $N_{1:k} = \sum_{l=1}^k N_l$ points. We denote the set of sampling locations by

$$\mathcal{S}_{1:k} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{k-1}, \mathcal{S}_k\}, \quad k = 1, 2, \dots, \quad (2)$$

with associated salinity measurements

$$\mathbf{y}(\mathcal{S}_{1:k}) = \mathbf{y}_{1:k} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{k-1}, \mathbf{y}_k), \quad k = 1, 2, \dots \quad (3)$$

At each stage k , the agent computes expected rewards for staying on the same trajectory and for changing its path to another direction. Higher rewards are attained for design directions that have large expected derivatives. Design paths form transects consisting of M_k^j new locations in a set $\mathcal{P}_k^j = \{\mathbf{p}_{k,l}^j; j = 1, \dots, J, l = 1, \dots, M_k^j\}$ with single locations $\mathbf{p}_{k,l}^j \in \mathcal{D}$. For the number of directions J , we use a spider-leg formation in the lateral domain. The number of transect points M_k^j is fixed at all stages and for all designs, except at boundary locations. The spacing between single locations along each design direction is also fixed, and it is determined by the agent's speed

and sampling frequency. In doing so, measurements and plan evaluations are easily comparable.

Figure 2 shows the situation with an agent path consisting of 5 stages. At the current location (blue circle), it makes a decision about where to go next for stage 6. The agent is not constrained to travel the entire segment of length M_j . Instead, it

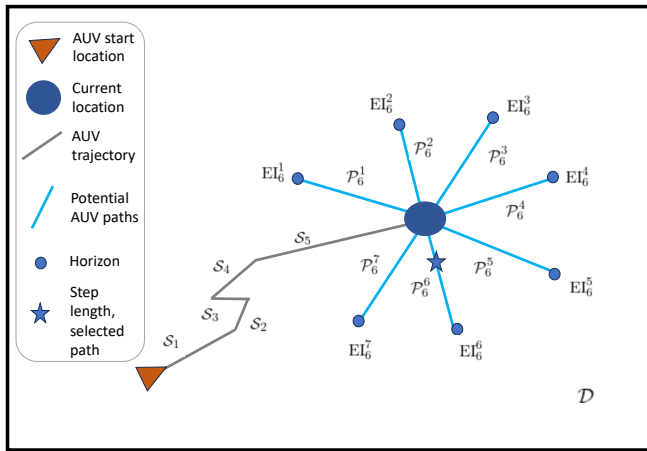


Figure 2: Illustration of an AUV trajectory made up of segments over 5 earlier stages. At the current location, the AUV will choose one of the 7 possible designs. The design selection criteria is expected improvement (EI) for the salinity derivative.

conducts new design evaluations after moving one step-length down the best segment. In this illustration, design 6 is selected and the agent moves a step-length in this direction (marked by a star).

In summary, bringing the model and design choices back to the context of AUV sampling, we assume that:

- The AUV moves significantly faster than the ocean phenomenon develops over time, and this means that the AUV is able to detect changes in space. Hence we focus on differences in space, and not time.

- Salinity changes are often most difficult to characterize in the lateral domain, so without loss of generality, we focus on differences in the east, north plane, assuming the AUV is at a fixed depth (set to 1 meter in the field deployment).
- The AUV is able to maintain a nearly constant velocity. Because the sensors are sampling at a constant frequency, the relatively close locations \mathbf{s} and \mathbf{s}' are at the same distance during the operation.

Note that the suggested approach does not rely on the usual concepts of a predefined waypoint graph or grid for the path-planning. Instead, data points and variables are allocated to continuous space-time locations, and this occurs when the design criterion is computed according to a spider leg design. Hence, the discretization occurs only along transect lines, and it is formed during the mission, not before the deployment starts. One benefit of this approach is that the AUV maintains a model with relatively few points compared to a waypoint graph, but we can still have a high level of detail close to where we sample. In practice, one might miss the phenomenon by placing a strict waypoint graph onboard the AUV model. Here, the AUV is more free to follow where it is most interesting to sample.

3. Spatio-temporal Gaussian random fields

The agent has an onboard spatio-temporal model which is updated with the data that is gathered. The model is also used to compute the expected rewards along potential design trajectory and decisions for adaptive sampling. For ease of data assimilation and real-time decision-making, a GRF is used onboard the agent. Notably, derivatives or differences are then also Gaussian distributed.

3.1. Gaussian random fields

A GRF is fully described by its mean $\mu(\mathbf{s}) = \mathbb{E}[x(\mathbf{s})]$, $\mathbf{s} \in \mathcal{D} \times \mathcal{R}^+$ and a covariance function $C(\mathbf{s}, \mathbf{s}') = \text{Cov}(x(\mathbf{s}), x(\mathbf{s}'))$, see e.g. Cressie and Wikle (2015). In our case

study, the mean function is specified from physical oceanography modeling for the domain of interest, which involves a spatio-temporally varying function in the initial (prior) mean. The covariance function is specified from multiple ocean models as well as previously acquired data from the domain of interest. For the space-time covariance, we assume a separable model so that

$$C(\mathbf{s}, \mathbf{s}') = \sigma^2 \exp\left(-\left(\frac{d(\mathbf{s}, \mathbf{s}')}{\phi_s}\right)^2\right) \exp\left(-\left(\frac{|t_1 - t_2|}{\phi_t}\right)^2\right), \quad (4)$$

and as we limit scope to a fixed operational depth, we do not incorporate any kind spatial anisotropy which would be relevant to enhance smaller correlation in depth than in the lateral domain.

For any set of N space-time locations $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$, the random vector $\mathbf{x}_{\mathcal{S}} = (x(\mathbf{s}_1), x(\mathbf{s}_2), \dots, x(\mathbf{s}_N))$ is then Gaussian distributed with mean vector

$$\boldsymbol{\mu}_{\mathcal{S}} = (\mu(\mathbf{s}_1), \mu(\mathbf{s}_2), \dots, \mu(\mathbf{s}_N)), \quad (5)$$

and a symmetric positive semi-definite covariance matrix

$$\boldsymbol{\Sigma}_{\mathcal{S}} = \begin{pmatrix} C(\mathbf{s}_1, \mathbf{s}_1) & C(\mathbf{s}_1, \mathbf{s}_2) & \dots & C(\mathbf{s}_1, \mathbf{s}_N) \\ C(\mathbf{s}_2, \mathbf{s}_1) & C(\mathbf{s}_2, \mathbf{s}_2) & & C(\mathbf{s}_2, \mathbf{s}_N) \\ \vdots & & \ddots & \\ C(\mathbf{s}_N, \mathbf{s}_1) & C(\mathbf{s}_N, \mathbf{s}_2) & & C(\mathbf{s}_N, \mathbf{s}_N) \end{pmatrix}. \quad (6)$$

In short notation, we write this as

$$\mathbf{x}_{\mathcal{S}} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathcal{S}}, \boldsymbol{\Sigma}_{\mathcal{S}}). \quad (7)$$

3.2. Directional differences and GRFs

Because $g(\mathbf{s}, \mathbf{s})$ in Equation (1) is a linear combination of two Gaussian distributed variables, and hence $g(\mathbf{s}, \mathbf{s})$ is also Gaussian. In particular, we have mean

$$E[g(\mathbf{s}, \mathbf{s}')] = \frac{\mu(\mathbf{s}') - \mu(\mathbf{s})}{d(\mathbf{s}, \mathbf{s}')}, \quad (8)$$

and variance

$$\begin{aligned}\text{Var}(g(\mathbf{s}, \mathbf{s}')) &= \frac{1}{d(\mathbf{s}, \mathbf{s}')^2} (\text{Var}(x(\mathbf{s})) + \text{Var}(x(\mathbf{s}')) - 2\text{Cov}(x(\mathbf{s}), x(\mathbf{s}'))) \\ &= \frac{1}{d(\mathbf{s}, \mathbf{s}')^2} (C(\mathbf{s}, \mathbf{s}) + C(\mathbf{s}', \mathbf{s}') - 2C(\mathbf{s}, \mathbf{s}')).\end{aligned}\quad (9)$$

In doing so, one can further take the difference between any pairs of variables along a transect line in the spatial domain. Then the random vector of directional differences is multivariate Gaussian distributed.

When studying properties of such derivatives we can see one of the main benefits of using a Gaussian covariance function. Figure 3 shows three different 1D GRF realizations, one using a Gaussian covariance function, one using an exponential covariance function and one using a Matérn covariance function. For each of the covariance functions the correlation at 300 is 0.05. The realization using the Gaussian covariance function is the smoothest out of the three, and the gradients are also smooth. For the exponential the derivatives are extremely large. This means that one should impose a smooth correlation function when the goal is to search for hot-spots in gradients. Hence, even though we regard more complicated spatial or spatio-temporal correlation functions as promising models, such as the one with spatially varying anisotropy by Berild and Fuglstad (2023) or the advection-diffusion model of Foss et al. (2022) that have been applied to coastal domain ocean modeling, we did not pursue complex covariance models here.

3.3. Conditioning to in-situ observations

As the agent gathers data, it will update the on-board model. In doing so, it needs a model for the data. The measurement model is here defined via

$$y(\mathbf{s}) = x(\mathbf{s}) + \epsilon(\mathbf{s}) \quad \epsilon(\mathbf{s}) \sim \mathcal{N}(0, \tau^2), \quad (10)$$

where the errors at different locations are assumed to be independent. The Gaussian assumption for the error terms crucially means that measurements are jointly Gaus-




Figure 3: One realization using Gaussian, Matérn (smoothness $\nu = 3/2$) and exponential covariance function. The Gaussian covariance function gives a much smoother realization and gradient. The plotting scale for the gradients is truncated because the gradients for the exponential become extremely large.

sian distributed. Hence, given observations $\mathbf{y}_{1:k}$ in Equation (3), we can compute the conditional model using properties of the Gaussian distribution.

For any location set $\mathcal{P} \in \mathcal{P}_{k+1}^j$ among all possible design sets at stage $k + 1$, we denote the associated variable $\mathbf{x}_{\mathcal{P}}$. Given the currently available data, the mean vector and covariance matrix are computed as follows:

$$\mathbf{m}_{\mathcal{P}} = \boldsymbol{\mu}_{\mathcal{P}} + \boldsymbol{\Sigma}_{\mathcal{P}, \mathcal{S}_{1:k}} (\boldsymbol{\Sigma}_{\mathcal{S}_{1:k}} + \mathbf{T}_{\mathcal{S}_{1:k}})^{-1} (\mathbf{y}_{\mathcal{S}_{1:k}} - \boldsymbol{\mu}_{\mathcal{S}_{1:k}}), \quad (11)$$

$$\boldsymbol{\Psi}_{\mathcal{P}} = \boldsymbol{\Sigma}_{\mathcal{P}} - \boldsymbol{\Sigma}_{\mathcal{P}, \mathcal{S}_{1:k}} (\boldsymbol{\Sigma}_{\mathcal{S}_{1:k}} + \mathbf{T}_{\mathcal{S}_{1:k}})^{-1} \boldsymbol{\Sigma}_{\mathcal{P}, \mathcal{S}_{1:k}}^T, \quad (12)$$

where $\mathbf{T}_{\mathcal{S}_{1:k}} = \tau^2 \mathbf{I}_{N_{1:k}}$ is the measurement noise covariance matrix and $\boldsymbol{\Sigma}_{\mathcal{P}, \mathcal{S}_{1:k}}$ is the cross-covariance between variables at locations \mathcal{P} and those variables at former data

locations $\mathcal{S}_{1:k}$.

In particular, at any two points \mathbf{s} and \mathbf{s}' , the variable $(x(\mathbf{s}), x(\mathbf{s}'))$ has a joint bivariate Gaussian distribution conditional on the observations $\mathbf{y}_{1:k}$. Their scaled difference is then Gaussian distributed, similar to what we have in Equation (8)-(9), and the same holds for variables along a transect.

3.4. Efficient matrix calculations

Matrix inversion or factorization can take a considerable amount of time when the number of data increases. Say, in Equation (11) and (12), one must invert the matrix $\Sigma_{\mathcal{S}_{1:k}} + \mathbf{T}_{\mathcal{S}_{1:k}}$ of dimension $N_{1:k} \times N_{1:k}$, which is of order $O(N_{1:k}^3)$ calculations, and could quickly stall the agents' computing performance. We utilize the structure with batch data collection, and then rely on a block version of the Sherman-Woodbury-Morrison formula for efficient matrix factorization (Petersen et al., 2008). In particular, we have that

$$\begin{bmatrix} \Sigma_{\mathcal{S}_{1:k}} + \mathbf{T}_{1:k} & \Sigma_{\mathcal{S}_{1:k}, \mathcal{P}} \\ \Sigma_{\mathcal{P}, \mathcal{S}_{1:k}} & \Sigma_{\mathcal{P}} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{B} & -\mathbf{A}\Sigma_{\mathcal{S}_{1:k}, \mathcal{P}}\mathbf{C} \\ -\mathbf{C}\Sigma_{\mathcal{P}, \mathcal{S}_{1:k}}\mathbf{A} & \mathbf{C}^{-1} \end{bmatrix}, \quad (13)$$

where $\mathbf{A} = [\Sigma_{\mathcal{S}_{1:k}} + \mathbf{T}_{1:k}]^{-1}$ is assumed to be available from the previous stage, and

$$\mathbf{B} = \mathbf{A} + \mathbf{A}\Sigma_{\mathcal{S}_{1:k}, \mathcal{P}}\mathbf{C}^{-1}\Sigma_{\mathcal{P}, \mathcal{S}_{1:k}}\mathbf{A}, \quad \mathbf{C} = \Sigma_{\mathcal{P}} - \Sigma_{\mathcal{P}, \mathcal{S}_{1:k}}\mathbf{A}\Sigma_{\mathcal{S}_{1:k}, \mathcal{P}}.$$

This calculation is used both to evaluate many designs $\{\mathcal{P}_{k+1}^j\}_{j=1,2,\dots,J}$ and to update the mean and covariance in the data assimilation step. The required inversion is for \mathbf{C} is of moderate size as it only involves the variables at the new batch or the potential transect locations.

This trick in Equation (13) allows efficient computing onboard an agent. However, over time the mere size of the covariance matrix leads to evaluation challenges. To approach this challenge further, we implement an on-board algorithm that reduces the data size over space-time by thinning data from far away/long ago. The agent

needs to make a decision in a reasonable amount of time, therefore we set a threshold time for how long the data assimilation and prediction stage should take. If the total time is larger than this threshold time then we thin the data points in memory. This will remove half the points in memory. The points that are thinned are mostly points that are far away in time because these points have a low correlation with the points we want to predict. After the thinning the inverse must be re-computed without using the recursive formula in Equation (13), but with the reduced number of data points this will take a shorter time than the threshold time.

4. Adaptive sampling design

For adaptive selection of designs as illustrated by the spider leg design in Figure 2, this shows the agent at some stage k deciding to take one of 7 possible paths. In order to decide what path to take one needs to have an objective function. The agent will then make an optimal decision (direction) from the highest expected reward. We outline expected improvement (EI) in directional differences as our reward function. The onboard algorithm is summarized with the GRF and EI calculations, and involves some tuning parameters that we have tailored for the application.

4.1. Expected improvement for spider legs transects

The spider leg designs shown in Figure 2 illustrate potential sampling designs for the agent at the current stage. For each of these transects, we calculate the conditional distribution, and base the path selection on the optimal expected reward. The expectation for a particular set of design points \mathcal{P} is based on the conditional mean and covariance in Equation (11)-(12). In addition, we are mainly interested in the directional differences, see Equation (8)-(9).

Let g_{\max} denote the largest absolute directional difference observed thus far in the sampling. We first study the probability of finding a larger directional gradient than this g_{\max} along a transect. This probability of improvement (PoI) is chance of

having a difference that is larger than g_{\max} . Note that the maximum difference is in the observations and not in the true field. Nevertheless, we compare designs in the variables of interest because it is comparable between the different design transects at this stage. For a transect path $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n\}$, two-neighbour locations define the difference. The difference is Gaussian distributed as in Equation (8)-(9). To simplify the notation, we let the conditional distribution of the difference $g_i = g(\mathbf{p}_i, \mathbf{p}_{i+1}) \sim \mathcal{N}(\zeta_i, \eta_i^2)$, given the available data $\mathbf{y}_{1:k}$. Then the probability that $|g(\mathbf{p}_i, \mathbf{p}_{i+1})|$ is larger than g_{\max} is

$$\text{PoI} = P(|g(\mathbf{p}_i, \mathbf{p}_{i+1})| \geq g_{\max}) = 1 - \Phi\left(\frac{g_{\max} - \zeta_i}{\eta_i}\right) + \Phi\left(\frac{-g_{\max} - \zeta_i}{\eta_i}\right). \quad (14)$$

Now we compute this for all two-neighbor locations along a transect and compare the different transects to create a decision rule for which direction the agent should choose. The best transect according to the largest PoI is

$$\text{Best direction}_{\text{Prob}} = \arg \max_{j \in \{1, 2, \dots, J\}} \max_{\mathbf{p}_i \in \mathcal{P}_k^j} P(|g(\mathbf{p}_i, \mathbf{p}_{i+1})| \geq g_{\max}). \quad (15)$$

Note that even if g_{\max} is the highest derivative in the field the PoI still gives a value larger than zero along each transect provided that the variance is larger than zero.

An alternative approach that accounts for the expected gain in the difference is available via the EI which has been used much in for instance the design of complex optimization problems (Zhan and Xing, 2020). The EI has a closed form solution for Gaussian distributions, see e.g. Gramacy and Apley (2015). Let

$$I(g_i) = \max(|g_i| - g_{\max}, 0), \quad (16)$$

then the expected value of this improvement becomes

$$\begin{aligned} E[I(g_i)] &= (\zeta_i - g_{\max}) \left(1 - \Phi\left(\frac{g_{\max} - \zeta_i}{\eta_i}\right)\right) + \eta_i \phi\left(\frac{g_{\max} - \zeta_i}{\eta_i}\right) \\ &+ (-\zeta_i - g_{\max}) \Phi\left(\frac{-g_{\max} - \zeta_i}{\eta_i}\right) + \eta_i \phi\left(\frac{-g_{\max} - \zeta_i}{\eta_i}\right). \end{aligned} \quad (17)$$

Now we compute the EI for all two-neighbor locations along a transect and compare the different transects to create a decision rule for which direction the agent should choose. The best transect according to the largest EI is then

$$\text{Best direction}_{EI} = \arg \max_{j \in \{1, 2, \dots, J\}} \max_{\mathbf{p}_i \in \mathcal{P}_k^j} E(I(g_i)). \quad (18)$$

Later in Section 5 we will compare PoI against EI in a simulation study.

4.2. Algorithm

Along with the statistical model and the objective functions some other details are needed to fully describe the algorithm. The algorithm works in a sequential loop; sampling data, data assimilation, predicting gradients along possible paths and then use the objective function to choose one of these paths. The algorithm using EI is shown in Algorithm 1. When the AUV arrives at a waypoint (WP) it first assimilates the new data into the model (Section 3), then it finds 7 possible new paths it can take as defined by the spider-web legs. For each of these paths it predicts the salinity changes for several points along this path. It uses this prediction and EI (described in Section 4) to choose which transect is the best. The last step is to move along this transect, and it samples salinity data with a frequency of 1 Hz. The AUV does not move all the way until the end of the predicted transect, rather it uses a longer horizon and then it moves a shorter step-length. This is illustrated by the star in Figure 2. The reason for predicting far into the future and only moving a short path is that we can react quickly to the new measurements while still looking for derivatives far away.

Figure 4 illustrates some parts of the algorithm from the simulation study. In the figure the AUV has sampled for 20 steps and is making a decision on where to move next. There are 8 candidate transects for the AUV to move in. The next WP is set back towards the where the salinity change is large. Along the observed path

the conditional mean is closer to the true field. The conditional mean for the whole field is not computed during the mission, but here it is included only for illustration. The algorithm does require that we choose a step-length, horizon and a number of transects that fit the application.

Algorithm 1 shows the main steps in the sequential procedure. The AUV is guided by setting target waypoints (WP_k) for stage *k*.

Algorithm 1 Sampling for derivatives.

Require:

$\mu(\mathbf{s})$, $C(\mathbf{s}, \mathbf{s}')$, operational domain \mathcal{D} .

$\mathcal{S}_0 = \emptyset$, $\mathbf{y}_0 = \emptyset$, $g_{\max} = 0$, $\text{WP}_0 = \mathbf{s}_{\text{start}}$.

repeat For each time $k = 1, \dots$:

 Define spider legs \mathcal{P}_k^j for $j = 1, \dots, J$ transects.

 Define $\text{WP}_{k-1} = \text{WP}_k$ and $\text{EI}_{\max} = 0$

repeat For each spider-leg \mathcal{P}_k^j , $j = 1, \dots, J$

 Predict $\mathbf{m}_{\mathcal{P}_k^j}$ and $\Psi_{\mathcal{P}_k^j}$ from $\mathcal{S}_{1:k-1}$ and $\mathbf{y}_{1:k-1}$ ▷ Eq. (11)-(12).

 Compute $\text{EI}_k^j = \max_{\mathbf{p}_i \in \mathcal{P}_k^j} E[I(g_i)]$ ▷ Eq. (17).

if $\text{EI}_j > \text{EI}_{\max}$ **then**

$\text{EI}_{\max} = \text{EI}_k^j$

 Set WP_k one step-length down \mathcal{P}_k^j

end if

until

 AUV moves from WP_{k-1} to WP_k , and it gathers data values \mathbf{y}_k at points \mathcal{S}_k

 Update $\mathbf{y}_{1:k} = (\mathbf{y}_{1:k-1}, \mathbf{y}_k)$ and $\mathcal{S}_{1:k} = \{\mathcal{S}_{1:k-1}, \mathcal{S}_k\}$

 Update maximum derivative $g_{\max} = \max(g_{\max}, g_k)$

if Update time > max update time **then**

 Points in memory are thinned.

end if

until

5. Simulation study

Before deploying the AUV in the ocean we want to test different strategies in a simulated case. Here we generate a replicate study where we know the true salinity field. In this section we will go over the setup of the simulation study, the different




Figure 4: Illustration of spider leg designs over; (a) one realization of the true salinity field, (b) the conditional mean salinity field and (c) the prior mean field. All three plots show the observed path of the AUV in black. The conditional mean changes compared with the prior mean close to the observed path, and becomes closer to the true field. The conditional mean for the full field is only included for illustration here. The possible paths are shown in dark green and the best path in light green. The next waypoint is played one step-length down the best path.

metrics used for evaluating the different strategies, and the conclusions that become recommendations for the real-world setup.

5.1. Simulation setup

For the simulation study we use a setup with a square 2D field as shown in Figure 4, the field is of size $1\text{km} \times 1\text{km}$. For the simulation we need to have a true-field and a prior mean, both fields are static. The prior should capture some characteristics of the true field, but will not be completely accurate. For all the tests we will run 100 replicates, where each replicate gives a unique realization of the true field.

The starting location for the AUV is different for each replicate and will always be several step-lengths away from the interesting high gradient regions. The speed of the AUV is set to be 1 m/s with a sampling frequency of 1 Hz . For all experiments the AUV will run for a total of 5000 m . The step-length for the AUV is set to be 100 m , and the horizon is 500 m .

5.2. Evaluation metrics

The evaluation metrics look at how large the largest gradient observed is and how many important regions the AUV is able to visit.

One way to check if one strategy is better than another is to see the absolute gradient g_{\max} measured after sampling for a time t . Ideally we would want the g_{\max} to be as large as possible for each stage k . We will use this to ascertain how well a strategy is performing. We can define $g_{\max}(t)$ as

$$g_{\max}(t) = \max_{t_i < t} |g(\mathbf{s}_i, \mathbf{s}_{i-1})|. \quad (19)$$

This metric will be a good indicator of whether we have found some large gradient during the mission. But the metric will not tell us how well the AUV is able to explore gradients in different regions of the field. For this purpose we need another evaluation metric. The AUV should be able to explore several regions where the gradient is high. We refer to these regions containing large gradients by important regions. One aims to visit many important regions during the mission. We split the region \mathcal{D} into $N = 400$ equal regions, and then we look at the 20 % regions with the largest absolute gradients. In Figure 5 we illustrate the important regions (marked with red in the right plot) for




Figure 5: One realization showing absolute gradient (left plot) and the corresponding important regions (right plot). The corresponding regions are the 20 % regions with the highest gradients.

one replicate realization of the field. These important regions are different for each

replicate, but remain static for the whole simulation time. It is considered better if the AUV visits more of these regions. Hence, we construct a performance measure by counting how many of these regions the AUV visits on its exploration of the field. We count this over distance traveled (which is proportional to time).

5.3. Simulations results

We have proposed two different objective functions that can be used in looking for large gradients; PoI (14) and EI (17). In addition to these, we compare performance with 3 other strategies; one goes in the direction with the largest expected gradient, one goes to the largest variance in the derivatives and the last is a random walk. All strategies use the same algorithm as described in Algorithm 1 with the statistical model described in Section 3, the only difference is that the objective function is swapped.

The results for running simulations with these 5 different objective functions using a step-length of 100 m and a horizon of 500 m are shown in Figure 6. We first study the AUV's ability to detect large directional derivatives. In Figure 6 Left) we show the maximum derivative discovered by each of the five criteria over time. The thick line is the mean over the 100 replicates while the shaded region represents 2 standard errors in this mean. We clearly see that the PoI and EI work better than just looking for the max gradient, the reason for this is that it might get stuck in a local maxima, and it will not move away. This plot shows that EI does the best out of the five objective functions, although the gap between PoI and EI closes towards the end of the simulation. The random strategy and the one going for largest variance in the field are not performing so well, compared with the others.

In Figure 6 right) we observe how many important regions the AUV has visited after running a given distance. The thick lines indicate the mean number of regions visited for a given distance, and the shaded region represents 2 standard errors in

this mean. The ranking of the different objective functions is the same as in the left plot. The relative performances can be viewed a bit differently. Just looking for the maximum gradients does not really explore the field all that well. Rather, it can get stuck in local minima. Focusing on the maximum variance in the gradients, will eventually sample in the important regions, but it spends too much time in the exploration. Lastly, the gap between PoI and EI does not close down towards the end here. From this test, the conclusion is that EI works best overall.




Figure 6: Compare the performance of different objective functions with 100 realizations for each. (a) increase in mean g_{\max} and (b) increase in important regions visited. Both (a) and (b) show that expected improvement works best, with probability of improvement following.

To inspect further we take a look at worst-case and best-case outcomes of the simulations to see how they differ for the different objective functions. This can be important because it can be difficult to conduct many experiments in the ocean. Table 1 shows the 5th, 50th and 95th percentiles for the two metrics at the end of the simulation. For the 95th and 50th percentile both EI and PoI have a similar performance, but for the 5th percentile the difference is large. In the 5th percentile

EI is able to sample in 39 (out of 80) important regions, while PoI is only able to sample in 19. This means that EI is much better at exploring the important parts of the field in the worst cases.

Objective function	g_{\max}			Important Regions		
	5 %	50%	95 %	5 %	50%	95 %
EI	1.52	1.81	2.13	39	47	56
PoI	1.49	1.83	2.18	19	45	55
Max Gradient	0.61	1.78	2.15	0	33	47
Max Variance	0.87	1.50	1.94	2	14.5	31
Random	0.61	0.81	1.54	0	1.0	13

Table 1: Percentile table for the different objective functions with the two different evaluation metrics. These percentiles are calculated at the end of the simulations.

Regarding algorithmic parameter tuning we conducted some other tests to evaluate other aspects of the algorithm. One thing to test was how the horizon affected the performance, not surprisingly the longer the horizon the better. The main reason to limit the horizon is the computational cost for each iteration. It was also important that the step length was not too long as the AUV tends to be overstepping the phenomena in that case. It must travel the long way back, and this wastes time.

6. Case study

The suggested algorithm was tested in the Trondheim fjord on June 22. 2023. The AUV ran the adaptive mission for 2 hrs 10 min starting at 11:00 am. We first describe the parameter specification in the spatio-temporal model using numerical ocean data. We then describe the AUV setup and finally show results of the deployment.

6.1. Prior model specification based on SINMOD

We have access to a numerical ocean model for the fjord called SINMOD developed by SINTEF Ocean (Slagstad and McClimans, 2005). In our case the model simulates

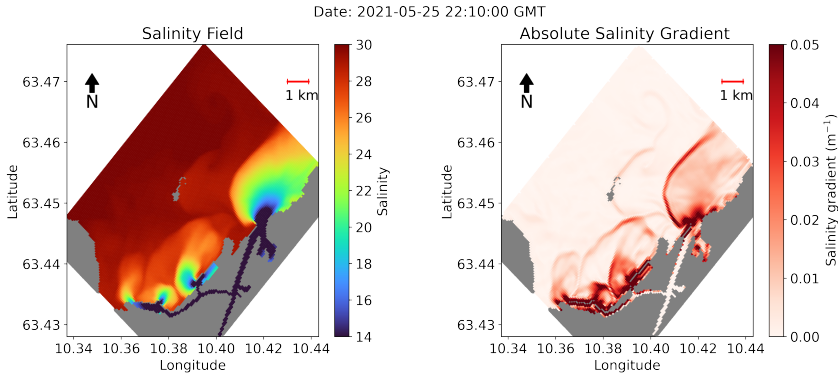


Figure 7: SINMOD simulation for the surface level of the river plume in the Trondheim fjord. Left plot shows the salinity level and the right shows the absolute salinity gradient. Left plot shows that the river has 4 outlets. The right plot shows clear river fronts in dark red. The tide is going from high tide to low tide.

several features of the fjord like currents, temperature, and salinity, but we are mostly interested in the salinity and the spatial salinity changes. Figure 7 shows a snapshot from a simulation, the left plot shows the salinity level and the right plot shows the absolute gradient. At this time in the simulation the water level goes from high tide to low tide¹ The dark red regions in the right plot show the river front. This is the region that is most interesting to sample.

These SINMOD simulations are computationally heavy to run. Therefore we use simulations done some time before the mission. We can use the outcome of the simulation as the prior mean for the surrogate GRF model of the salinity field.

The spatial and temporal covariance parameters ϕ_t and ϕ_s , and the sill σ in Equation (4) are estimated from a variogram analysis of SINMOD data. The parameters are $\phi_t = 5400s$, $\phi_s = 530m$ and $\sigma = 2$. The measurement noise for the salinity sensor is estimated from previous AUV deployments in the same location. The variance for

¹Tide data gathered from <https://www.kartverket.no/>

the difference between two measurements is

$$\text{Var}[y(\mathbf{s}_i) - y(\mathbf{s}_{i-1})] = \text{Var}[x(\mathbf{s}_i) - x(\mathbf{s}_{i-1})] + \text{Var}[\epsilon(\mathbf{s}_i) - \epsilon(\mathbf{s}_{i-1})].$$

The salinity sensor has a frequency of 1 Hz and the AUV maintains a speed of 1.6 m/s, therefore two consecutive measurements are done within 1.6m and 1s of each-other. Large depth changes are also filtered out, because that salinity change in depth is much larger. Then we assume that $\text{Cov}(x(\mathbf{s}_i), x(\mathbf{s}_{i-1})) \approx \sigma^2$, therefore $\text{Var}[x(\mathbf{s}_i) - x(\mathbf{s}_{i-1})] \approx 0$. Then

$$\text{Var}[y(\mathbf{s}_i) - y(\mathbf{s}_{i-1})] \approx \text{Var}[\epsilon(\mathbf{s}_i) - \epsilon(\mathbf{s}_{i-1})] = 2\tau^2. \quad (20)$$

We can also get that $E[y(\mathbf{s}_i) - y(\mathbf{s}_{i-1})] \approx 0$ for such close measurements. This means that we can estimate τ by using

$$\hat{\tau}^2 = \frac{1}{2(n-1)} \sum_{i=1}^{n-1} (y(\mathbf{s}_i) - y(\mathbf{s}_{i-1}))^2. \quad (21)$$

We get $\tau = 0.27$.

The algorithm contains other tuning parameters, the step-length to be 250 m, and the prediction horizon to be 1000 m. The AUV will evaluate at most 7 transects, and the maximum planning time is set at 5 s. The target depth layer for the deployment is set at 1 m.

6.2. AUV Setup

In this field experiment, a Light Autonomous Underwater Vehicle (LAUV) from NTNU's Applied Underwater Robotics Laboratory (AURLab) was employed. Pre-launch protocol consisted of standard remote control verification (Figure 8).

The primary computational unit of the LAUV is the NVIDIA Jetson TX2. The vehicle's onboard algorithmic capabilities are augmented through the integration of an adaptive sampling framework (Mo-Bjørkelund et al., 2020), which mediates message exchange between the Robot Operating Systems (ROS) (Quigley, 2009) and



Figure 8: The AUV named Thor is heading towards the river mouth area where potential high gradient might exist.

DUNE (DUNE: Unified Navigation Environment(Pinto et al., 2013)). Communication among the vehicle’s components utilizes the Inter-Module Communication (IMC) protocol (LSTS, 2022). The integration follows the scheme outlined in Ge et al. (2023), providing additional insights into the ROS-IMC bridge.

The AUV maneuvered at a depth of 1 m, where the salinity variance is large. Also, the AUV is less prone to colliding with small boats when keeping this depth, so it induces less risk. The AUV is programmed to try to maintain a speed of 1.6 m/s, which is much faster than the dynamics of the plume phenomenon observed in the SINMOD results. It was configured to re-surface at 10-minute intervals for navigational adjustments.

6.3. Results

The trajectory made by the AUV is shown in Figure 9. The left plot shows the measured salinity along the path, the middle plot shows the absolute directional

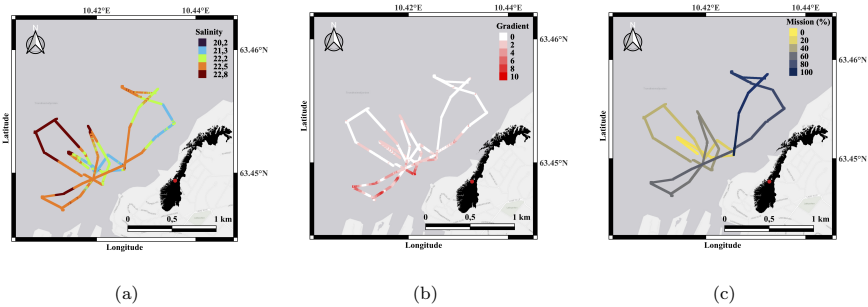


Figure 9: (a) Observed salinity along the AUV trajectory. (b) The observed absolute gradient along the trajectory. (c) The mission accomplished in time steps. The AUV takes 39 steps, each step is around 250m.

derivatives, and the right plot shows where the AUV is at any point in the mission. We notice that the AUV spends most of the time close to the river outlet, this is where the salinity change is expected to be the largest. The AUV measures the largest salinity changes in the south-west region of the map, this is around mid-way through the mission.

How the maximum gradient changes over time is shown in Figure 10. This display shows the increase in g_{\max} , the observed absolute gradient and the predicted absolute gradients. There are four distinct increases in the g_{\max} , in the start, 0.2 hours, 0.75 hours and 1.1 hours into the mission. At around 1 hour to 1.5 hours into the mission the AUV samples a region with a lot of salinity change, this is where the largest gradient is found. After this the AUV does not measure any very large salinity gradients. It is also interesting to look at what gradients the AUV predicts during the mission. The figure shows that when the AUV observes a large gradient, it also predicted a large gradient, however the prediction is often much larger. There are also some points where the model predicts a large gradient, but no large gradient is observed. This means that the model does perform one of the most important tasks which is to guide the agent towards large salinity changes




Figure 10: Increase in maximum derivative for the duration of the mission. The display also shows the measured absolute derivative during the mission and the predicted derivative.

During the mission the AUV takes 39 steps. In Figure 11 the value for the EI for each of the transects for a given step. The EI usually starts out with large values while the as g_{\max} is low, this can be seen in step 4. Towards the end the EI drops closer to zero, the AUV will start to explore different regions where it does not necessarily predict that the gradient will be large, but rather that the variance is large.

We next study how predictions for the gradients correlate with the measured gradients. For each step k we predict the salinity distribution along the path \mathcal{P}_k , but because of currents and other navigational errors the points we measure \mathcal{S}_k will not be exactly the same, and it can be a large difference. We use the model estimated in step $k - 1$ to predict the points \mathcal{S}_k that will be measured in step k . Then we look at how well the model can predict the next transect, these contain 160 - 300 data points. The results from this are shown in Figure 12. There is a correlation between the predicted gradient and the observed gradient. Ideally we would like the values to




Figure 11: Expected improvement for each step of the mission. The red line represents the highest expected improvement, and this will be the direction the AUV moves in. The blue dots are the expected improvement values for the other possible directions.

lay along the 1:1 line, but the model tends to predict larger gradients than what is observed. This might be because the variance in the salinity for the prior model is larger than what is observed, leading the model to predict larger gradients.

In summary, the AUV was able to sample and measure salinity changes in the river plume. EI worked rather well at finding large gradients and for exploring different parts of the river plume. The onboard model gives a reasonable prediction for the gradients that are going to be observed, but the gradients predicted by the model are larger than what is observed.

7. Closing remarks

We have presented an approach for constructing sampling designs by an agent moving in a spatio-temporal domain. The goal is to provide valuable designs, which in our case involves locating regions that exhibit large spatial changes. For our application in oceanography, such locations could indicate transitional zones in water masses which are potentially indicative of much biological activity. The approach for adaptive sampling is based on a Gaussian random field model, and the directional




Figure 12: Correlation between the predicted and observed salinity gradients. There is a correlation between the predicted and the observed gradients, but the model predicts larger gradients than what is observed.

changes in the field are then also Gaussian distributed. By using a Gaussian model, one facilitates efficient calculations on the agent's limited computing resources. With the sampling design setting of a moving agent, we suggest a spider leg design at each stage of the adaptive operation, and we use expected improvement in directional differences to guide the adaptive sampling. There is hence no operational grid such as a waypoint graph. Instead, the prior model assumption is effectively updated with data at each stage, limited only by the size of the data vector. In long-term operations, storage problems can occur, and we suggest to fade distant data (in space and time). This allows long-time operations in large spatial domains. We demonstrate the merits of the approach in a simulation study and in a field deployment running an autonomous underwater vehicle in a Norwegian fjord.

The field of robotics and embedded computing is growing quickly with the technological advances in small-size computing units and the current societal focus on AI. Statistics should play a substantial role in the development of new algorithms in this field. We have shown one example of added value in using spatial statistics and spatial design for underwater robotics. There are plenty of other applications where statistics can contribute. In our experience working on this, the engineers see much merit in more formalized statistical methodologies. They are however striving for efficiency and real-time operation, so rather than overly complex statistical models or methods, there seems to be a need for fast and robust systems that still have reasonably good statistical properties.

For future work, we want to investigate more nuanced algorithms where one can tune the distance and design parameters to automatically capture the right scales on the fly. We also aim to look at multivariate fields which requires a re-formulation of the derivatives used here. Rather than just derivatives, one is often interested in volumes (spatial integrals). In oceanographic applications volumes of relevance include high biomass, oxygen production, net primary production, etc., see e.g. Wu et al. (2022). Integral expressions are linear operators and it is hence Gaussian distributed if the variable of interest is Gaussian. Many of the methods described in this paper can hence be used for such applications. We used a relatively standard spatio-temporal model here. It can be extended to more complex temporal dynamics as well as non-stationary spatial elements. Staying within the Gaussian model class, we can for instance build on advection-diffusion processes (Sigrist et al., 2015; Foss et al., 2022) or use links to stochastic partial differential equations in the spatial domain (Berild and Fuglstad, 2023).

Acknowledgments

This work is funded by the Research Council of Norway, project number 305445. The authors thank Ingrid Ellingsen and SINTEF for supplying the SINMOD simulations used for parameter estimation and the prior mean. We also thank the people in the Aurlab ² for making the deployment in the fjord possible, especially Karoline Barstein for operating the AUV during the deployment.

References

- Belkin, I. M., Cornillon, P. C., and Sherman, K. (2009). Fronts in large marine ecosystems. *Progress in Oceanography*, 81(1-4):223–236.
- Berild, M. O. and Fuglstad, G.-A. (2023). Spatially varying anisotropy for gaussian random fields in three-dimensional space. *Spatial Statistics*, 55:100750.
- Brus, D. J. (2022). *Spatial sampling with R*. CRC Press.
- Catto, J. L. and Pfahl, S. (2013). The importance of fronts for extreme precipitation. *Journal of Geophysical Research: Atmospheres*, 118(19):10–791.
- Cressie, N. and Wikle, C. K. (2015). *Statistics for spatio-temporal data*. John Wiley & Sons.
- Dhingra, S., Madda, R. B., Gandomi, A. H., Patan, R., and Daneshmand, M. (2019). Internet of things mobile–air pollution monitoring system (iot-mobair). *IEEE Internet of Things Journal*, 6(3):5577–5584.
- Dunbabin, M. and Marques, L. (2012). Robots for environmental monitoring: Significant advancements and applications. *IEEE Robotics & Automation Magazine*, 19(1):24–39.

²<https://www.ntnu.edu/aur-lab>

- Fedorov, K. N. (1986). *The physical nature and structure of oceanic fronts*, volume 333. Springer.
- Fonseca, J., Bhat, S., Lock, M., Stenius, I., and Johansson, K. H. (2023). Adaptive sampling of algal blooms using autonomous underwater vehicle and satellite imagery: Experimental validation in the baltic sea. *arXiv preprint arXiv:2305.00774*.
- Foss, K. H., Berget, G. E., and Eidsvik, J. (2022). Using an autonomous underwater vehicle with onboard stochastic advection-diffusion models to map excursion sets of environmental variables. *Environmetrics*, 33(1):e2702.
- Fossum, T. O., Norgren, P., Fer, I., Nilsen, F., Koenig, Z. C., and Ludvigsen, M. (2021). Adaptive sampling of surface fronts in the arctic using an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 46(4):1155–1164.
- Fuhg, J. N., Fau, A., and Nackenhorst, U. (2021). State-of-the-art and comparative review of adaptive sampling methods for kriging. *Archives of Computational Methods in Engineering*, 28:2689–2747.
- 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.
- Gramacy, R. B. (2020). *Surrogates: Gaussian process modeling, design, and optimization for the applied sciences*. CRC press.
- Gramacy, R. B. and Apley, D. W. (2015). Local gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, 24(2):561–578.
- Hopkins, J., Challenor, P., and Shaw, A. G. (2010). A new statistical modeling approach to ocean front detection from sst satellite images. *Journal of Atmospheric and Oceanic Technology*, 27(1):173–191.

- Krause, A., Singh, A., and Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9(2):235–284.
- Lermusiaux, P. F. (2006). Uncertainty estimation and prediction for interdisciplinary ocean dynamics. *Journal of Computational Physics*, 217(1):176–199.
- LSTS (2022). Inter module communication protocol.
- Manohar, K., Brunton, B. W., Kutz, J. N., and Brunton, S. L. (2018). Data-driven sparse sensor placement for reconstruction: Demonstrating the benefits of exploiting known patterns. *IEEE Control Systems Magazine*, 38(3):63–86.
- Mateu, J. and Müller, W. G. (2012). *Spatio-temporal design: Advances in efficient data acquisition*. John Wiley & Sons.
- 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.
- Petersen, K. B., Pedersen, M. S., et al. (2008). The matrix cookbook. *Technical University of Denmark*, 7(15):510.
- Pinto, J., Dias, P. S., Martins, R., Fortuna, J., Marques, E., and Sousa, J. (2013). The lsts toolbox for networked vehicle systems. In *2013 MTS/IEEE OCEANS - Bergen*, pages 1–9.
- Quigley, M. (2009). Ros: an open-source robot operating system. In *ICRA 2009*.
- Schranz, M., Di Caro, G. A., Schmickl, T., Elmenreich, W., Arvin, F., Şekerciöğlü, A., and Sende, M. (2021). Swarm intelligence and cyber-physical systems: concepts, challenges and future trends. *Swarm and Evolutionary Computation*, 60:100762.

- Shadrin, D., Menshchikov, A., Somov, A., Bornemann, G., Hauslage, J., and Fedorov, M. (2019). Enabling precision agriculture through embedded sensing with artificial intelligence. *IEEE Transactions on Instrumentation and Measurement*, 69(7):4103–4113.
- Sigrist, F., Künsch, H. R., and Stahel, W. A. (2015). Stochastic partial differential equation based modelling of large space–time data sets. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 77(1):3–33.
- 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.
- Wang, J.-F., Stein, A., Gao, B.-B., and Ge, Y. (2012). A review of spatial sampling. *Spatial Statistics*, 2:1–14.
- Wang, Y., Le, N. D., and Zidek, J. V. (2020). Approximately optimal spatial design: How good is it? *Spatial Statistics*, 37:100409.
- Wu, J., Goes, J. I., do Rosario Gomes, H., Lee, Z., Noh, J.-H., Wei, J., Shang, Z., Salisbury, J., Mannino, A., Kim, W., Park, Y.-J., Ondrusek, M., Lance, V. P., Wang, M., and Frouin, R. (2022). Estimates of diurnal and daily net primary productivity using the geostationary ocean color imager (goci) data. *Remote Sensing of Environment*, 280:113183.
- Zhan, D. and Xing, H. (2020). Expected improvement for expensive optimization: a review. *Journal of Global Optimization*, 78(3):507–544.
- Zhang, Y., Rueda, C., Kieft, B., Ryan, J. P., Wahl, C., O’Reilly, T. C., Maughan, T., and Chavez, F. P. (2019). Autonomous tracking of an oceanic thermal front by a wave glider. *Journal of Field Robotics*, 36(5):940–954.

Paper V

Long-Horizon Informative Path Planning with Obstacles and Time Constraints

Yaolin Ge, André Julius Hovd Olaisen, Jo Eidsvik, R. Praveen Jain, Tor Arne
Johansen

Published to IFAC-PapersOnLine

Long-Horizon Informative Path Planning with Obstacles and Time Constraints

Yaolin Ge* André Julius Hovd Olaisen* Jo Eidsvik*
R. Praveen Jain** Tor Arne Johansen**

* Department of Mathematical Sciences, Norwegian University of
Science and Technology (NTNU) (email: {yaolin.ge, andre.j.h.olaisen,
jo.eidsvik}@ntnu.no.)

** Department of Engineering Cybernetics, NTNU (email:
{ravinder.p.k.jain, tor.arne.johansen}@ntnu.no.)

Abstract: We apply non-myopic informative path planning in a simulated river plume case study with several constraints on our agent. A cost valley philosophy is proposed to guide the agent through the field. The purpose of this path planner is to reveal the river plume front with the long-horizon while safely returning home in time. Among others, we employ RRT*, a variant of RRT (rapidly-exploring random trees), as the path planner to determine the least-cost path between locations. The distance budget from start to end destination, the obstacle constraint, and directional change are penalties, whereas the reduced variance of the field and an excursion set are the two rewards. The cost valley is then computed by superimposing those five fields. The simulation results demonstrate the efficiency of such a strategy. They show that the suggested approach balances exploitation and exploration while bearing in mind the go-home constraint.

Copyright © 2022 The Authors. This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Keywords: Path Planning, Autonomous underwater vehicles, Gaussian Random Field, and adaptive sampling

1. INTRODUCTION

Autonomous Underwater Vehicles (AUVs) have been used extensively for the investigation of different oceanographic phenomena (Hwang et al., 2019). AUV adaptive sampling has gained more interest in oceanographic surveying (Zhang et al., 2020; Fossum et al., 2019). Plume and ocean front investigation with AUVs has formed scenarios to validate adaptive sampling methods and algorithms (Fossum et al., 2021; Berget et al., 2018; Fossum et al., 2018). In using only one AUV to conduct the adaptive sampling, dominating methods can be grouped into either myopic (greedy) or non-myopic approaches. Myopic schemes guide the agent (AUV) towards the most informative location selected from a subset of candidate locations within the myopic neighborhood radius (Fossum et al., 2021; Berget et al., 2018). The greediness of such algorithms can make it too short-sighted, and it can fail at revealing other interesting areas. Non-myopic strategies can alleviate such challenges by expanding its horizon to a longer-stage (Bai et al., 2021). Xiao and Wachs (2022) show the effectiveness of such algorithms in a small-scale case where the shapes of the unknown objects are revealed by a robot arm, using global kriging variance reduction as the main criterion. However, the computational cost associated with such non-myopic algorithms is usually very high, and it might not apply to larger-scale onboard computation in the ocean. Suh et al. (2017) provide an idea of using cost-aware RRT* to generate sampling paths using cross-entropy as the cost function.

In our case, we want the agent to continuously determine a path for collecting valuable information from the salinity

field to reveal the river plume front, with the goal being to map unknown ocean properties while considering the remaining distance budget and avoiding the potential risks of collision with mapped obstacles. Going beyond common pre-scripted planners and greedy strategies, we propose a long-horizon path planner to solve this problem by constructing a cost valley built on multiple penalty and reward fields.

In Section 2 we describe the background and the simulation case which motivates the study of long-horizon path planning for sampling the river plume front. In Section 3 we outline our modeling approach and the statistical methods for informative sampling with operational constraints. In Section 4 we present the algorithms associated with long-horizon path planning. In Section 5 we show results from a simulation study comparing standard pre-scripted lawnmower and myopic (greedy) strategies. In Section 6 we conclude and point to future research directions.

2. PROBLEM STATEMENT

We consider the challenge of AUV adaptive sampling to map a spatial salinity field to uncover the river plume front in a specified domain. The focus is on the ability of the AUV to conduct adaptive sampling when there are constraints at the start and end point of the deployment, limited distance budget for the mission, and static obstacles in the field.

In this simulated case, we consider a two-dimensional domain. The variable of interest is salinity, which is assumed to vary smoothly in the domain. Fig. 1 illustrates a syn-




Fig. 1. Exemplary true salinity field on the unit square domain. The starting point is $(0, 0)$ and the endpoint is $(0, 1)$. The obstacle is in the center, marked by red dashed lines.

thetic case for the unit square domain. Here, the surface salinity field is generated from a Gaussian random field (GRF) model. In this salinity realization, there is a very low salinity region to the east and also some low salinity values to the north. In a practical case, such regions could represent water masses from river plumes. The red polygon inside the field shows an obstacle that might represent an island. The agent is deployed at the red dot in the south-west corner (coordinates $(0, 0)$), and one must retrieve the agent after a certain time at the black star in the north-west corner (coordinate $(0, 1)$). The perimeter of the domain is 4 units, and the distance budget for the AUV is set to be 5 units. To achieve its task, the agent needs to conduct adaptive path-planning and make sure that it is maximizing information-gathering objectives while getting back to the desired destination in time and avoiding the collision.

3. STATISTICAL MODELS AND METHODS FOR AUV SAMPLING

3.1 On-board computing with GPs

We denote the salinity variable by $\xi_{\mathbf{u}}$, with spatial variable $\mathbf{u} \in \mathcal{M} \subset \mathcal{R}^2$, where \mathcal{M} denotes the domain of interest. We assume that the salinity field is represented by a GRF. Similar assumptions of Gaussianity on salinity variables have been used in e.g. Das et al. (2013) and Binney et al. (2010). The initial specification then includes the estimation of underlying trends, variability, and spatial dependence. We assume that they can be extracted from ocean model data (Slagstad and McClimans, 2005).

The GRF modeling assumptions enable fast onboard data assimilation and AUV adaptive sampling efforts. For onboard implementation and computing, the domain is \mathcal{M} which is discretized to a set of n grid locations; $\{\mathbf{u}_1, \dots, \mathbf{u}_n\}$. The prior GRF model at these grid locations is denoted by

$$\boldsymbol{\xi} = (\xi_{\mathbf{u}_1}, \dots, \xi_{\mathbf{u}_n})^T, \quad \boldsymbol{\xi} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \quad (1)$$

with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. We assume a Matern correlation function so that $\Sigma(i, i') = \sigma^2(1 + \phi h(i, i')) \exp(-\phi h(i, i'))$, with variance σ^2 , correlation de-

cay parameter ϕ and Euclidean distance $h(i, i')$ between sites \mathbf{u}_i and $\mathbf{u}_{i'}$.

The measurement y_j at each stage $j = 1, \dots, N_{\text{steps}}$ is modeled by

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

where the vector \mathbf{f}_j defines the sampling indices at this stage of operation and r is the salinity measurement noise standard deviation. The sampling design \mathbf{D}_j at this stage j , say directions north, east, west, or south, determines the 0 and 1 structure in vector \mathbf{f}_j because it directly defines the measurement location. This aspect will be important for design evaluation in what follows.

Starting with $\mathbf{m}_0 = \boldsymbol{\mu}$ and $\mathbf{S}_0 = \boldsymbol{\Sigma}$, Bayes' rule is used to achieve data assimilation at stages $j = 1, \dots, N_{\text{steps}}$. This gives the updated Gaussian model with mean and variance given by

$$\begin{aligned} \mathbf{G}_j &= \mathbf{S}_{j-1} \mathbf{f}_j (\mathbf{f}_j^T \mathbf{S}_{j-1} \mathbf{f}_j + r^2)^{-1} \\ \mathbf{m}_j &= \mathbf{m}_{j-1} + \mathbf{G}_j (y_j - \mathbf{f}_j^T \mathbf{m}_{j-1}) \\ \mathbf{S}_j &= \mathbf{S}_{j-1} - \mathbf{G}_j \mathbf{f}_j^T \mathbf{S}_{j-1}. \end{aligned} \quad (3)$$

3.2 Information criterion for sampling

Based on our problem statement in Section 2, we have chosen to use numerous information criteria in the objective function which determines the AUV sampling design. The overall function is defined by a sum of normalized versions of the following criteria that we describe next:

Integrated variance reduction (IVR) uses the latter part of the posterior covariance calculation in (3). As an information criterion, the goal now is to provide maximum reduction of the marginal variances at all spatial locations in the grid, see also Binney et al. (2010) and Fossum et al. (2018). For a particular AUV sampling design \mathbf{D}_j , defined via the sampling design vector \mathbf{f}_j , the variance reductions at this stage are given by the diagonal entries of $\mathbf{R}_j = \mathbf{S}_{j-1} \mathbf{f}_j (\mathbf{f}_j^T \mathbf{S}_{j-1} \mathbf{f}_j + r^2)^{-1} \mathbf{f}_j^T \mathbf{S}_{j-1}$. The sum of these represents an information measure.

Expected integrated Bernoulli variance reduction (EIBV) uses the classification of salinity above and below a threshold t according to an excursion set $\text{ES} = \{\mathbf{u} \in \mathcal{M} : \xi_{\mathbf{u}} \leq t\}$, and the goal is to increase the classification accuracy of salinity levels according to this threshold. The Bernoulli variance (BV) at location \mathbf{u} is $p_{\mathbf{u}}(1 - p_{\mathbf{u}})$, $p_{\mathbf{u}} = P(\xi_{\mathbf{u}} \leq t)$, which we aim to reduce by the data gathering. The goal of a sampling design for data y_j is then to minimize the expected spatially integrated Bernoulli variance, and for the Gaussian model, there is a closed-form solution (Fossum et al., 2021).

Obstacle avoidance ensures that the AUV does not crash into land or islands. There is an infinity cost penalty if the AUV runs into an obstacle. In practice, a finite cost penalty can be applied to account for uncertainty.

Directional change sets a penalty for sharp AUV turns. Beyond the 90° limit, there is an increasing penalty for high turning angles, so only very rarely (in situations with conflicting objectives), will we see abrupt angle changes in the AUV path.

Budget limitations include constraints that make sure that the AUV gets to its destination in time. The budget penalty will start when the mission approaches the end, and form an elliptical region away from the current AUV location towards its destination. There is infinity loss outside this region. The cost can be finite in practice as mentioned above.

The objective function guiding the AUV sampling is composed of all these measures. In doing so, we aim to balance exploration for uncertainty and salinity boundaries with operational constraints for the vehicle. At each stage $j = 1, \dots, N_{\text{Steps}}$, the AUV updates its calculation of a *Cost Valley* based on evaluating all these criteria for feasible designs. The results are used to compare designs and for selecting the optimal sampling design D_j at the current stage. The procedures involving model updating, design criterion calculation, and data gathering are summarized in Algorithm 1.

Algorithm 1 Informative sampling algorithm

Require: Initial mean m_0 and covariance S_0
 Set start waypoint $D_1 = \{u_1\}$
 $\mathcal{Y}_0 = \emptyset$
 $j = 1$
while $j \leq N_{\text{Steps}}$ **do**
 Act :
 Go to waypoint D_j .
 Sense :
 Gather data y_j . $\mathcal{Y}_j = (\mathcal{Y}_{j-1}, y_j)$.
 $G_j = S_{j-1} f_j (f_j^T S_{j-1} f_j + R_j)^{-1}$
 $m_j = m_{j-1} + G_j (y_j - f_j^T m_{j-1})$
 $S_j = S_{j-1} - G_j f_j^T S_{j-1}$
 $j = j + 1$
 Plan :
 Budget = Budget - $\|D_j - D_{j-1}\|_2$
 $\mathbf{CV} = \text{updateCostValley}(m_j, S_j, \text{Budget}, u_j, u_{j-1})$
 $u_j = \text{argmin}_{u \in \mathcal{M}}(\mathbf{CV})$
 $T_j = \text{RRT}^*(\mathbf{CV}, D_{j-1}, u_j)$
 $D_j = T_{j1}$
end while

4. PATH PLANNER

We design a path planning algorithm to find the cheapest path through a cost valley. This path should take the AUV from the current location to the lowest point in the cost valley. The cost valley is defined by combining information criteria, direction criteria, and budget limitations. The cost valley is updated with data as the AUV gathers salinity information from the field. The cost valley philosophy relies on a long-horizon path plan, which is not only considering the next step but anticipating future steps.

We build such a cost valley based on the overlay of five different information criteria. The weights among them are equal, but before merging them, all the costs are normalized to be within range $[0, 1]$, except those costs which are ∞ . The RRT* path planner, see e.g. Karaman and Frazzoli (2011) and Hollinger and Sukhatme (2014), is used to determine the optimal least-cost path from the current AUV location to the end destination. From

this calculation, the agent selects the next optimal design location. It measures the salinity y_j at this design location and the entire GRF model is updated. This means that the cost valley is also updated, and hence the new starting location will be fed to the RRT* path planner.

The building blocks of the core algorithm are presented in Algorithm 2-8.

Algorithm 2 RRT*, called by Algorithm 1

Require: $\mathbf{CV}, D_{j-1}, u_j$
for $k \in 1 \dots K$ **do**
 Generate random location u_{rand} within constraints.
 $u_{\text{nearest}} \leftarrow \text{Nearest}(G = (V, E), u_{\text{rand}})$
 $u_{\text{new}} \leftarrow \text{Steer}(u_{\text{nearest}}, u_{\text{rand}})$
 if $\text{ObstacleFree}(u_{\text{nearest}}, u_{\text{new}})$ **then**
 $U_{\text{near}} \leftarrow \text{Near}(G = (V, E), u_{\text{new}}, R)$
 Add new node: $V \rightarrow V \cup \{u_{\text{new}}\}$
 $u_{\text{min}} \leftarrow u_{\text{nearest}}$
 $c_{\text{min}} \leftarrow \text{getPathCost}(\mathbf{CV}, u_{\text{min}}, u_{\text{new}})$
 end if
 for $u_{\text{near}} \in U_{\text{near}}$ **do**
 $c_{\text{near}} \leftarrow \text{getPathCost}(\mathbf{CV}, u_{\text{near}}, u_{\text{new}})$
 if $c_{\text{near}} < c_{\text{min}}$ **then**
 $u_{\text{min}} \leftarrow u_{\text{near}}$
 $c_{\text{min}} \leftarrow c_{\text{near}}$
 end if
 end for
 Add new edge: $E \leftarrow E \cup \{(u_{\text{min}}, u_{\text{new}})\}$
 for $u_{\text{near}} \in U_{\text{near}}$ **do**
 $c_{\text{temp}} \leftarrow \text{getPathCost}(\mathbf{CV}, u_{\text{new}}, u_{\text{near}})$
 if $c_{\text{temp}} < c_{\text{near}}$ **then**
 $c_{\text{near}} \leftarrow c_{\text{temp}}$
 $u_{\text{parent}} \leftarrow \text{Parent}(u_{\text{near}})$
 $E \leftarrow (E \setminus \{(u_{\text{parent}}, u_{\text{near}})\}) \cup \{(u_{\text{new}}, u_{\text{near}})\}$
 end if
 end for
 if isArrived **then**
 $\text{Parent}(u_j) \leftarrow u_{\text{new}}$
 end if
end for
 $T = []$
while $\text{Parent}(u_k) \neq \emptyset$ **do**
 $T.append(\text{Parent}(u_k))$
 $u_k = \text{Parent}(u_k)$
end while
return T

Algorithm 3 getPathCost, called by Algorithm 2

Require: \mathbf{CV}, u_1, u_2
 $\text{Cost}_1 = \mathbf{CV}(u_1)$
 $\text{Cost}_2 = \mathbf{CV}(u_2)$
 $\text{Cost}_{\text{path}} = \text{Cost}(u_1) + \|u_1 - u_2\|_2 + (\text{Cost}_1 + \text{Cost}_2) / 2 \cdot \|u_1 - u_2\|_2$

5. SIMULATION RESULTS

To present the performance of the algorithm, we discuss the step-wise behavior of the algorithm to better demonstrate the capability of achieving the information-gathering goal while keeping an eye on the remaining distance budget and avoiding obstacles throughout the entire process.

Algorithm 4 updateCostValley, called by Algorithm 1

Require: $m_j, S_j, \text{Budget}, u_j, u_{j-1}$
 $\text{Cost}_{\text{EI}} = \text{getEIField}(m_j, S_j)$
 $\text{Cost}_{\text{Budget}} = \text{getBudgetField}(\text{Budget}, u_j)$
 $\text{Cost}_{\text{Obstacle}} = \text{getObstacleField}$
 $\text{Cost}_{\text{Direction}} = \text{getDirectionalField}(u_j, u_{j-1})$
 $\text{Cost}_{\text{Valley}} = \sum \text{Cost}_{\text{EV}, \text{Budget}, \text{Obstacle}, \text{Direction}}$
return CostValley

Algorithm 5 getEIField, called by Algorithm 4

Require: m_j, S_j
 $\text{EIBV} = 0^{n \times 1}$
 $\text{IVR} = 0^{n \times 1}$
for $i \in 1 \dots n$ **do**
 $f_j = 0^{n \times 1}$, and $f_j[i] = 1$
 $R_j = S_{j-1} f_j (f_j^T S_{j-1} f_j + r^2)^{-1} f_j^T S_{j-1}$
 $\text{IVR}[i] = \sum_{i=1}^n \text{diag}(R_j)$
 $\text{EIBV}[i] = \sum_{i=1}^n \Phi_2 \left(\begin{bmatrix} t \\ -t \end{bmatrix}; \begin{bmatrix} m_{j-1}(i) \\ -m_{j-1}(i) \end{bmatrix}, W_j(i, i) \right)$,
 where, $W_j(i, i) = \begin{bmatrix} T(i, i) & -R_j(i, i) \\ -R_j(i, i) & T(i, i) \end{bmatrix}$
 given, $T(i, i) = S_{j-1}(i, i) + R_j(i, i)$
end for
 $\text{Cost}_{\text{EI}} = \text{norm}(\text{EIBV}) + 1 - \text{norm}(\text{IVR})$
return Cost_{EI}

Algorithm 6 getBudgetField, called by Algorithm 4

Require: Budget, u_j
 Form a budget ellipse with a, b, c
 $a = \text{Budget}/2$
 $c = \|\mathbf{u}_{\text{goal}} - \mathbf{u}_j\|_2/2$
 $b = \sqrt{a^2 - c^2}$
 $\text{Cost}_{\text{Budget}} = \infty^{n \times 1}$
for $i \in 1 \dots n$ **do**
 $\omega = \frac{\mathbf{u}_{ix}^2}{a^2} + \frac{\mathbf{u}_{iy}^2}{b^2}$
if $\omega < 1$ **then** Cost_{Budget}[i] = 0
end if
end for
return Cost_{Budget}

Algorithm 7 getObstacleField, called by Algorithm 4

$\text{Cost}_{\text{obstacle}} = \infty^{n \times 1}$
for $i \in 1 \dots n$ **do**
if $u_i \in \text{Cfree}$ **then**
 $\text{Cost}_{\text{obstacle}}[i] = 0$
end if
end for
return Cost_{obstacle}

We first show the results of the algorithm for one realized salinity field. We then study the performance of the algorithm over replicate data from the GRF model, where we also compare results with that of a lawnmower algorithm and a myopic exploration algorithm with time operation constraints.

5.1 Simulation setup

The prior mean is produced by Equation (4), with a lower expected salinity near location $\mathbf{u} = (u_1, u_2) = (1, 0.5)$

Algorithm 8 getDirectionalField, called by Algorithm 4

Require: u_{j-1}, u_j
 $\mathbf{b}_1 = u_j - u_{j-1}$
 $\text{Cost}_{\text{direction}} = 10^{n \times 1}$
for $i \in 1 \dots n$ **do**
 $\mathbf{b}_2 = u_i - u_j$
if $\mathbf{b}_1 \cdot \mathbf{b}_2 \geq 0$ **then**
 $\text{Cost}_{\text{direction}}[i] = 0$
end if
end for
return Cost_{direction}




Fig. 2. The prior mean of the salinity at the initial stage.

and visualized in Fig. 2. At each step, the cost valley will be reconstructed based on the overlay of five cost fields including the EIBV cost field, IVR cost field, directional-changing cost field, budget cost field, and obstacle cost field. Then a new tree will be built based on the new cost valley. The ending location will be the place where it has a minimum value in the cost valley. Then the path planner will decide where to go next.

$$\mu_0(\mathbf{u}) = 31 - \exp\left(-\frac{(u_1 - 1)^2 + (u_2 - 0.5)^2}{0.07}\right) \quad (4)$$

5.2 Simulation results

In this part, the ground truth is as shown in Fig. 1. Results at several stages of the AUV operation are displayed in Fig. 3 ~ Fig. 5. In each figure, four illustrations are showing various parts of the cost valley calculation. The conditional mean in salinity is shown to the left (in [mg/g]). The normalized cost field based on EIBV and the normalized cost field based on IVR are in the middle. The cost valley with random tree paths for the RRT* algorithm is shown to the right. Over the different data gathering stages $j = 1, \dots, 55$, the sequence of figures shows the balance between different cost fields contributing to the cost valley calculations. The EIBV and IVR cost fields are important when there are no obstacles nearby or still much budget left. EIBV plays a more important role than IVR sometimes when it is more important to exploit, whereas IVR plays a more important role than EIBV when it is more important to explore. When the budget is running out, the other parts of the cost valley calculation are undoubtedly dominating, since it has ∞ penalty outside

the budget region. The directional change penalty plays a continuous role in path planning since it guides the agent to move forward with a smooth path, such path planning can avoid sharp turns in practice and hence reduce navigational inaccuracy.

5.3 Replicate study

To remove random effects, 50 replicated simulation results are averaged and shown in Fig. 6. During each iteration, IBV, RMSE, and EV (Expected Variance) are monitored for the comparison of the three strategies including myopic(greedy), RRT*, and pre-scripted lawnmower. Myopic strategy and the RRT* both choose the next candidate locations based on the cost associated with them from the cost valley, whereas the pre-scripted lawnmower just moves along according to its pre-designed paths. The result shows that the RRT* planner outperforms the other two, and all the indicators for RRT* including IBV, RMSE, and EV decrease fastest among the others.

6. CONCLUSION

In our simulated case, the agent can explore the field adaptively, and it achieves low spatial variance and precise river plume water classification while avoiding the island obstacle and reaching the destination in time. Via comparison with existing approaches, we learned that the suggested approach has good-quality performance metrics and it satisfies our goal of achieving exploratory path planning with a constraint.

In the future, we aim to conduct similar experiments in the field. This entails fine-tuning the non-myopic path planning strategies requiring faster computation and cautious implementation of such algorithms with a need to be well designed onboard the AUV.

ACKNOWLEDGEMENT

We acknowledge support from the Norwegian Research Council (RCN) through the MASCOT project 305445 and NTNU AMOS grant No. 223254.

REFERENCES

- Bai, S., Shan, T., Chen, F., Liu, L., and Englot, B. (2021). Information-driven path planning. *Current Robotics Reports*, 2(2), 177–188. doi:10.1007/s43154-021-00045-6.
- 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. doi:https://doi.org/10.1016/j.ifacol.2018.09.509. 11th IFAC Conference on Control Applications in Marine Systems, Robotics, and Vehicles CAMS 2018.
- Binney, J., Krause, A., and Sukhatme, G.S. (2010). Informative path planning for an autonomous underwater vehicle. In *2010 IEEE International Conference on Robotics and Automation*, 4791–4796. IEEE.
- Das, J., Harvey, J., Py, F., Vathsangam, H., Graham, R., Rajan, K., and Sukhatme, G.S. (2013). Hierarchical probabilistic regression for auv-based adaptive sampling of marine phenomena. In *2013 IEEE International Conference on Robotics and Automation*, 5571–5578. IEEE.
- 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. doi:10.1126/scirobotics.aav3041.
- 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. doi:https://doi.org/10.1002/rob.21805.
- 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. doi:10.1214/21-AOAS1451.
- Hollinger, G.A. and Sukhatme, G.S. (2014). Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9), 1271–1287.
- Hwang, J., Bose, N., and Fan, S. (2019). Auv adaptive sampling methods: A review. *Applied Sciences*, 9(15), 3145.
- Karaman, S. and Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7), 846–894.
- 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.
- 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. doi:10.1109/TRO.2017.2738664.
- 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. doi:10.1109/LRA.2022.3141458.
- 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. doi:10.1109/JOE.2019.2920217.




Fig. 3. Step 1, at first, the agent believes the northeast side has very interesting information and hence plans a long almost straight path towards the hotspot. At this step, both EIBV and IVR cost fields play important roles to guide the agent. The trees are distributed in a way that the major stems tend to align along with the low-cost areas.




Fig. 4. Step 43, at this stage, the agent moves towards the home direction while still trying to collect as much information as it can. So it tries to reach the boundary since it still has a low cost. The red ellipse shows the remaining distance budget. The penalty outside this budget is infinity, which induces the chaotic behavior of the trees, whereas the trees inside the boundary remain optimally distributed.




Fig. 5. Step 55, the agent finally decides to go home, and there appears to have a sharp turn in the trajectory since it has to prioritize the mission of going home rather than getting a high penalty from the directional restrictions.




Fig. 6. Simulation results compare myopic, RRT*, and pre-scripted lawnmower strategies in terms of IBV, RMSE, and EV. RRT* performs better than the other two in all of those metrics.