

Spatio-Temporal Sampling Designs for Autonomous Marine Vehicles

André Julius Hovd Olaisen

June 2026

Abstract

This thesis aims to improve the ability of autonomous marine vehicles to monitor important ocean processes. Ocean monitoring is increasingly important due to climate change, anthropogenic pressure, and biodiversity loss. At the same time, the number of capable autonomous marine vehicles is growing, offering new opportunities to monitor marine environments more effectively. These marine vehicles are equipped with a variety of sensors, navigation capabilities, and onboard computational facilities. By embedding onboard statistical models and decision-making algorithms, this thesis develops intelligent agents that adapt to data and decide where and when to sample. The core contribution of this thesis is the design, implementation, and evaluation of these agents.

This thesis investigates two ocean phenomena: (1) oceanic frontal systems, particularly river plumes, and (2) biomass hotspots composed of phytoplankton and/or zooplankton. Both phenomena are represented using spatio-temporal statistical models that assimilate data, incorporate prior knowledge, produce predictions, and quantify uncertainty. For this purpose, this thesis uses Gaussian random fields (GRFs) and latent GRFs. Furthermore, this thesis formalizes sampling objectives mathematically; these can include targeting large gradients in an ocean front or locating high concentrations of plankton. To achieve these goals, this thesis evaluates sequential sampling designs using acquisition functions such as expected improvement (EI) and expected integrated Bernoulli variance (EIBV) to select sampling designs that best advance the sampling objectives. Through simulation studies, this thesis identifies data-driven agents that account for the practical limitations of real autonomous marine vehicles. Lastly, these methods are implemented on autonomous marine vehicles, typically autonomous underwater vehicles (AUVs), and are tested in field experiments.

This text is divided into two sections. The first section provides background information that is relevant to the research contribution. The second section presents the research contributions themselves, consisting of several papers. Each paper examines a specific phenomenon and then develops agents that are evaluated through simulations and field experiments.

Acknowledgments

I would like to extend my greatest gratitude to my main supervisor, Jo Eidsvik. You have been extremely supportive, helpful, and insightful during the entire PhD. You have been a joy to work with. I also want to thank my co-supervisor, Martin Ludvigsen. We started working together a bit later in the PhD, but you have been very helpful in opening doors and fostering new collaborations.

Further, I would like to thank my fellow PhD researchers in the MASCOT project, Yaolin Ge and Matin Outzen Berild. I would also like to thank Geir-Arne Fuglstad and Ingrid Ellingsen, who worked on the project.

During my time here, I have worked with AurLab, which has been an excellent partner, and support was always available when needed. Here, I would like to thank my fellow PhD researchers: Tore Mo-Bjørklund, Karoline Barstein, Ahmed Abdelgayed, and Erik Liu. In addition, I would like to thank many people who helped make the field experiments possible: Pedro De La Torre, Martin Bredesen, Sebastian Sikora, Kay Arne Skarpnes, and Halvar Gravråkmø. Through the lab, I have also been in contact with people working at SINTEF; here, I would especially like to thank Ralph Stevenson-Jones and Muriel Dunn.

I have had several research visits to the LSTS lab in Porto. Here, I have received great help from many people, especially Renato Mendes, João Sousa, and Paulo Dias.

During my time working here, my professional home has been the Department of Mathematical Sciences, and in particular, the statistics division. I am grateful to everyone who has contributed to making it such an enjoyable and welcoming workplace. I especially want to thank my office mates: Hajg Jasa, Kenneth Aase, Martin Ludvigsen (different Martin), Andrea Leone, Samir Aghayev, Temesgen Melese, and Håvard Utne Terland.

I would also like to thank my family: my parents, Rut and Vidar; my siblings, Johanne, Emil, and Elise; and my aunt Karin and cousin Matias for their constant support. Lastly, but definitely not least, I would like to thank my girlfriend Åshild, who has been incredibly supportive and patient over the years.

Contents

1	Introduction	7
1.1	Projects	8
1.2	Thesis Outline	8
2	Background	10
2.1	Ocean phenomena of interest	10
2.1.1	Ocean fronts and river plumes	11
2.1.2	Plankton observation	14
2.2	Ocean observing platforms	16
3	Modeling and observations	19
3.1	Basics of Gaussian Random Fields	19
3.2	Conditioning on data	22
3.3	Computational considerations	24
3.4	Parameter estimation	27
4	Path planning and acquisition function	37
4.1	Sense-Plan-Act	37
4.2	What is interesting? Where to sample?	38
4.2.1	Finding maximum values	39
4.2.2	Classifying water-masses	44
4.2.3	Sampling high gradients	48
4.3	Path planning	49
5	Experiments	54
5.1	Simulation study	54
5.2	Field experiment	56
5.3	Experiments timeline	57
6	Paper overview	59
6.1	Paper A	59
6.2	Paper B	61
6.3	Paper C	62
6.4	Paper D	63

7	Conclusion and Further Work	64
7.1	Closing remarks	64
7.2	Further work	65

List of Abbreviations

ASV autonomous surface vehicle.

AUV autonomous underwater vehicle.

CDOM chromophoric dissolved organic matter.

CTD conductivity–temperature–depth.

DVM diel vertical migration.

EIBV expected integrated Bernoulli variance.

HAB harmful algal bloom.

IBV integrated Bernoulli variance.

IVR integrated variance reduction.

LAUV light autonomous underwater vehicle.

MCMC Markov chain Monte Carlo.

RMSE root mean square error.

ROV remotely operated vehicle.

RRT* goal oriented rapidly exploring random tree.

UAV unmanned aerial vehicle.

VR variance reduction.

Chapter 1

Introduction

I never see the truth as a fixed
star on the horizon but always as
an activity, a search, an
approximation.

Werner Herzog
Every Man for Himself and God
Against All: A Memoir

The ocean is essential for life on Earth. It covers about 70% of the planet's surface and plays a key role in regulating the climate. Approximately half of the Earth's oxygen is generated in the ocean (NOAA, 2026a), and the ocean absorbs around one third of anthropogenic CO₂ emissions (IPCC, 2013). This only begins to illustrate why preserving healthy oceans is crucial for all life on Earth. To achieve this, we must monitor the ocean effectively.

Humans have a long history of surveying the ocean, but the emergence of oceanography as a formal scientific discipline is commonly traced to the HMS Challenger expedition of 1872–1876 (Yonge, 1972). This expedition is frequently regarded as the first voyage conducted with explicitly scientific objectives, yielding foundational discoveries in oceanography. Since that time, ocean observation capabilities have expanded substantially: contemporary observing systems include research vessels, moored and drifting buoys, satellite remote sensing platforms, and a variety of in situ instrument arrays.

In recent decades, autonomous marine vehicles have played an increasingly prominent role in ocean observation. Equipped with onboard navigation, sensing, and computational resources for real-time decision-making, these platforms can execute surveys with minimal or no human oversight. One common operational paradigm is to predefine a survey plan that the vehicle subsequently attempts to follow. An alternative, more adaptive paradigm is to permit the vehicle to modify its sampling strategy in real time in response to the stream of observational data. Realizing this capability requires onboard ocean models

and principled algorithms for sequential decision-making.

This thesis seeks to integrate biological oceanography, robotics, and statistical modeling to enhance the efficiency and effectiveness of autonomous ocean sampling. The proposed sampling design algorithms are driven by concrete scientific objectives, including, for example, the characterization of ocean fronts and the detection of regions with high plankton abundance. At the same time, the constraints and capabilities of autonomous robotic platforms fundamentally shape which sampling designs are practically implementable.

1.1 Projects

Throughout my PhD, I have been funded by the Norwegian Research Council (RCN) through the Maritime Autonomous Sampling and Control (MAS-COT) project 305445. The goal of this project is to use statistical models and decision-making algorithms on autonomous ocean robots for sampling various ocean variables. A big focus has been on developing spatio-temporal models that are useful and can run on edge in autonomous vehicles. The models are often informed by physics-based ocean models. The statistical models are then coupled with path planning algorithms that enable informative path planning. Importance has been placed on simulation based testing and subsequently running field experiments. The main problem investigated was river plumes, which resulted in several field campaigns in rivers in Trondheim, Norway, and Porto, Portugal. Here we used AUVs to investigate salinity and temperature fronts arising from the river water mixing with the ocean.

During these projects, there have been continuous collaborations between the Department of Mathematical Sciences and the Applied Underwater Robotics Laboratory (AURLab) in the Department of Marine Technology at NTNU. There has also been collaboration with the research institute SINTEF Ocean and the marine robotics lab Laboratório de Sistemas e Tecnologia Subaquática (LSTS) in Porto.

In addition, I have been involved in the NRC SFI Harvest project 309661, which is a collaboration between several departments at NTNU, research institutes, and industrial partners. Specifically, I was involved in two field campaigns in Frohavet, outside of the Trondheimsfjord. The field campaigns focused on finding and measuring high concentrations of the zooplankton species *Calanus finmarchicus*. This was done using several autonomous surface vehicles and autonomous underwater vehicles. These vehicles are equipped with heterogeneous sensors that enable the sampling of different variables in different ways.

1.2 Thesis Outline

The thesis is split into two parts; the first part is an introduction to the problems and subsequent techniques to solve them, and the second part consists of the scientific publications.

In Chapter 2, there is a brief introduction to the problems; studying river plumes and evaluating the abundance of plankton. Different methods and sensors that can investigate the phenomena are described. The ocean observation platforms are introduced, specifically autonomous ocean vehicles, which are the primary platform for the thesis. The main research question is how to design algorithms that can most effectively investigate the phenomena of interest. In Chapter 3, we introduce the statistical models used for modeling different ocean variables. These models are implemented on autonomous vehicles, with an emphasis on being easily adaptable and providing quick computations. In Chapter 4, the focus shifts towards algorithms that enable autonomous agents to collect the most valuable information. This includes describing several acquisition functions for different applications. It also includes how these acquisition functions are used to enable informative path planning. In Chapter 5, we outline how simulations are used to determine which acquisition functions and path planning algorithms are best for different scenarios. Then, the best algorithms are applied in the field; here, there is a list of these field experiments. In Chapter 6, there is an overview of the papers that are listed; the papers are presented in full in the second part of the thesis. In Chapter 7, a short summary of the main contributions to the thesis is provided, along with some possible improvements and new directions for the work.

Chapter 2

Background

“Overhead the albatross hangs
motionless upon the air
And deep beneath the rolling
waves in labyrinths of coral caves
The echo of a distant time comes
willowing across the sand
And everything is green and
submarine”

Pink Floyd
Echoes

In this chapter, background information is provided on the two primary problems addressed in this thesis: detecting fronts and classifying water masses, and identifying regions with high concentrations of biomass in the form of phytoplankton or zooplankton. For this, there will be short explanations of the different scientific phenomena, as well as the techniques and sensors used to measure them. There are significant variations in terms of scale of detail, cost, spatio-temporal resolution, and many other features; therefore, we will try to highlight some of the advantages of different methods as well as some drawbacks. Lastly, the ocean platforms used to collect the data are discussed.

2.1 Ocean phenomena of interest

We aim to investigate biological-oceanographical phenomena that occur in the upper water column. This part of the ocean is often very dynamic, with currents, mixing zones, and an abundance of marine life (Lalli and Parsons, 1997). Due to these conditions, the areas we are interested in monitoring are constantly shifting and changing. Therefore, we need a host of sensors and platforms to accurately describe the phenomena of interest. Here, we aim to give a short introduction to the phenomena in conjunction with how we observe and investigate them.

2.1.1 Ocean fronts and river plumes


From a distance, the ocean looks like a fairly uniform surface, and in many places, it is. One can travel long distances while experiencing little change in parameters like salinity and temperature. But there are some places where this is not the case; there are locations where the properties of the water change sharply over relatively short horizontal distances. These places are commonly called ocean fronts. In oceanography, a front is typically defined as a region or boundary where two different water masses meet (Ferrari, 2011). Fronts are characterized by a large horizontal variability in an ocean variable, often salinity (haline fronts) or temperature (thermal fronts), but they can also involve other physical or biochemical properties.

Fronts appear on several different scales, from planetary frontal systems to smaller river fronts. The mixing zones arising from these fronts produce some of the most productive regions in the ocean and are "hot spots" for marine life (Belkin et al., 2009). Reasons why these regions can be productive include the transportation of nutrients to the sun rich surface; they can also act as accumulation zones (Olson et al., 1994). With these important properties and complex behaviors, these phenomena are essential to investigate.


A specific type of front arises in the estuarine environment where a river flows into the ocean or other bodies of water. When this occurs, a front forms between the fresh river water and the saline ocean water; this is the river plume (Lalli and Parsons, 1997). The fresher river water has a lower density than the saline ocean water; thus, river plumes are usually on the surface and can cover large areas in the ocean (Kang et al., 2013). The shape and characteristics of the plume are influenced by many factors: the river flow rate, ocean currents, salinity, temperature, wind, tides, seabed and outlet topography, and other processes (Horner-Devine et al., 2015). In Fig. 2.1, we see satellite images of two river plumes, the Douro in Portugal and the Gaula in Norway. While these images mostly show suspended sediment, the plume shapes clearly differ between the two rivers.

It is not only temperature and salinity that distinguish the water masses; the nutrient load can also differ. Rivers transport nutrients from agricultural land and other sources, which are then mixed into the coastal ocean, affecting the marine ecosystem (Han et al., 2012). In addition, river flow can induce upwelling that brings nutrient-rich water to the surface. In some cases, the resulting high nutrient levels can trigger harmful algal blooms (HAB) (Wang et al., 2019). This is one of the reasons why river plumes are important to investigate.

There are many ways to investigate river fronts; these include, among others, using physics based simulations, remote sensing, and in-situ measurements. Using fluid dynamic principles, seabed topography, and often other inputs like daily water discharge, wind, and tide, one can construct numerical solutions of complex differential equations to simulate several characteristics of the plume. There are several models that can do this, for example, (Mendes et al., 2016) used numerical models to investigate interactions between Douro and Minho



(a) Douro River plume (Portugal).



(b) Gaula River plume (Norway).

Fig. 2.1: True-colour Sentinel-2 images of river plumes: (a) Douro River, Portugal (24 March 2026); (b) Gaula River, Norway (28 March 2026). Images courtesy of the E.U. Copernicus Marine Service.

river plumes. For the Nidelva River, the numerical model SINMOD (Slagstad and McClimans, 2005) can be applied to simulate several variables; Fig. 2.2 illustrates the spatial distribution of surface salinity. While numerical models are important, they need to be validated and can often be improved by using other sensor types.

Remote sensing can often be used to monitor the characteristics of river plumes. Satellites use optical spectrometers to measure reflected light from the ocean. The reflective properties vary with the conditions in the ocean. A common objective is to investigate spatio-temporal variability (see, e.g., Thomas and Weatherbee (2006), Lihan et al. (2008)), which can be used to validate results from numerical simulations (Mendes et al., 2016). This can be done in

several ways; for instance, the suspended sediment in the plume was investigated using satellite images by Ody et al. (2022). Zhu et al. (2011) estimated chromophoric dissolved organic matter (CDOM) in the Mississippi and Atchafalaya river plume regions, and chlorophyll concentration variability was investigated by Mustapha (2011). A key strength of using satellites is that one can monitor

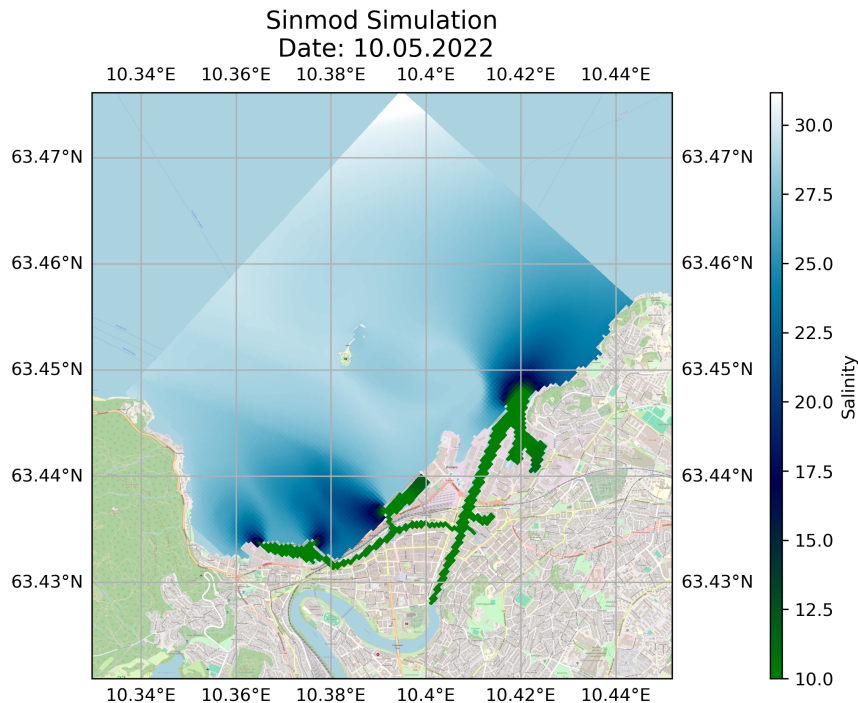


Fig. 2.2: Surface salinity from the SINMOD simulation of the Nidelva river plume.

phenomena over a long duration and at a large scale. However, there are also some weaknesses. The first issue is that images are often not available; cloud coverage and orbits mean that data is not always accessible. The second issue is that the satellite relies on optical sensors, which cannot detect every characteristic; they can only observe the surface, while many important processes are subsurface (Klemas and Yan, 2014). In addition, the estimates from the satellites should be corrected and verified by measurements made in-situ and using water samples.

In-situ measurements can be done by several ocean going platforms with different sensors, depending on the application. The different types of platforms will be discussed a bit later, but they can provide both surface measurements and subsurface measurements. There are many suitable and valuable sensors for investigating river plumes; these include salinity, temperature, turbidity, and

chlorophyll *a* (chl *a*). However, they do have some weaknesses compared with satellites; an ocean platform can only be in one place at a time, and moving them around is often not possible, slow, or dependent on currents.

The advantages and limitations of various approaches to observing and studying river plumes are also relevant to many other oceanographic questions. The best way to alleviate the weaknesses of one type of observation is to combine it with others. When many different observation platforms for different scales observe different physical and biochemical properties, we get the best possible understanding of the phenomena in question; this approach can be called an observational pyramid (SmallSat Lab, 2026).

2.1.2 Plankton observation

The ocean is populated with tiny creatures called plankton; these are characterized by their size and limited ability for horizontal movement, and therefore move with the currents (NOAA, 2026b). Plankton are usually divided into two sub-groups; phytoplankton and zooplankton. Phytoplankton are plant-like organisms that obtain energy through photosynthesis, while zooplankton are animal-like organisms that graze on phytoplankton or prey on other zooplankton. Together, they form the basis of the ocean food web and are vital to all life on earth. Further, we will discuss each of these types of plankton, and importantly, how we can monitor them.

Phytoplankton

Phytoplankton is the dominant primary producer in the ocean, making it the base of the ecosystem. These tiny creatures account for approximately half the global oxygen production (McQuatters-Gollop et al., 2011). They play a key role in carbon sequestration (Polimene et al., 2017). They can also grow excessively and form HABs, which can negatively affect other organisms (Smayda, 1997). Monitoring these organisms is important, but it can also be challenging.

In the field, multiple water samples can be taken from various locations and depths. These samples will contain a wide range of phytoplankton species, as well as information about other aspects of the samples biochemical properties (Karlson et al., 2010). Laboratory analysis can subsequently be used to quantify abundance and identify species, among other information. This approach is time-consuming and expensive, but it remains the most effective method for obtaining a highly detailed picture. One common laboratory method is pigment analysis. Phytoplankton rely on photosynthetic pigments to convert sunlight into chemical energy. The main pigment is chl *a*, supported by various accessory pigments; observing these pigments can be used as a proxy for observing phytoplankton. chl *a* is particularly valuable for estimating phytoplankton biomass because it occurs in all phytoplankton groups. It is not only in the lab that this pigment can be detected. A wide range of other techniques exists for measuring this pigment, including remote sensing from satellites and in-situ fluorometry, among others. An example of a processed ocean color satellite image can be

seen in Fig. 2.3¹. This satellite product observes the ocean color for chl *a*, which is the proxy for phytoplankton. Here we can observe the complex and patchy distribution of chl *a*. It should also be noted that each pixel in the satellite image is roughly 4x4 km in size; thus, the spatial resolution is limited. For greater resolution, in-situ measurements might be required. The ease of measurement and prevalence of chl *a* in phytoplankton means that chl *a* is widely regarded as the most suitable proxy for primary production in the ocean (Huot et al., 2007).

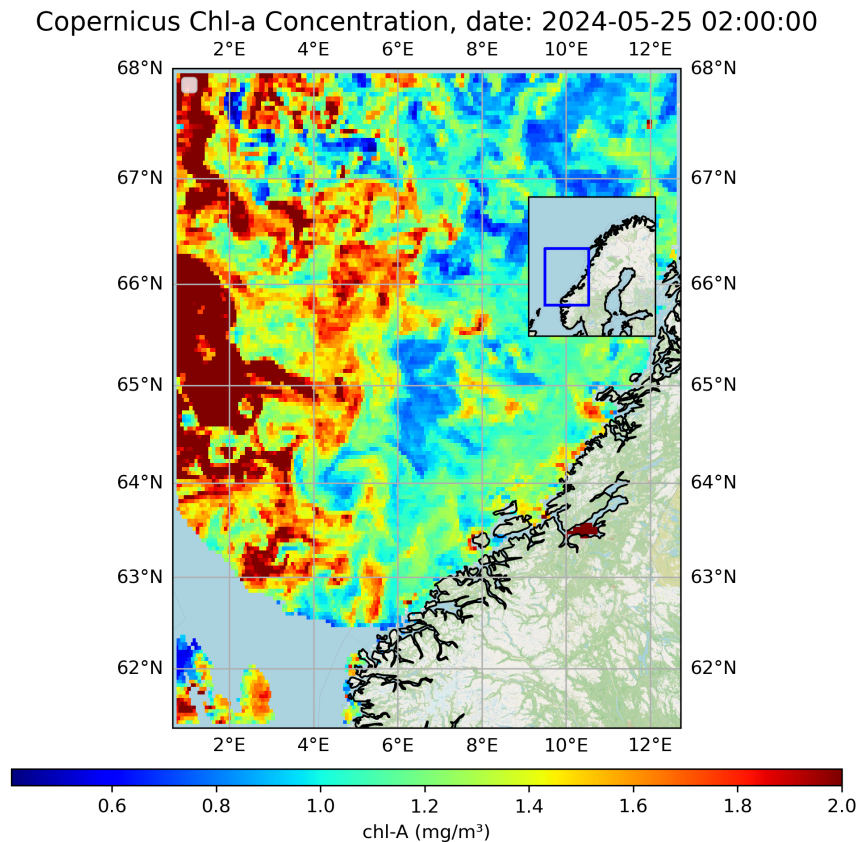


Fig. 2.3: Satellite-derived chlorophyll *a* along the Norwegian coast (clear-sky day). Data courtesy of the E.U. Copernicus Marine Service.

Shifting attention to the Norwegian coastline, seasonal changes in sunlight cause recurring phytoplankton blooms along the coast (Dale et al., 1999). These blooms emerge in different coastal areas, depending on a variety of environmen-

¹Global Ocean Color Plankton and Reflectances MY L3 daily observations. E.U. Copernicus Marine Service Information (CMEMS). Marine Data Store (MDS). DOI: <https://doi.org/10.48670/moi-00282>

tal conditions; a major bloom typically develops in spring around May, followed by a weaker bloom in autumn. The spring bloom is in turn followed by a bloom of the zooplankton species *Calanus finmarchicus* (Broms and Melle, 2007), which is a key dominant zooplankton species.

Zooplankton

Zooplankton get their energy from grazing on phytoplankton and preying on other zooplankton. They can be further divided into two groups; holoplankton, which spend their whole life as plankton, and meroplankton, which only spend part of their life as plankton (Lalli and Parsons, 1997). The latter group includes fish eggs and several marine creatures that spend part of their lives as larvae. Zooplankton are also an important food source for fish and other aquatic animals (O'Brien, 1979).

An interesting behavior of zooplankton is that many species engage in a daily vertical migration, this is called diel vertical migration (DVM). The plankton move up and down in the water column at different times of day; most commonly, they go to the surface during nighttime and dive down during the day. Many point to avoiding visual predators as the main reason for this behavior (Lampert, 1993), but there are other explanations. Before the phenomenon was widely known, the migration created a perplexing result for vessels with echo-sounders: in the echogram, a phantom bottom would appear during the day, but at night, it would rise and disperse (Dietz, 1962). It turned out that these layers in the echogram were the result of so many small aquatic creatures migrating up and down that they were detectable by acoustic-sensors.

Using acoustics has, in turn, become one of the more robust approaches for detecting zooplankton (Lavery et al., 2010). However, numerous other observation techniques exist. A traditional method involves towing a fine-meshed net through the water, where the mesh size determines which species are retained. Additional approaches rely on optical detection using particle-imaging cameras. One example is the SilCam, developed by SINTEF (Davies, 2026), which is a backlight imaging system that can be mounted on various platforms.

Zooplankton, however, are more difficult to detect with satellite imagery. Unlike phytoplankton, which contain distinct pigments for photosynthesis, zooplankton are mostly transparent and may also reside at depths beyond the reach of satellite sensors. Nonetheless, satellite detection is possible for certain species. For instance, studies suggest that the red-colored *C. finmarchicus* can be observed using satellite data (Basedow et al., 2019). This zooplankton species is the most abundant in the North Atlantic (Broms and Melle, 2007) and is the primary focus of the sampling campaigns.

2.2 Ocean observing platforms

We have already talked a bit about the different types of sensors used to observe the ocean; now we will look at the platforms on which these sensors are mounted.

Ocean observing platforms cover a range of functions and costs. The historical way of observing the ocean is to use a research vessel. These boats can host scientists, all kinds of sensors, and laboratories for analysis. They are often the basis for launching the smaller observation platforms. These smaller platforms include fixed stations called Moored buoys, drifting platforms further divided into surface drifters and Argo floats, unmanned aerial vehicles (UAV), remotely operated vehicles (ROVs), and lastly, autonomous vehicles. The moored buoys stay fixed in one place and collect data for one location for a long duration; sometimes they are equipped with sensors at different depths or sensors that can be moved up and down. The drifters are moved by currents and can therefore move over large distances; Argo floats even have the ability to move up and down in the water column. ROVs have the ability to move freely in the water column but are usually tethered and use a human operator. Lastly, we have the autonomous vehicles, which are the focus of this thesis. These are platforms equipped with some form of propulsion, and therefore movements are not fully dependent on currents. Importantly, for our aims, what matters is their capacity to integrate real-time sensor inputs with onboard computing in order to change the sampling design. Equipped with various sensors, their own propulsion systems, and onboard computers, autonomous ocean vehicles can modify their sampling strategies on the fly in response to new measurements, all without human intervention. In a variable and rapidly changing ocean, this capability is extremely valuable. It may, for example, allow a vehicle to track a river plume or locate and follow dense aggregations of plankton. Autonomous ocean vehicles are generally classified into two main categories: autonomous surface vehicles (ASVs) and autonomous underwater vehicles (AUVs). This work focuses on the latter, and in particular on AUVs that use propellers for propulsion. Nevertheless, the algorithms presented are not restricted to AUVs; although certain components may be vehicle-specific, the overall methods are broadly applicable to autonomous platforms facing a sampling design problem.

The AUVs employed in this thesis were developed by OceanScan and are referred to as light autonomous underwater vehicles (LAUVs). Similar to many other AUVs, they feature a torpedo-shaped hull; an example of such a vehicle is shown in Fig. 2.4 with some key components highlighted. We are particularly focused on the sensors: at the front, there is the CTD sensor, which records salinity, temperature, and depth; in the same area, there is also the fluorometer, which measures chl *a*. Slightly further back, we have the echo-sounder, used to detect zooplankton and other types of biomass, and located behind it is the SilCam, which acquires images of zooplankton during sampling. Although these LAUVs share a largely common design, they differ in terms of payload and technical specifications. A crucial property is that each vehicle is light enough to be carried by two people and deployed and recovered from a small boat, eliminating the need for costly research vessels. They are positively buoyant and must remain in motion in order to stay submerged.

A fundamental challenge for all such systems is that electromagnetic radiation does not propagate far underwater, which means that GPS and most common communication methods are ineffective below the surface. While sub-

merged, the AUV must therefore maintain an internal estimate of its position. This positioning is often incorrect and increases with the duration underwater. Alternative approaches to underwater positioning rely on acoustics, either by detecting known features underwater or by getting the range to other known positions (Vickery, 1998; Zhou et al., 2023). Because sound waves travel more effectively in water than electromagnetic waves, they can be used for communication between vehicles by transmitting acoustic signals and listening for them. Underwater communication is thus possible, but it remains constrained relative to communication at the surface. However, when the AUV surfaces, it can use its antenna for communication and to acquire a GPS fix. Taken together, these capabilities allow the AUV to support a wide range of applications and address diverse and interesting sampling design problems.




Fig. 2.4: LAUV with key components for propulsion/control, communication, and sensing.

We are interested in using these autonomous platforms to examine dynamic ocean processes, such as ocean fronts and plankton distributions. We have several methods available for this, but our primary emphasis here is on the AUV. We will explore how these phenomena can be modeled in a way that incorporates incoming data, allowing the system to adapt and make improved decisions about where to sample. The first step is to look at ways to model the ocean.

Chapter 3

Modeling and observations

“We demand rigidly defined areas
of doubt and uncertainty!”

Douglas Adams
The Hitchhiker’s Guide to the
Galaxy

Ocean models usually aim to describe one or more physical or biochemical features. These models vary greatly in complexity, spatio-temporal resolution, and assumptions, each designed for different purposes. Increasing complexity and resolution can enhance the skill of the model, but it often comes at the cost of expensive computations. In this chapter, the goal is to describe the models we used onboard the agents that enable them to make better decisions on the sampling design. Usually, we are only interested in one or a small selection of features to model at a time. The agents are exploring relatively small areas, but we are interested in high resolution in these areas. The model needs to be very reactive to the continuously incoming data and use that to make predictions and model the uncertainty of the predictions. Finally, the model should be able to run onboard the agent, which means that computations have to be fast. The core method for this section is to model the ocean variable as a Gaussian random field (GRF)

3.1 Basics of Gaussian Random Fields

We consider an ocean variable x defined in spatial-temporal locations $\mathcal{S} = (\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N)$. For the applications here, the locations are defined by the lateral components s_x and s_y , the depth s_z , and the temporal component t , or in some cases, a subset of these components. This variable can be salinity, temperature, chl a , etc. It is common to model the ocean variable using a GRF. A GRF is fully characterized by its mean vector $E[x(\mathbf{s})] = \mu(\mathbf{s})$ and

Model	Covariance
Exponential	$C(h) = \sigma^2 \exp\left(-\frac{h}{\varphi}\right)$
Matérn 3/2	$C(h) = \sigma^2 \left(1 + \frac{h\sqrt{3}}{\varphi}\right) \exp\left(-\frac{h\sqrt{3}}{\varphi}\right)$
Gaussian	$C(h) = \sigma^2 \exp\left(-\frac{h^2}{2\varphi^2}\right)$

Table 3.1: Three common covariance functions. Here h refers to the distance between two points \mathbf{s} and \mathbf{s}' ($h = \|\mathbf{s} - \mathbf{s}'\|$), and φ is a model parameter.

its covariance $\text{cov}(x(\mathbf{s}), x(\mathbf{s}')) = c(\mathbf{s}, \mathbf{s}')$, see, e.g., Cressie and Wikle (2015). Represented on the domain \mathcal{S} , we write

$$\mathbf{x}_{\mathcal{S}} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathcal{S}}, \boldsymbol{\Sigma}_{\mathcal{S}}), \quad (3.1)$$

where $\mathcal{N}(\cdot, \cdot)$ denotes the multivariate normal distribution with mean vector $\boldsymbol{\mu}_{\mathcal{S}}$ and covariance matrix $\boldsymbol{\Sigma}_{\mathcal{S}}$.

When modeling the ocean, the prior mean $\boldsymbol{\mu}$ represents our best understanding prior to deployment in the ocean. This knowledge can be derived from various sources, including rule-based models, physics-based models (E.g Fig. 2.2), and previously collected survey or satellite data (E.g Fig. 2.3).

The covariance matrix describes how each pair of variables in locations \mathbf{s}_i and \mathbf{s}_j varies together. The covariance between two locations can be written as $\text{cov}(x(\mathbf{s}_i), x(\mathbf{s}_j)) = c(\mathbf{s}_i, \mathbf{s}_j)$. The constructed covariance matrix with each element is

$$\boldsymbol{\Sigma}_{\mathcal{S}} = \begin{pmatrix} c(\mathbf{s}_1, \mathbf{s}_1) & c(\mathbf{s}_1, \mathbf{s}_2) & \dots & c(\mathbf{s}_1, \mathbf{s}_N) \\ c(\mathbf{s}_2, \mathbf{s}_1) & c(\mathbf{s}_2, \mathbf{s}_2) & & c(\mathbf{s}_2, \mathbf{s}_N) \\ \vdots & & \ddots & \\ c(\mathbf{s}_N, \mathbf{s}_1) & c(\mathbf{s}_N, \mathbf{s}_2) & & c(\mathbf{s}_N, \mathbf{s}_N) \end{pmatrix}. \quad (3.2)$$

The diagonal elements of this matrix are the variances for each location, and the non-diagonal elements represent the covariance between two locations. There are several ways to construct the covariance matrix; here we use parametric covariance functions that depend on some distance h between two locations.

In Table 3.1 we list some commonly used covariance kernels. These functions describe how much the variables change together based on distance h . Fig. 3.1 shows how the covariance drops with the distance h , the variance parameter is $\sigma = 1$ and the correlation decay parameter φ set in such a way that $C(300; \varphi) = 0.3$. From this figure, we can see that the exponential covariance function drops the fastest, the Gaussian stays high for longer, and the Matérn (3/2 and 5/2) are somewhere in the middle. This affects the "roughness" or smoothness of the realizations. To illustrate the roughness, we can simulate 3 fields using 3 different covariance functions with zero mean. This is defined in a 2D domain in a square with a size of 1 km \times 1 km. The distance h is defined as the Euclidean distance between two locations. The three resulting fields are shown in Fig

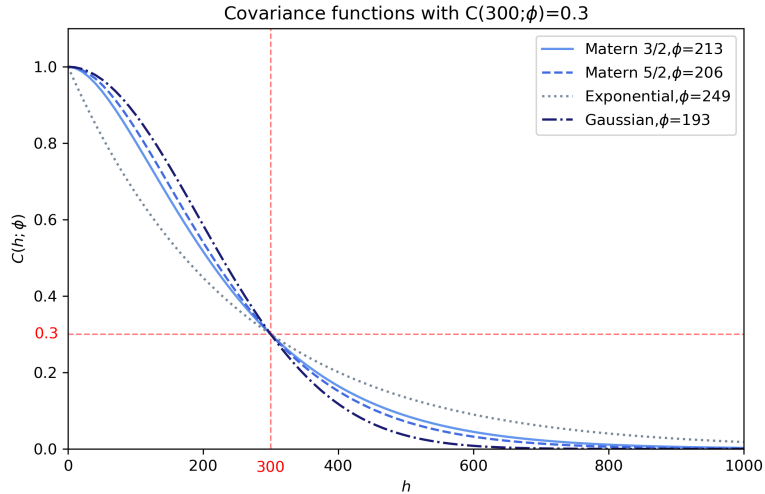


Fig. 3.1: Different parametric covariance functions with. Here the field variance is $\sigma^2 = 1$, changing this will quadratically scale the height of the functions.

3.2. From left to right, the fields appear rougher on the left and smoother on the right. Choosing the least smooth (Exponential) adapts easily to the data, but the correlations drop quickly and do not provide accurate predictions far away. The Gaussian covariance function is very smooth, often resulting in high correlation at long distances. Due to the smoothness, it can make the field unrealistically smooth.




Fig. 3.2: Different realizations of a field using a different covariance function. a) Uses an exponential covariance function, b) uses a Matérn ($\nu = 3/2$) and c) uses a Gaussian covariance function. The realizations become smoother from left to right.

This example shows 3 different covariance functions; there are many more, but these are very commonly used and well studied. It could also be noted

that the Matérn covariance function has a smoothness parameter ν , and the exponential and Gaussian are cases of this covariance function with parameters $\nu = 1/2$ (exponential) and $\nu \rightarrow \infty$ (Gaussian).

Separable Kernel

So far, we have looked at isotropic covariance functions, which means that the covariance between points \mathbf{s}_i and \mathbf{s}_j depends only on the distance and not the direction. In the fields in Fig. 3.2 there is no form of orientation of the field. This assumption does not work well for the ocean, which is not isotropic, especially when also considering the depth. In some sense, the ocean is flat; usually, the vertical variability is much higher than the lateral (Talley, 2011). A model needs to take this into account to be able to model both lateral variation and depth variation. In addition to depth, we do not want the model to be static over time. For a shorter time-frame, this might be a serviceable assumption. However, for longer time periods, one might want to take this into account.

One of the simpler ways to model this anisotropy is by using a separable covariance function. With a separable covariance function,

$$c(\mathbf{s}_i, \mathbf{s}_j) = \sigma^2 \cdot C_{xy}(h_{xy}; \varphi_{xy}) \cdot C_z(h_z; \varphi_z) \cdot C_t(h_t; \varphi_t). \quad (3.3)$$

The functions $C_d(\cdot; \varphi_d)$ represent different covariance functions for different dimensions. These functions can, for example, be chosen from the suggested functions in Table 3.1, with a common variance σ^2 . The distance values for h_{xy} , h_z , and h_t represent lateral, depth, and temporal separation between locations \mathbf{s}_i and \mathbf{s}_j . These distances are defined as

$$h_{xy} = \|(s_{x,i}, s_{y,i}) - (s_{x,j}, s_{y,j})\|, \quad h_z = |s_{z,i} - s_{z,j}|, \quad h_t = |s_{t,i} - s_{t,j}|.$$

In Fig. 3.3 we generated a field using a separable covariance function. It is perhaps not obvious to observe, but the effective correlation range differs across different dimensions.

3.2 Conditioning on data

One of the stated goals of the model is that it should be able to be updated using data from the onboard sensors. During a field experiment, the agent collects data; observed data for a location \mathbf{s} is denoted as $y(\mathbf{s})$ and comes from the true field value $x(\mathbf{s})$ plus independent Gaussian noise $\epsilon(\mathbf{s})$. The observation model for a point \mathbf{s} can be written as

$$y(\mathbf{s}) = x(\mathbf{s}) + \epsilon(\mathbf{s}), \quad \epsilon(\mathbf{s}) \sim \mathcal{N}(0, \tau^2), \quad (3.4)$$

where τ^2 is the observation noise variance. After observing the data $\mathbf{y}_{\mathcal{S}}$, we can derive the posterior distribution of $\mathbf{x}_{\mathcal{P}} | \mathbf{y}_{\mathcal{S}}$, where \mathcal{P} denotes the set of points of interest. This set $\mathcal{P} = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_M)$ may coincide with the sampled points $\mathcal{S} = (\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N)$, or it could represent a regular grid, or even a trajectory

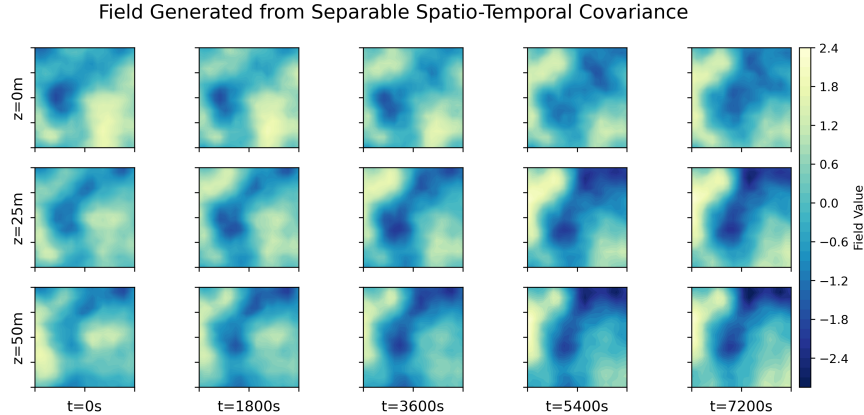


Fig. 3.3: Field simulated using a separable covariance function in the lateral dimension, depth dimension and temporal dimension.

along which the agent is considering making observations. The conditional model is again Gaussian, with mean $\mathbf{m}_{\mathcal{P}|\mathcal{S}}$ and covariance matrix $\mathbf{V}_{\mathcal{P}|\mathcal{S}}$. This is derived from the following equations

$$\begin{aligned}
 \mathbf{x}_{\mathcal{P}}|\mathbf{y}_{\mathcal{S}} &\sim \mathcal{N}(\mathbf{m}_{\mathcal{P}|\mathcal{S}}, \mathbf{V}_{\mathcal{P}|\mathcal{S}}), \\
 \mathbf{m}_{\mathcal{P}|\mathcal{S}} &= \boldsymbol{\mu}_{\mathcal{P}} + \boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}(\boldsymbol{\Sigma}_{\mathcal{S}} + \mathbf{R})^{-1}(\mathbf{y}_{\mathcal{S}} - \boldsymbol{\mu}_{\mathcal{S}}), \\
 \mathbf{V}_{\mathcal{P}|\mathcal{S}} &= \boldsymbol{\Sigma}_{\mathcal{P}} - \boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}(\boldsymbol{\Sigma}_{\mathcal{S}} + \mathbf{R})^{-1}\boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}^t,
 \end{aligned} \tag{3.5}$$

where \mathbf{R} is the measurement noise covariance matrix $\mathbf{R} = \tau^2 \mathbf{I}_{\mathcal{S}}$ and $\boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}$ is a cross covariance matrix between the variables at the points we are interested in \mathcal{P} and the data points \mathcal{S} . If we want the posterior distribution of $\mathbf{x}_{\mathcal{S}}$ on $\mathbf{y}_{\mathcal{S}}$, then we simply set $\mathcal{P} = \mathcal{S}$. Often, we work with a regular vector grid field $\mathbf{G} = (\mathbf{g}_1, \mathbf{g}_2, \dots, \mathbf{g}_M)$, and the observed values are assigned to nodes on the grid. Then we can set $\mathcal{P} = \mathbf{G}$.

Using Eq. (3.5), we can simulate a 2D field with $\mathcal{D} = [0, 1000] \times [0, 1000]$, a Matérn (3/2) covariance function with a variance parameter $\sigma^2 = 2^2$, a correlation decay parameter $\varphi = 200$, and a prior mean shown in Fig. 3.4 a). One realization of this field is generated and shown in Fig. 3.4 b), then we observe data along the red path. The path is then \mathcal{S} and the data along this path is $\mathbf{y}_{\mathcal{S}}$, the points we are interested in predicting is a regular grid inside \mathcal{D} . Using Eq. (3.5) we can Krieger the rest of the field conditional on the observed data, this can also be called the posterior mean. The resulting field is shown in the left field (Fig. 3.4 c). The plot shows that the conditional field is closer to the true field after sampling, especially near the observed path. This is further highlighted in Fig. 3.5. The left plot shows the absolute error between the prior mean and the true field, in contrast, the middle plot is the posterior error. It is clear that the error is greatly reduced after adding the data. The right plot

(Fig. 3.5 c)) shows how the variance, or the diagonal terms of matrix $V_{\mathbf{G}|\mathcal{S}}$, is lower closer to the path.

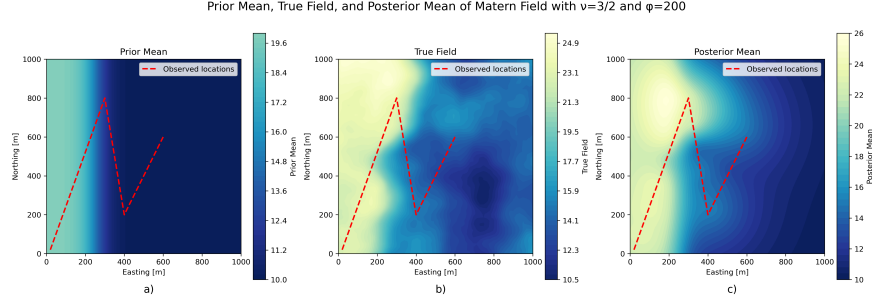


Fig. 3.4: a) Prior mean of the field, b) the true field and c) the posterior mean. All plots shows how the agent moves in the field. The posterior mean is closer to the true field close to the agents path.




Fig. 3.5: a) Prior absolute error, b) posterior error and c) the posterior variance. All plots show the observed path. The absolute error is reduced when conditioning on the observed locations.

3.3 Computational considerations

A requirement for the models is that the necessary computations should be fast enough to run on a computer inside the vehicle with a reasonable deliberation time. The main issue arises from the number of points conditioned on M and the number of predicted points N resulting in large matrix multiplications and inversions. A normal way to model the field is to use regular grid nodes \mathbf{G} that extend the lateral plane of the domain and relevant depths. The spacing between the nodes and the size of the domain then determines the number of predicted points N . As the number of observed grid nodes M increases, Eq. (3.5) becomes more and more expensive. Because data points are usually assimilated in batches during sampling, the computation time for each batch can be greatly

reduced by using a batch-sequential formula and updating the conditional mean and variance for the field at each step; see Chevalier et al. (2013) or Ge et al. (2023). However, the computation time is affected by the size and resolution of the grid, and sometimes the onboard computer cannot handle the desired size and resolution in a reasonable time. Another approach is to not use a grid and only predict the parts of the field that are relevant at any one time. That means that the size of the model in memory is the same as the number of observed points. This also means that one does not need to consider the size, shape, and resolution for the field in each different experiment, making it a bit more flexible to quickly adapt to new conditions. An illustration of the difference between these models can be seen in Fig. 3.6, in Fig. 3.6 a) a normal regular grid is shown, in Fig. 3.6 b) a model using only points along the observed path is shown. This illustration does not show the depth, which increases the number of points for the regular grid, but not for the updating model.

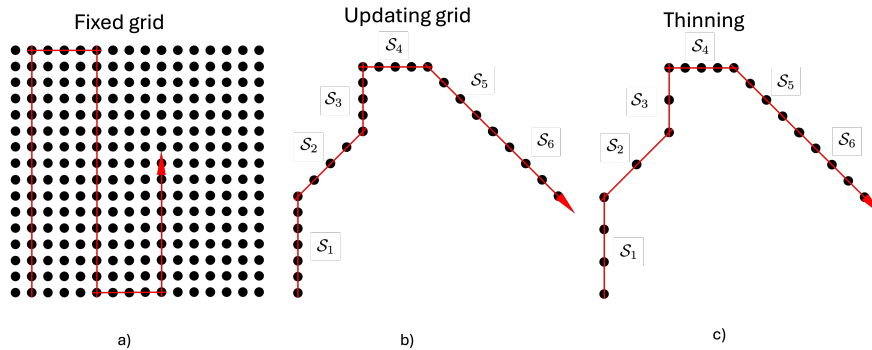


Fig. 3.6: The black dots are the nodes in the model, the red line follows the path of the agent and the red arrow is the agent. a) Illustrates a regular grid, b) Illustrates a model with points only along the path and c) Points being thinned from the model.

This still results in computational challenges. The data is assimilated into the model in batches, the full list of points is written as $\mathcal{S}_{1:J} = (\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_J)$, where each \mathcal{S}_j is a batch of points. For each batch added to the model, the Equations in (3.5) become more expensive, especially the inverse part. Using a block version of the Sherman-Woodbury-Morrison formula (Petersen et al., 2008) this can be done more efficiently. See Olaisen et al. (2025) for how this is done. In summary, this trick revolves around only computing the inverse for the new batch, and store the inverse for the other points from the previous stage.

Despite the trick with matrix inversion, the computation time will increase for each batch added, and there is a chance that the computations will become too expensive. The method is to simply reduce points in the model that are far away in space and time (Olaisen et al., 2025), this thinning will have a limited effect on nearby predictions due to the spatial screening effect (Stein, 2002). An illustration of this can be found between Fig. 3.6 b) and Fig. 3.6 c). When

removing the grid, we have assumed that predicting the whole grid is not always relevant, but sometimes predicting the entire field is valuable; in these cases, the grid can be split into batches $\mathbf{G} = (\mathbf{G}_1, \dots, \mathbf{G}_k)$. Then Eq. (3.5) is used for each batch, note that much of the computations here are only done once. This has the benefit of runtime scaling linearly with the size of the grid instead of cubically; it does, however, come at the cost of not having access to the full covariance matrix.

Ultimately, there are often several ways to control computational demand, so it is crucial to understand the specific requirements of the problem being addressed and work from there.

log-Normal and Poisson extension

The GRF is often straightforward to handle, but it does not always provide a good representation of the variable we are interested in. Therefore, we can extend the framework to include latent GRFs. Certain variables can be patchy; therefore, it may be suitable to model the field variable as a log-normal distribution. Then we can change the observation model to become

$$\log(y(\mathbf{s})) = x(\mathbf{s}) + \epsilon(\mathbf{s}), \quad \epsilon(\mathbf{s}) \sim \mathcal{N}(0, \tau^2). \quad (3.6)$$

Here, the latent field \mathbf{x} becomes the latent GRF. There is also an added benefit that the field values on the exponential scale are always positive.

Another issue arises when dealing with discrete data while the GRF takes continuous values. For instance, when using a particle camera, we may wish to model the number of plankton visible in an image. This plankton count can be treated as Poisson distributed with intensity or mean $\lambda(\mathbf{s}) = \exp(x(\mathbf{s}))$, where x denotes the latent intensity field of the plankton at a location \mathbf{s} . The observation model conditioned on this latent intensity field can then be expressed as

$$\begin{aligned} y(\mathbf{s}) &\sim \text{Poisson}(\lambda(\mathbf{s})), \\ \lambda(\mathbf{s}) &= \exp(x(\mathbf{s})), \end{aligned} \quad (3.7)$$

where the latent intensity field \mathbf{x} here is a GRF. Thus, we can use the same covariance functions for x and similar equations for conditional predictions as shown in Eq. (3.5). The modified equations can be written as

$$\begin{aligned} \mathbf{x}_S | \mathbf{y}_S &\approx \mathcal{N}(\hat{\mathbf{m}}_{\mathbf{x}_S | \mathbf{y}_S}, \hat{\mathbf{V}}_{\mathbf{x}_S | \mathbf{y}_S}), \\ \hat{\mathbf{m}}_{\mathbf{x}_S | \mathbf{y}_S} &= \boldsymbol{\mu}_S + \boldsymbol{\Sigma}_S (\boldsymbol{\Sigma}_S + \mathbf{R})^{-1} (\mathbf{z} - \boldsymbol{\mu}_S), \\ \hat{\mathbf{V}}_{\mathbf{x}_S | \mathbf{y}_S} &= \boldsymbol{\Sigma}_S - \boldsymbol{\Sigma}_S (\boldsymbol{\Sigma}_S + \mathbf{R})^{-1} \boldsymbol{\Sigma}_S^T. \end{aligned} \quad (3.8)$$

Here we need to estimate $\hat{\mathbf{m}}_{\mathbf{x}_S | \mathbf{y}_S}$, $\hat{\mathbf{V}}_{\mathbf{x}_S | \mathbf{y}_S}$ and \mathbf{R} , this can be done using Newton-Raphsons Algorithm shown in Alg. 1.

Algorithm 1 Newton-Raphson Algorithm for posterior approximation in the situation with a Poisson observation model.

Require:

Observations \mathbf{y}_S from a path S , prior-mean $\boldsymbol{\mu}_S$, and covariance matrix $\boldsymbol{\Sigma}_S$ and a tolerance δ

$\mathbf{x}^{(0)} = \log(\mathbf{y}_S + \mathbf{1} \cdot \epsilon_{\text{off}})$ as the initial guess (minuscule $\epsilon_{\text{off}} > 0$)

repeat For each time $k = 0, 1, \dots$: ▷ Loop until convergence

Get \mathbf{z} with entries $z_i = \left(y_i - \exp(x_i^{(k)}) + x_i^{(k)} \cdot \exp -x_i^{(k)} \right) \cdot \exp(x_i^{(0)})$

Get \mathbf{R} with entries $R_{i,i} = \exp\left(x_i^{(k)}\right)$ and $R_{i,j} = 0$ when $i \neq j$

Get $\hat{\mathbf{m}}_{\mathbf{x}_S|\mathbf{y}_S}^{(k)}$ and $\hat{\mathbf{V}}_{\mathbf{x}_S|\mathbf{y}_S}^{(k)}$ from Equation (3.8)

$\mathbf{x}^{(k)} = \hat{\mathbf{m}}_{\mathbf{x}_S|\mathbf{y}_S}^{(k)}$

if $\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| < \delta$ **then**

Return $\hat{\mathbf{m}}_{\mathbf{x}_S|\mathbf{y}_S}^{(k)}$, $\hat{\mathbf{V}}_{\mathbf{x}_S|\mathbf{y}_S}^{(k)}$, \mathbf{z} and \mathbf{R}

end if

until

3.4 Parameter estimation

The choice of parameters will be important for the quality of the predictions and for how well the model fits the data. To illustrate this a bit, we use the previous example of the agent moving in the field (Fig. 3.4). The conditional field generated on the right is done using the same parameters with which the field was generated, but we can test what would happen if we set the correlation decay parameter φ to large or too small when computing the conditional. The conditional mean along the path is shown in Fig. 3.7 a). This display shows that the correct choice of the correlation decay parameter φ allows the posterior to follow the variation of the true field without overly reacting to measurement noise. It is not only the posterior along the path that matters; more importantly, the predictions further away from the path are interesting. If we run 50 experiments with new true fields and a new random path of the agent, we can see what the squared error is a certain distance from the path, this is shown in Fig. 3.7 b). By choosing the correct value for the correlation decay parameter φ , the predictions remain as accurate as possible at a distance. This is a simple experiment where we only set one parameter, but it can illustrate how it affects the models ability to describe the field variance and predict new locations correctly.

Variogram

A popular method commonly used in geostatistics is to use variograms (Hilal et al., 2024), or more commonly semi-variograms, for both model selection and parameter estimation. The semi-variogram can be understood as how much

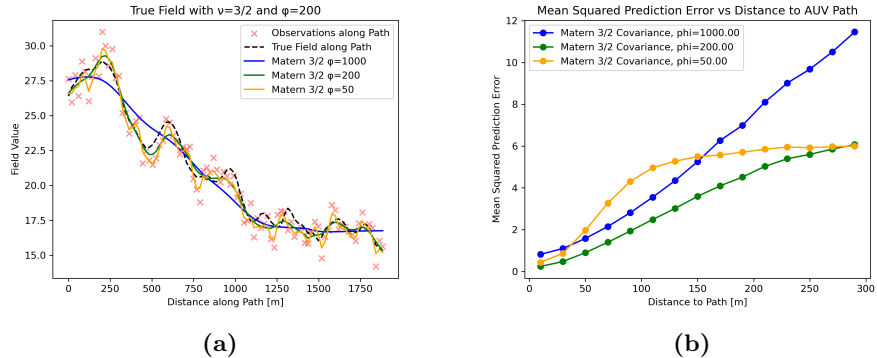


Fig. 3.7: a) Conditional mean along the agent path using different values for the correlation decay parameter φ against the observed points and the true mean, b) Mean squared prediction error based on the distance form the observed path

change there is at a certain distance h . The semi-variogram can be defined as

$$\gamma(h_{ij}) = \frac{1}{2} \text{Var} [(x(\mathbf{s}_i) - x(\mathbf{s}_j))] = \frac{1}{2} E [(x(\mathbf{s}_i) - x(\mathbf{s}_j))^2]. \quad (3.9)$$

An approach is to estimate an empirical non-parametric semi-variogram by taking the average squared difference between observation pairs at a distance h . To do this, distance bins are defined $\mathbf{H} = (\mathbf{h}_1^*, \mathbf{h}_2^*, \dots, \mathbf{h}_K^*)$ around center points such that $\mathbf{h}_k^* = (h_k^* - d, h_k^* + d)$. Then, we can define a set of point pairs where the distance is within that range $\mathcal{S}_k = \{(\mathbf{s}_i, \mathbf{s}_j); \|\mathbf{s}_i - \mathbf{s}_j\| \in \mathbf{h}_k^*, \mathbf{s}_i, \mathbf{s}_j \in \mathcal{S}\}$. Then, the empirical semi-variogram from the midpoints h_k^* is defined as the average squared difference between the point pairs in \mathcal{S}_k . This can be written as

$$\hat{\gamma}(h_k^*) = \frac{1}{2 \cdot N_{\mathcal{S}_k}} \sum_{(\mathbf{s}_i, \mathbf{s}_j) \in \mathcal{S}_k} (y(\mathbf{s}_i) - y(\mathbf{s}_j))^2, \quad (3.10)$$

where $N_{\mathcal{S}_k}$ is the number of pairs in \mathcal{S}_k . An example of the empirical variogram is plotted in Fig. 3.8; the blue line represents the empirical variogram for different distances. In this example, $N = 1000$ points are drawn from a field similar to that shown in Fig. 3.2 with a Matérn($\nu = 3/2$) covariance function, with the correlation decay parameter $\varphi = 50$, variance parameter $\sigma^2 = 1.2^2$, and iid measurement noise $\epsilon(\mathbf{s}_i) \sim \mathcal{N}(0, 1^2)$. In the figure, one can see that there is a jump from 0 up to the first point; this value is called the nugget and can be interpreted as measurement noise in this context. Secondly, the variogram flattens out at a certain point; this is the distance where two points are uncorrelated, and this is often called the effective range, which is often defined as the distance where the variogram reaches 95 % of the plateau height. Lastly, it is the plateau itself; this height is called the sill. There are several considerations when constructing the semi-variogram, such as the number of bins and maximum distance, and good practices depending on the dataset; see, e.g., Müller

and Zimmerman (1999) for a much deeper discussion.

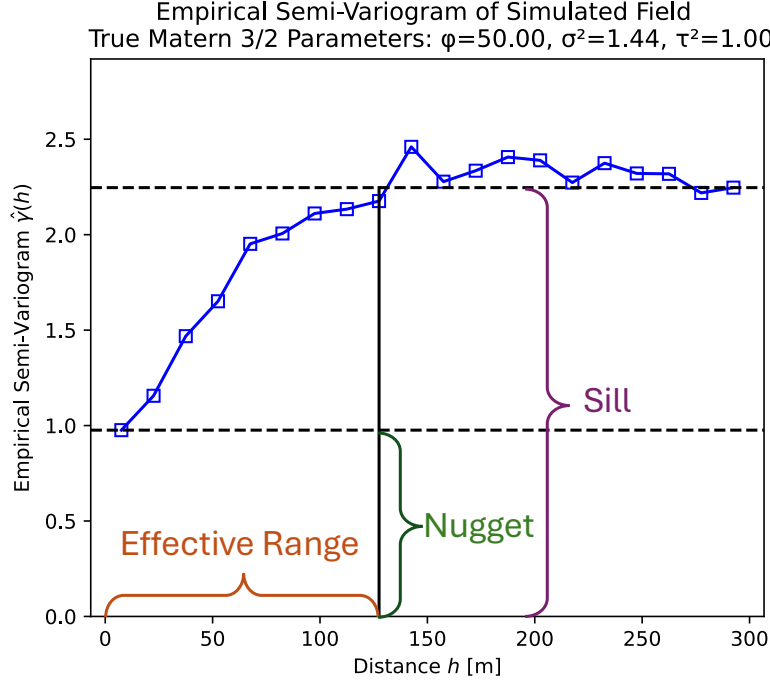


Fig. 3.8: Empirical semi-variogram with the effective range, nugget and sill.

Variograms can be defined with the parametric covariance functions in Tab. 3.1 and using Eq. (3.9). This gives the semi-variogram

$$\gamma(h; \boldsymbol{\theta}) = \tau^2 + (\sigma^2 - C(h; \varphi)), \quad (3.11)$$

where $\boldsymbol{\theta} = (\tau^2, \sigma^2, \varphi)$. We can relate parameters with features in the non-parametric case in Fig. 3.8, which results in $\gamma(0) = \tau^2$ (nugget) and $\gamma(h \rightarrow \infty) = \tau^2 + \sigma^2$ (sill) and effective range can be found by solving $\gamma(h; \boldsymbol{\theta}) = 0.95 \cdot (\tau^2 + \sigma^2)$ for h . Similarly to the non-parametric case we want to estimate these parameters using the data. Utilizing the empirical semi-variogram and minimizing the weighted least squares (WLS) between the empirical semi-variogram defined in Eq. (3.10) and the parametric semi-variogram in Eq. (3.11), this results in the following optimization problem

$$\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \frac{1}{N_p} \sum_{h^*} N_{S_k} \cdot (\hat{\gamma}(h_k^*) - \gamma(h_k^*; \boldsymbol{\theta}))^2, \quad (3.12)$$

where N_{S_k} is the number of pairs in S_k , and N_p is the total number of pairs in \mathcal{S} . To continue with the example from earlier in the section, we can use the

empirical semi-variogram to estimate the parameters $\boldsymbol{\theta} = (\tau^2, \sigma^2, \varphi)$. This is illustrated in Fig. 3.9. Here we see the empirical variogram and the estimated empirical variogram with the red line, and the true variogram in green. The resulting variogram is very close to the true variogram, this is not always the case. Some factors that can make estimation more difficult are having fewer points, unevenly distributed points, or a more complicated underlying covariance structure.

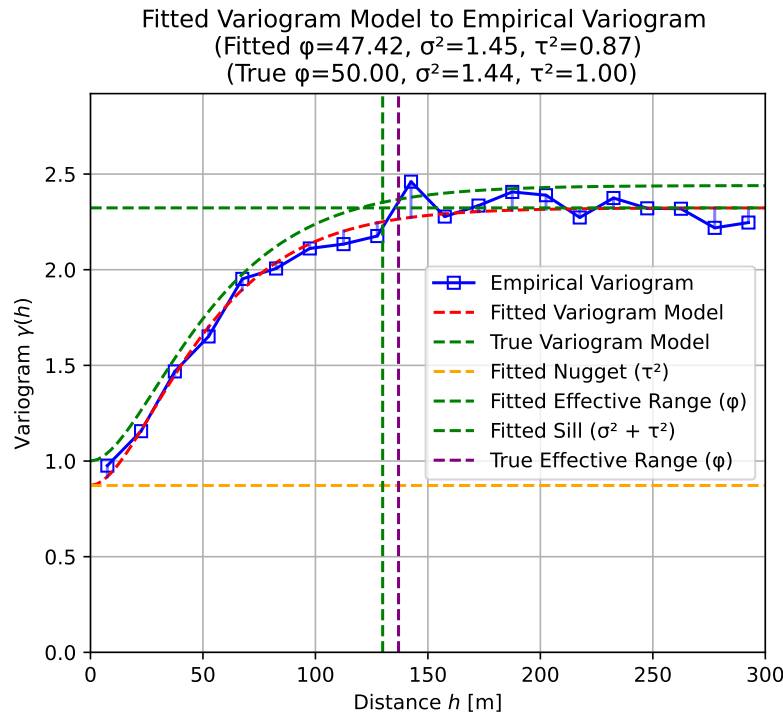


Fig. 3.9: Fitted parametric variogram with the help of the non-parametric variogram using WLS.

This method can also be used for model selection to determine which model gives the lowest WLS in the end. Naturally, there are a range of variations on how to fit these variograms and use them for model selection. For a better discussion, see, e.g., Müller and Zimmerman (1999) or Mahdi et al. (2020).

Maximum (log-)likelihood

Another approach is to use maximum likelihood estimation, or log-likelihood for practical reasons, to estimate the parameters. Provided a dataset $\mathbf{y}(\mathcal{S})$ is drawn from a multivariate normal distribution with $\mathbf{y}(\mathcal{S}) \sim \mathcal{N}(\mathbf{m}_{\mathcal{S}}, \mathbf{K}_{\mathcal{S}})$, where

$\mathbf{K}_S = \Sigma_S + \mathbf{I}\tau^2$, we can define the likelihood as

$$L(\mathbf{y}; \mathbf{x}, \boldsymbol{\theta}) = \frac{1}{(2\pi)^{N/2} \cdot \det(\mathbf{K}_S)^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{y}_S - \mathbf{m}_S)^T \mathbf{K}_S^{-1}(\mathbf{y}_S - \mathbf{m}_S)\right), \quad (3.13)$$

where $\boldsymbol{\theta} = (\tau^2, \sigma^2, \varphi)$ and N is the number of points. In this case, we are interested in finding the parameters that maximize the likelihood function. Practically, it is easier to optimize the log-likelihood, this is then defined as

$$\log L(\mathbf{y}; \mathbf{x}, \boldsymbol{\theta}) = \frac{N}{2} \ln 2\pi - \frac{1}{2} \ln \det(\mathbf{K}_S) - \frac{1}{2}(\mathbf{x}_S - \mathbf{m}_S)^T \mathbf{K}_S^{-1}(\mathbf{x}_S - \mathbf{m}_S). \quad (3.14)$$

There is no closed form solution to this problem, but there are many options for obtaining numerical optimization. The most basic approach is a simple grid search for all probable parameter values, but a method like the Newton-Raphson algorithm can be executed much faster. The Newton-Raphson algorithm relies on finding the gradient vector \mathbf{z} , also called the score vector, and the Hessian \mathbf{H} . The score vector entries are defined by

$$z_i = \frac{\partial}{\partial \theta_i} \log L(\mathbf{y}; \mathbf{x}, \boldsymbol{\theta}), \quad (3.15)$$

and the entries of the Hessian can be written as

$$H_{ij} = \frac{\partial^2}{\partial \theta_i \partial \theta_j} \log L(\mathbf{y}; \mathbf{x}, \boldsymbol{\theta}). \quad (3.16)$$

Then the iterative equations can be written as

$$\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}^{(i)} - \mathbf{H}^{-1} \mathbf{z} \cdot \lambda, \quad (3.17)$$

where λ is a step-length parameter. This iteration can be done until it converges.

We can use the Newton-Raphson method for the same parameter estimation as in Fig. 3.9. In Fig. 3.10, the path-steps for $\boldsymbol{\theta}^{(i)}$ are shown for the correlation decay parameter φ and variance parameter σ^2 . The measurement noise parameter τ^2 is also included in the algorithm but is not shown in the figure. This path is plotted on top of the log-likelihood surface with a marker for the true parameters; the algorithm converges to a point close to the true values, similar to the parameters found using the semi-variograms. The log-likelihood surface is plotted with τ^2 at its final value after the optimization, not the true value.

In some cases, the score vector and Hessian in equations (3.15) and (3.16) are cumbersome to compute analytically. In such situations, numerous methods and solvers instead approximate these quantities numerically (see, e.g., Broyden, 1970).

Cross-validation with spatio-temporal blocks

The last approach discussed here is using cross-validation (CV) with spatio-temporal blocks. This is also a more computationally expensive method than

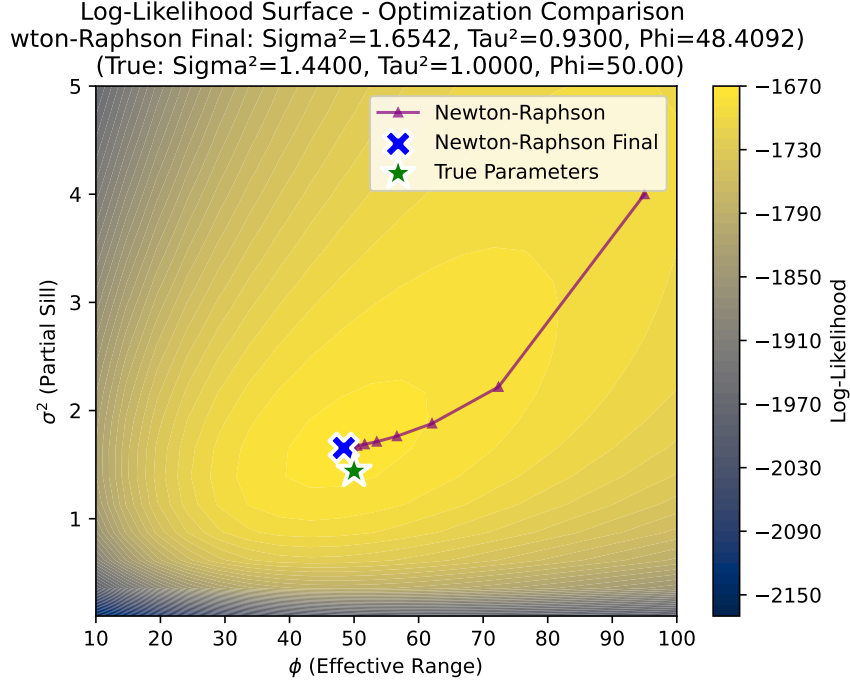


Fig. 3.10: Newton-Raphson for finding the parameters $\theta = (\tau^2, \sigma^2, \varphi)$. The path of the evaluated parameters, a mark for both the final parameters and the true parameters are shown. The log-likelihood surface is plotted using the final value for the noise variance τ^2 .

the previously mentioned methods, but in some cases, the other two do not work as well. If the field is very anisotropic or the log-likelihood and derivatives are hard to compute, one may want to use other approaches.

The general idea of cross-validation splits the data \mathbf{y} into K disjoint sets $\mathbf{y} = (\mathbf{y}^{(\kappa_1)}, \dots, \mathbf{y}^{(\kappa_K)})$, where κ_j is an indicator for the j th set. These sets are often called folds. Then, for each fold κ_j , we use the data *not* in the j th set $\mathbf{y}^{(\kappa_{-j})}$ to estimate a predictor $\hat{f}_{(\kappa_{-j})}(\cdot)$. We then use the predictor to predict the data points \mathbf{y}^{κ_j} , and finally, we compute a score for this fold $L(\mathbf{y}^{(\kappa_j)}, \hat{f}_{(\kappa_{-j})}(\mathcal{S}^{(\kappa_j)}))$ (Hastie et al., 2009). When this is done for each fold, we can compute an average score. The average score can be used to find which hyperparameters work best or to estimate the accuracy of the model. In the model described above, the predictor is the conditional distribution shown in equations (3.5) $\mathbf{y}_{\mathcal{S}^{(\kappa_j)}} | \mathbf{y}_{\mathcal{S}^{(\kappa_{-j})}} \sim \mathcal{N}(\mathbf{m}_{\mathcal{S}^{(\kappa_j)}} | \mathcal{S}^{(\kappa_{-j})}, \mathbf{V}_{\mathcal{S}^{(\kappa_j)}} | \mathcal{S}^{(\kappa_{-j})} + \tau^2 \mathbf{I})$. The scoring function L can just be the mean-squared error,

$$\text{MSE}(\kappa_j; \boldsymbol{\theta}) = \frac{1}{N} \sum_{\mathbf{s} \in \mathcal{S}^{(\kappa_j)}} (y(\mathbf{s}) - m(\mathbf{s}))^2, \quad (3.18)$$

where m is the conditional mean for the point \mathbf{s} . However, there are ways to utilize the variance of the predictions. We can compute the negative log-likelihood for the out-of-sample points, this can be written as

$$nll(\kappa_j; \boldsymbol{\theta}) = -\frac{1}{N} \sum_{\mathbf{s} \in \mathcal{S}^{\kappa_j}} \ln p(\mathbf{y}(\mathbf{s}); \mathbf{y}_{\mathcal{S}^{(\kappa_{-j})}}, \boldsymbol{\theta}), \quad (3.19)$$

where $p(\cdot)$ is the likelihood for a normal distribution. The reason for using the negative here is simply that, with the rise of machine learning techniques, it has become more popular to show problems as minimization problems using loss functions. Here, we use a Gaussian likelihood function, but remember that the predictions can have a different distribution. In Eq. (3.8), we showed a model with Poisson data; thus, this would be the appropriate likelihood function.

We have not discussed how the data are split up, and this is the key difference between normal CV and CV with spatio-temporal blocks. Here we will split the data into K different folds with roughly the same size, but instead of distributing the data randomly into the folds, we use spatio-temporal blocks (Stock, 2025). The issue with splitting the data randomly can be seen in Fig. 3.11 a), due to the underlying spatial structure of the data; 2 close points will be highly correlated. This is also not representative of how the model will be used; in the experiments, the model will make predictions for locations that are further away from neighboring points. To mitigate these issues, we employ spatial blocks as shown in Fig 3.11 b). Note that this is just one of many different ways to split up the field; see Stock (2025) for more examples.

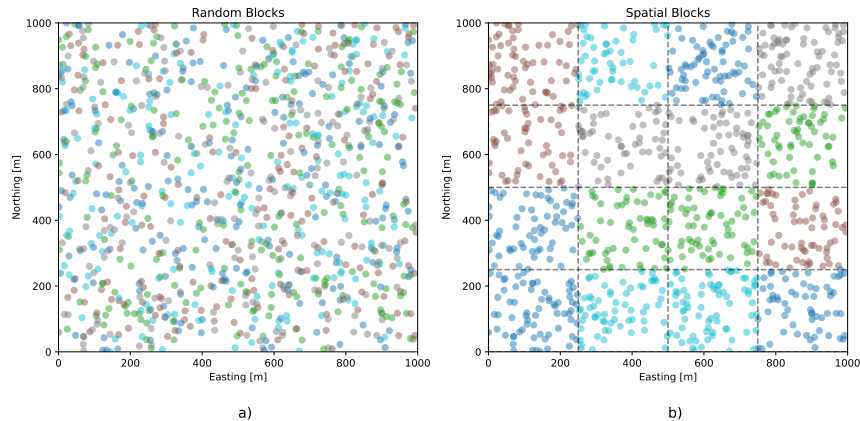


Fig. 3.11: a) Points randomly split into 5 folds, b) Points split into 5 folds using spatial blocks

Similarly to the semi-variogram and the MLE, we can use this method to estimate parameters using an optimizer. The results can be seen in Fig. 3.12, where we can observe how the parameters move towards the best estimate, noting that there is some backtracking. The algorithm estimates the noise variance parameter τ^2 as well, but this is not visualized. The surface is plotted using the final estimate for τ^2 .

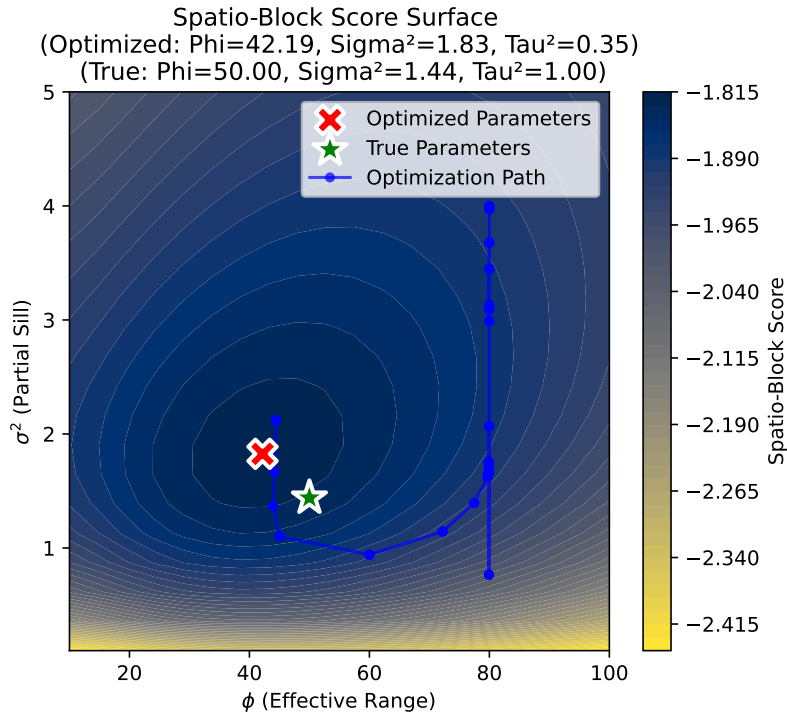


Fig. 3.12: The optimization path for the parameters for the CV. Note that the surface is plotted using the final estimated measurement variance parameter τ^2 .

For both MLE and CV, we have visualized the optimization surface for two variables; however, this is usually not available. Furthermore, there are often many more parameters. Thus, the methods can be harder to visualize. This demonstrates one of the advantages of the variogram: it is easy to interpret and compare with the data. However, with more complex covariance structures, it is not as easy.

Further remarks on parameter estimation

Here, three different approaches to parameter estimation are presented, each with its own advantages and drawbacks. A mixed approach can be used for different parameters in the model. Other common approaches that could be worth

using are Markov chain Monte-Carlo (MCMC), bootstrap methods (García-Soidán et al., 2014) - which can also be used for the methods described in this section to get some intervals for the estimated parameters, or integrated nested Laplacian approximation (Bakka et al., 2018) to name a few. The methods in this section are illustrated using a simple example, their performance in a real case will depend on how complex the model is and on the quality of the available data.

The quality of the data and where to get it are crucial considerations. To estimate parameters, one must rely on some type of data. Because the ocean is highly dynamic, obtaining data that match well to the specific target can be challenging. There are many different options discussed in the background Chapter 2, but usually only a few are available. For certain variables, satellite data will be widely available. There can be challenges with resolution and lack of depth perception, or it may be a different type of sensor than what is on the vehicles, but due to the availability, this is worth trying.

Numerical simulations can also play an important role. Even when they are unable to provide precise predictions because of uncertainties in the initial conditions, they can still reproduce key features of the system that are useful for estimating parameters.

Previous surveys in the region can also be used. It is not always easy to obtain the data used, but sometimes the parameters are published and can be adapted to different models. It is important to be cautious if the other survey is using different sensor types and that the properties of the ocean can be highly seasonal. Therefore, previous surveys should be a starting point but should be adjusted when there is sufficient new data.

Perhaps the best method is to estimate parameters using data from the current survey, either automatically on the fly by the vehicle or between deployments. When parameter estimation is done on the fly, some issues can arise. The computation onboard is often limited, and parameter estimation is one of the more computationally expensive tasks, which makes it challenging to achieve using edge computing resources. Another issue is that with the new parameters, the model will have to be estimated again, for example, using the sequential equations as shown by Chevalier et al. (2013); depending on the number of observed locations, this can be costly. Perhaps the largest issue is that the parameters should not be re-fitted until there is enough data available; however, this is harder to program and predict beforehand. In addition, several of the algorithms discussed here can diverge or create unreasonable outcomes, and it is challenging to implement sufficiently robust algorithms. Another approach is to send data from the vehicles during operation and then send the tuned parameters back to the vehicle. This can also be improved by combining sensor measurements from multiple vehicles. The common feature of all these approaches is that the algorithms for parameter estimation must be in place before the survey begins and must be tested to ensure that the parameters can be estimated accurately without too many issues.

In conclusion, good parameter estimation is important to make the models as accurate as possible, and there are many methods and data sources to achieve

this goal. However, the best method varies depending on the model and data availability for the specific survey.

Chapter 4

Path planning and acquisition function

How shall the question "Where does the river stop and the sea begin?", be answered?

Wassily Kandinsky
Point and Line to Plane

With limited time and resources for sampling, one needs to prioritize where and when to focus sampling efforts. It is hard to know, before the vehicles are deployed, exactly where sampling efforts should be focused. This is especially difficult if the feature that is interesting, such as ocean fronts, is moving throughout the experiment. With sensors, onboard computing, and the ability to move, we can design intelligent agents that make decisions on where to sample. Such agents can respond dynamically to incoming measurements and select subsequent sampling locations accordingly, a strategy known as adaptive sampling. Achieving this behavior requires methods for automated decision making and path planning. Although we focus primarily on AUVs in this chapter, the techniques presented are not limited to AUVs and are broadly applicable to other autonomous sampling platforms.

4.1 Sense-Plan-Act

An approach to choosing how the vehicle should behave is to use the "Sense→Plan→Act" cycle (Siegel, 2003), this can be used as a data driven approach to sampling. This cycle can be seen in Fig. 4.1 for adaptive sampling. An important part here is the onboard ocean model, as described in Chapter 3, which is used for assimilation and prediction. However, that is only one part of the decision making loop. We also need some way to suggest feasible paths or

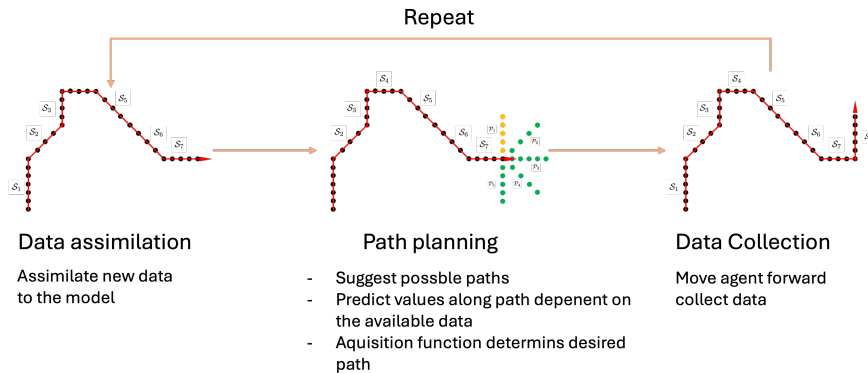


Fig. 4.1: Adaptive sampling loop. Each dot represents a measurement, the green dots are points in suggested paths, and the yellow dots are along the best path available.

designs $\mathbf{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_N)$, and then decide which path is the best path \mathcal{P}^* . The best path is the one that maximizes utility.

$$\text{BestPath} = \mathcal{P}^* = \arg \max_{\mathcal{P}_j \in \mathbf{P}} f(\mathcal{P}_j). \quad (4.1)$$

The function f is called the acquisition function; it is often also referred to as the utility function or the objective function (Garnett, 2023). This function attempts to find the path or sampling design $\mathcal{P}_j \in \mathbf{P}$ that would yield the most valuable samples. This acquisition function depends on the onboard model and what has been observed before. In the *act* part of the cycle, the agent will sample the path \mathcal{P}^* or parts of the path.

4.2 What is interesting? Where to sample?

With the model in place, we arrive at a central question: Where should we sample? If we revisit the true field in Fig. 3.4 b), it is not obvious which parts of the field are the most interesting. Here, one has to define some goals for the sampling; that is, features of the field that are considered valuable. A normal goal is to find the highest values in the field and, conversely, to find the lowest values in the field. Other times, like when we look at fronts, we want to separate the different water-masses, then we might be interested in sampling values close to a threshold. Another way to look at fronts is to find where the change in the field is the highest. These different considerations are illustrated in Figs. 4.2 and 4.3. The four regions illustrated are in different parts of the field; therefore, sampling efforts should differ depending on the goal. One can also observe that high gradient regions and the values close to the threshold are relatively close. The idea is to imagine that these are two ways of looking at a front.

Note that these illustrations only show the interesting regions for one variable; one could use different regions for different fields. To complicate it further, one could consider several types of interesting regions for each variable.

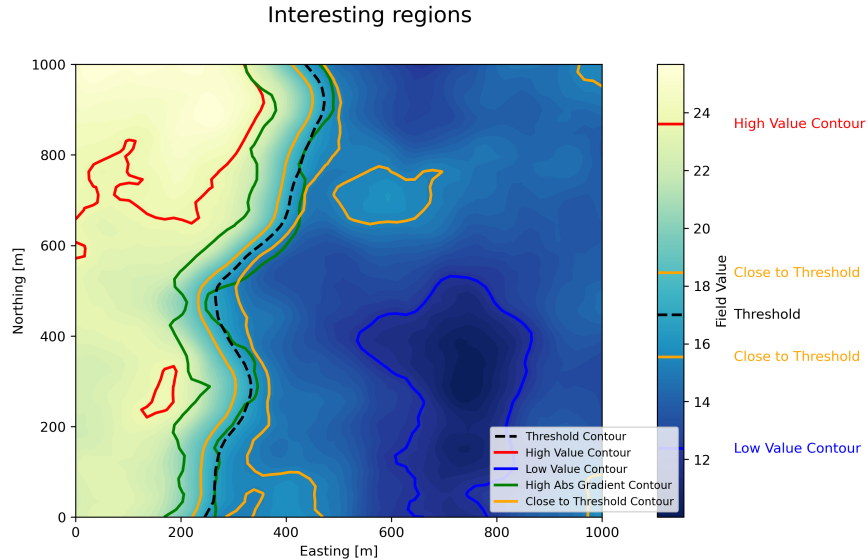


Fig. 4.2: The background field shows the true values, the colored contour lines shows the different interesting regions. The red line indicates where the highest values are, the blue line show where the lowest values are, the orange line shows values that are close to the threshold shown by the black line and green region shows where the change in field is high.

In this example, we pretend that we have knowledge of the whole field; in reality, this is never the case. Therefore, one needs to assess the value of sampling certain regions *before* the region is sampled, conditional on the values observed beforehand. The conditional model, as shown in Eq. (3.5), can be used to examine several sampling designs, and then we can use an acquisition function to determine what the best path is. Further in this chapter, we will show different acquisition functions that work for sampling different types of interesting regions.

4.2.1 Finding maximum values

One possible objective is to determine the maximum value of the field. To this end, we first focus on the value at a single point, denoted by $x(\mathbf{s})$. Under the assumption of a GRF, the conditional distribution of the unobserved variable can be expressed as

$$x(\mathbf{s}) \sim \mathcal{N}(m, v^2). \quad (4.2)$$

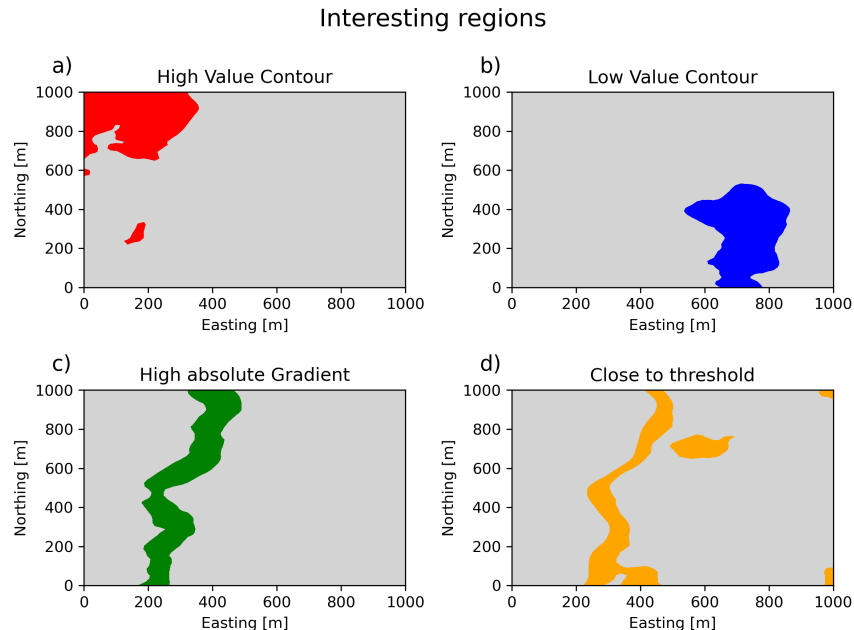


Fig. 4.3: Colored regions for the different interesting regions for the same field. a) 90th percentile field values, b) 10th percentile field values, c) where the field changes the most, and d) locations with field value closest to a threshold t .

With the goal of finding the maximum value in the field, we consider what the value of observing this location is. This obvious idea is to assign the value based on the predicted mean m . When considering multiple locations, the location with the highest mean is considered the best; we call this the maximum mean. This can work; the key weakness is that it can get stuck in local maxima and not explore potential new values. For the sake of illustration, we take another approach by going for the location with the highest uncertainty v . While this acquisition function does not target the maximum value, it does enable exploration and, therefore, can accidentally find the best value when exploring the field. A better acquisition function would target both where the predicted mean is high and explore the field. For this purpose, we use probability of improvement (PoI) and expected improvement (EI).

Probability of Improvement

Here we can define the improvement function $I(\cdot)$. This is defined as how much larger a new observation is compared with the best value observed before x_{\max} . This can then be written as

$$I(x(\mathbf{s})) = \max[0, x(\mathbf{s}) - x_{\max}]. \quad (4.3)$$

An observation is then valued if it improves on the previous best observation. An acquisition function can be defined as the probability that the improvement function is positive, this is then defined as

$$\text{PoI}(x(\mathbf{s})) = P(x(\mathbf{s}) \geq x_{\max}) = \Phi\left(\frac{m - x_{\max}}{v}\right). \quad (4.4)$$

Here Φ is the cumulative density function for the standard normal distribution. This does not take the magnitude of the improvement into account, just the probability of observing it.

Expected Improvement

EI is the expectation of the improvement function. Where the expectation is high, there is a good chance of observing a value that is better than we have observed before. EI is defined as

$$\text{EI}(x(\mathbf{s})) = \mathbb{E}[I(x(\mathbf{s}))] = \int_{-\infty}^{\infty} I(x) \cdot \pi(x; m, v^2) dx, \quad (4.5)$$

where $\pi(x; m, v^2)$ is the probability density function for a normal distribution with mean m and variance v^2 . We can define $z = (x - m)/v \sim \mathcal{N}(0, 1)$ as the standard normal distribution, with the corresponding probability density function $\phi(z)$. Then exchanging variables results in

$$\begin{aligned} \int_{-\infty}^{\infty} I(zv + m - x_{\max}) \cdot \phi(z) dz &= \underbrace{\int_{-\infty}^{(x_{\max}-m)/v} I(zv + m - x_{\max}) \cdot \phi(z) dz}_{=0} \\ &\quad + \int_{(x_{\max}-m)/v}^{\infty} I(zv + m - x_{\max}) \cdot \phi(z) dz \\ &= \int_{(x_{\max}-m)/v}^{\infty} (zv + m - x_{\max}) \cdot \phi(z) dz \\ &= v \int_{(x_{\max}-m)/v}^{\infty} z \cdot \phi(z) dz \\ &\quad + (m - x_{\max}) \underbrace{\int_{(x_{\max}-m)/v}^{\infty} \phi(z) dz}_{=\Phi((m-x_{\max})/v)}. \end{aligned}$$

The next trick depends on the fact that for a standard normal distribution $d/dz \phi(z) = z \cdot \phi(z)$, therefore

$$\int_{(x_{\max}-m)/v}^{\infty} z \cdot \phi(z) dz = \lim_{c \rightarrow \infty} [\phi(z)]_{(x_{\max}-m)/v}^c = 0 - \phi\left(\frac{x_{\max} - m}{v}\right). \quad (4.6)$$

When we put this together and let $\Phi(z)$ denote the cumulative density function of a standard normal, the EI can be written as

$$\text{EI}(x) = (m - x_{\max}) \cdot \Phi\left(\frac{m - x_{\max}}{v}\right) + v \cdot \phi\left(\frac{x_{\max} - m}{v}\right). \quad (4.7)$$

To illustrate the benefits of the 4 acquisition functions, we take one example where the previous best observation is $x_{\max} = 1$, and then we want to evaluate the value of observing a predicted variable $x \sim \mathcal{N}(m, v^2)$. In Fig. 4.4, we see the acquisition value for the four different acquisition functions: maximum mean, max variance, PoI, and EI when we change the predicted mean m and predicted variance v^2 . Firstly, we can look at the maximum predicted mean (Fig. 4.4 c); here, the maximum value observed before does not matter, only the maximum mean. The same applies to the variance in Fig. 4.4 d). The more interesting case is to compare and contrast EI and PoI. For PoI (Fig. 4.4 b), we only care about observing a value larger than x_{\max} ; therefore, the acquisition value increases with the uncertainty v when $m < x_{\max}$ and increases with v when $m > x_{\max}$. This can be a good quality when $x < x_{\max}$, it is better to look at an uncertain value when the next value probably will not be an improvement. However, when $x > x_{\max}$, this can be less than optimal; here, PoI prefers a certain small increase over an uncertain and potentially large increase. This changes with EI (Fig. 4.4 a); for values $m > x_{\max}$, EI prefers values with a higher uncertainty. This means that when looking for the maximum value of a GRF, EI will value exploration more than PoI.

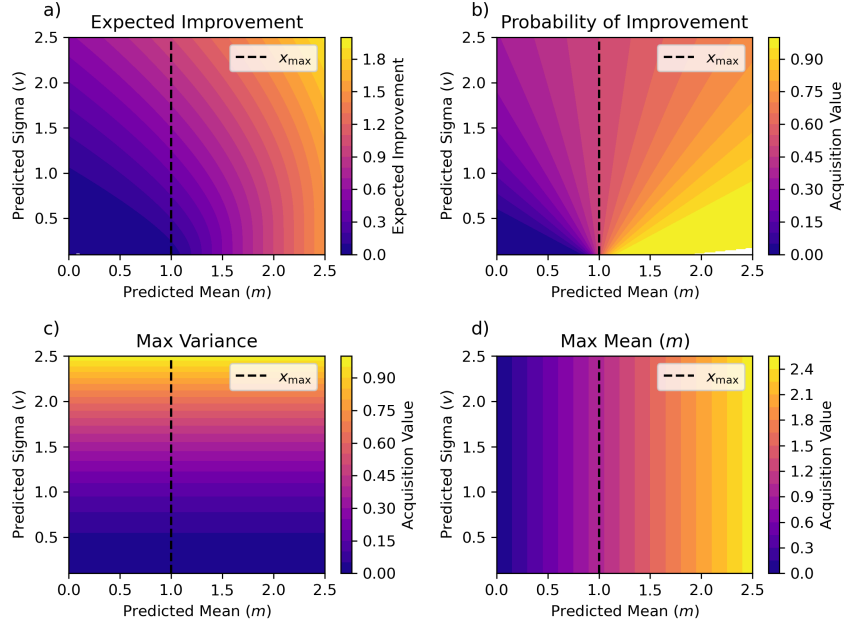


Fig. 4.4: Different acquisition functions for finding the maximum value. The bright yellow are considered the most valuable locations.

This can be illustrated a bit more with an example. In this example, we consider a 1D field $x(s)$ defined on $s \in \mathcal{D} = (0, 1000)$. We include 5 noisy observations and then compute the posterior mean and variance. The results are shown in Fig. 4.5 a). In the lower part of the figure (Fig. 4.5 b)), we can see a normalized acquisition value for the four functions based on the posterior mean and variance. The maximum variance gives the highest utility to the location with the highest uncertainty. Maximum mean and PoI are very similar; however, PoI favors the location with lower uncertainty a bit more. EI values the location close to the highest mean, but with more uncertainty.

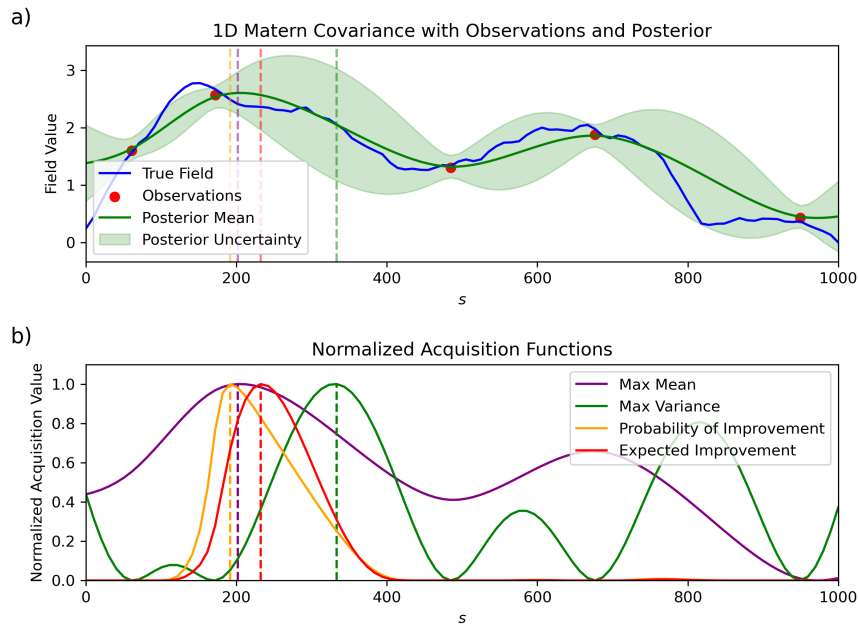


Fig. 4.5: a) Simulated field, observations and posterior field. b) shows the acquisition functions based on the posterior distribution.

This static example can be expanded into a small experiment where the acquisition functions are used iteratively to find the highest value of the field. This is done in a loop where we observe the value in the field with the highest utility, then we add the data point to the model and find the next best point to observe. This experiment is done 100 times with 10 iterations for each acquisition function. The result of this illustrative example can be seen in Fig. 4.6. The left plot shows one example for illustration; here we see the posterior mean for each of the acquisition functions, the points they observed, and the maximum value and location where it was observed. While each acquisition function did find a high value, the maximum mean acquisition function does get stuck. And while both EI and PoI did find the best value in the field, the points for EI are a bit more spread out. If we take a look at the right plot, we can see how the

four functions perform on average over 100 experiments. Here, it is clear that EI finds the highest values and finds them the quickest; As discussed earlier, maximum mean does better than maximum variance initially, but maximum variance overtakes it with more iterations. This is because maximum mean gets stuck in local maxima. Although the primary objective is to locate the maxi-




Fig. 4.6: Iterative simulation to find the highest value in the field. a) Example for one of the fields with 4 different acquisition functions. The dots shows which locations where samples and the lines are the posterior means, the true realizations is also shown. b) Mean normalized maximum value achieved at each iteration for different acquisition algorithms.

imum, one can also assess how accurately the field is approximated by computing the mean squared error (MSE) over 100 points in the field. Using the maximum mean criterion yields $MSE = 2.26$, maximum variance gives $MSE = 0.09$, PoI results in $MSE = 1.41$, and EI leads to $MSE = 0.73$. Thus, the maximum variance strategy provides the best approximation of the field, with EI as the next best. This simply serves to illustrate that different criteria emphasize different characteristics of the algorithm.

This little experiment aims to highlight some of the strengths of using EI as an acquisition function. This experiment also works as a primer for the full adaptive sampling strategies, which will be explained in more detail later, because in real cases, the agent cannot just sample any point in the field.

4.2.2 Classifying water-masses

In Section 2.1.1 the concept of an ocean front was discussed. That is, two water-masses that are differentiated based on some variable; a simulated front could be the orange region in Fig. 4.2. To separate the water-masses, we can define a threshold t and then aim to classify locations as above or below the threshold. Then we can define the excursion sets for the field as

$$ES = \{s \in \mathcal{D} : x(s) \leq t\}. \quad (4.8)$$

Repeating the example in Fig. 3.4, we can use it to define the prior excursion set, true excursion set, and posterior excursion set; this is shown in Fig. 4.7. The plot on the left shows the prior excursion set, the plot in the middle shows the true excursion set, and the plot on the right shows the posterior excursion set. The posterior excursion set is more accurate than the prior near where the AUV has sampled. In addition, we can look at the misclassifications set in Fig. 4.8. The left plot is the prior misclassification field, the right is the posterior, and the last plot shows where the classes changed. After sampling, the misclassification percentage is reduced, and importantly, the class changes happen close to where the AUV is sampling and the initial classification was wrong. The agent should sample where the potential to change the ES the highest.




Fig. 4.7: a) Excursion set for the prior mean, b) Excursion set for the true field and c) Posterior excursion set.

Excursion Set Misclassification of Prior and Posterior Mean, and Change in Classification




Fig. 4.8: a) the prior misclassifications field, b) the posterior classification field, and c) where the classification has changed. The misclassifications rate reduces after several locations are observed, and the misclassifications improves the most close to the observed locations.

Using the conditional mean and variance, we can assign the probability that

a location \mathbf{s} is in the excursion set, which is then defined as

$$P(x(\mathbf{s}) \in \text{ES}) = P(x(\mathbf{s}) \leq t) = p_{\mathbf{s}} = \Phi\left(\frac{t - m}{v}\right). \quad (4.9)$$

Here, the mean and variance are conditional on the previous observations' locations \mathcal{S} , but the notation is simplified to make it easier to read.

The more interesting locations are those where the probability is closer to 0.5 or where we are most uncertain about the classification. The Bernoulli variance is defined as $\text{BV}(\mathbf{s}) = p_{\mathbf{s}}(1 - p_{\mathbf{s}})$, and it is highest where the probability is equal to 0.5. This can be used to indicate the uncertainty for accurately classifying the field by integrating over all locations. This is called the integrated Bernoulli variance (IBV) and is defined as

$$\text{IBV} = \int_{\mathcal{D}} p_{\mathbf{s}}(1 - p_{\mathbf{s}}) d\mathbf{s}. \quad (4.10)$$

In practice, this will be computed discretely on a grid in the field $\mathbf{G} \subset \mathcal{D}$.

$$\text{IBV}_{\mathbf{G}} \approx \sum_{\mathbf{s} \in \mathbf{G}} p_{\mathbf{s}}(1 - p_{\mathbf{s}}) d\mathbf{s}. \quad (4.11)$$

The largest contributions to the IBV will come from the location where the BV is large, and a sampling strategy to reduce the IBV would improve the misclassification rate. In a sequential sampling design, we have observed the locations \mathcal{S} , and then we can choose between sampling different paths $\mathbf{P} = (\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_N)$. To reduce the IBV we choose the path that minimizes the expected integrated Bernoulli variance (EIBV) (Chevalier et al., 2014). The EIBV for a path \mathcal{P}_j can be defined as

$$\text{EIBV}(\mathcal{P}_j) = E \left[\int_{\mathcal{D}} \text{BV}(\mathbf{s}; \mathcal{P}_j) d\mathbf{s} \right]. \quad (4.12)$$

The expectation and integral can be exchanged, then we get the EBV for a point \mathbf{s} given that we sample a path \mathcal{P}_j as

$$\text{EBV}_{\mathcal{P}_j}(\mathbf{s}) = \Phi_2 \left(\begin{bmatrix} t \\ -t \end{bmatrix}; \begin{bmatrix} m(\mathbf{s}) \\ -m(\mathbf{s}) \end{bmatrix}, \mathbf{W}_{\mathcal{P}_j}(\mathbf{s}) \right), \quad (4.13)$$

where Φ_2 is the bi-variate Gaussian cumulative density function, and $m(\mathbf{s})$ is the posterior mean after sampling \mathcal{S} . The covariance matrix in the cumulative density function can be written as

$$\mathbf{W}_{\mathcal{P}_j}(\mathbf{s}) = \begin{bmatrix} \mathbf{V}_{\mathbf{s}|\mathcal{S}} & \mathbf{V}_{\mathbf{s}|\mathcal{S}} - \mathbf{V}_{\mathbf{s}|(\mathcal{S}, \mathcal{P}_j)} \\ \mathbf{V}_{\mathbf{s}|\mathcal{S}} - \mathbf{V}_{\mathbf{s}|(\mathcal{S}, \mathcal{P}_j)} & \mathbf{V}_{\mathbf{s}|\mathcal{S}} \end{bmatrix}, \quad (4.14)$$

where $\mathbf{V}_{\mathbf{s}|\mathcal{S}}$ is the conditional variance for $x(\mathbf{s})$ after sampling the locations \mathcal{S} from Eq. (3.5). $\mathbf{V}_{\mathbf{s}|(\mathcal{S}, \mathcal{P}_j)}$ is the conditional variance for $x(\mathbf{s})$ after sampling the locations \mathcal{S} and the proposed path \mathcal{P}_j . Note that the conditional variances are

not dependent on the data. Then, to approximate the EIBV, we compute the EBV for different locations in the field $\mathbf{s} \in \mathbf{G}$ as

$$\text{EIBV}(\mathcal{P}_j) \approx \sum_{\mathbf{s} \in \mathbf{G}} \text{EBV}_{\mathcal{P}_j}(\mathbf{s}). \quad (4.15)$$

We can illustrate how this looks with the example in Fig. 3.4. Consider that the AUV has already sampled the red dashed line, and now one can add one measurement anywhere in the field. In Fig. 4.9, the EIBV for the field is shown. The red line is the AUV path, the solid green line is the posterior threshold, and the dashed green line is the true threshold. The EIBV is lowest close to the posterior threshold and far from the AUV path, where the uncertainty is high. The best point to observe here, according to the EIBV, would be the dark red area.

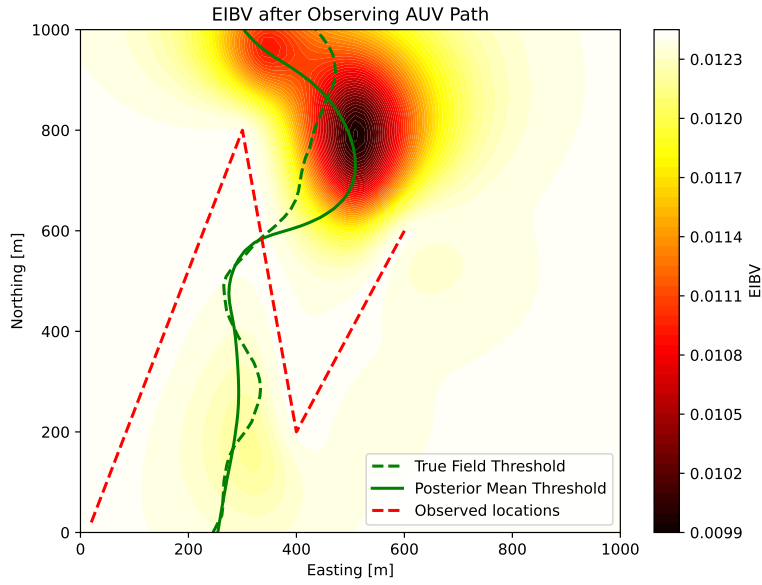


Fig. 4.9: The EIBV of the field. The red dashed line is the agents path, the green solid line is the posterior threshold, the green dashed line is the true threshold.

Here, the EIBV is illustrated for a 2D univariate field; this can be expanded to a 3D field, as in Ge et al. (2023), or with temporal effects Ge et al. (2025), or a bivariate field with two thresholds Fossum et al. (2021a).

4.2.3 Sampling high gradients

When working with ocean fronts, classifying something as above or below a certain threshold is a valuable approach, but it does have some weaknesses, especially if the threshold is not chosen properly. Another way to look at the fronts is to examine where the *change* in the field is large; therefore, we would like to use an acquisition function that enables sampling large gradients. Here we can define the directional gradient between two locations in the field as

$$g(\mathbf{s}_i, \mathbf{s}_j) = \frac{x(\mathbf{s}_i) - x(\mathbf{s}_j)}{d(\mathbf{s}_i, \mathbf{s}_j)}, \quad (4.16)$$

where $d(\mathbf{s}_i, \mathbf{s}_j)$ is defined as the lateral distance between two points. The distribution of the directional gradient $g(\mathbf{s}_i, \mathbf{s}_j)$ then becomes

$$\begin{aligned} g(\mathbf{s}_i, \mathbf{s}_j) &\sim \mathcal{N}(\zeta_{i,j}, \eta_{i,j}), \\ \zeta_{i,j} &= \frac{\mu(\mathbf{s}_i) - \mu(\mathbf{s}_j)}{d(\mathbf{s}_i, \mathbf{s}_j)}, \\ \eta_{i,j} &= \frac{1}{d(\mathbf{s}_i, \mathbf{s}_j)^2} (C(\mathbf{s}_i, \mathbf{s}_i) + C(\mathbf{s}_j, \mathbf{s}_j) - 2C(\mathbf{s}_i, \mathbf{s}_j)). \end{aligned} \quad (4.17)$$

For slightly simpler notation later in the text, we refer to $g \sim \mathcal{N}(\eta, \zeta)$.

With the goal of making the AUV sample large gradients, we need an acquisition function to evaluate what makes a potential design good to sample. Here we can reuse many of the ideas from the sampling of large values; the main difference here is that with directional gradients, both large positive and negative values are good. Looking for large absolute gradients or locations with large uncertainty is still a good baseline strategy. If we want to use better acquisition functions like PoI and EI, we need to make a few modifications. Firstly, we need to make some changes to the improvement function $I(\cdot)$ in Eq. (4.3); instead of just looking at the largest value observed before, we look at the largest absolute directional gradient observed g^* . The new improvement function becomes

$$I(g) = \max(|g| - g^*, 0). \quad (4.18)$$

To derive the equations for PoI we only need to change the limits to the integral a bit

$$\text{PoI}(g) = P(g \geq g^* \cup g \leq -g^*) = \Phi\left(\frac{\zeta - g^*}{\eta}\right) + \Phi\left(\frac{-\zeta - g^*}{\eta}\right). \quad (4.19)$$

A similar approach works for EI, this can be written as

$$\text{EI}(g) = \text{E}[I(g)] = \int_{-\infty}^{\infty} I(g) \cdot \pi(g; \eta, \zeta^2) dg, \quad (4.20)$$

where $\pi(g; \eta, \zeta^2)$ is the probability density function for a normal distribution with mean η and variance ζ^2 . The integral can be decomposed into three separate regions

$$\text{EI}(g) = \int_{-\infty}^{-g^*} I(g) \pi(g) dg + \underbrace{\int_{-g^*}^{g^*} I(g) \pi(g) dg}_{=0} + \int_{g^*}^{\infty} I(g) \pi(g) dg. \quad (4.21)$$

By following the same derivation steps used to obtain Eq. (4.7), the EI for the directional gradients is given by

$$\begin{aligned} \text{EI}(g) &= (\zeta - g^*) \left(1 - \Phi \left(\frac{g - \zeta}{\eta} \right) \right) + \eta \phi \left(\frac{g^* - \zeta}{\eta} \right) \\ &+ (-\zeta - g^*) \Phi \left(\frac{-g^* - \zeta}{\eta} \right) + \eta \phi \left(\frac{-g^* - \zeta}{\eta} \right). \end{aligned} \quad (4.22)$$

This can be used as an acquisition function to guide the agent towards large lateral gradients. Similarly to the small example shown in Fig. 4.6 using EI often balances exploration and exploitation quite well.

4.3 Path planning

Path planning is an important field of robotics research. This encompasses many branches and is widely researched regarding AUVs (Zeng et al., 2015). Some focus on the act of going from point A to point B on a local scale; that is, controlling actuators and propellers while monitoring the positioning to get to the desired location. Others look at larger scales, trying to find paths that navigate obstacles or find the most energy efficient approaches (Tsardoulis et al., 2016). Often, the vehicles are not alone; this can lead to path planning that avoids collisions (Vagale et al., 2021) or other risks (Bremnes et al., 2024), or it can be a more cooperative approach between autonomous vehicles (Wu et al., 2025). The type of path planning we are mainly interested in is informative path planning, that is finding the path that potentially gives the most information, the idea is roughly explained in Eq. (4.1). Now we will go into more detail on how to design these paths and how they should be evaluated.

The simplest approach is to use a myopic strategy or a greedy strategy. This approach evaluates a list of paths, samples the best path, and repeats. This approach only looks at alternative paths close to the vehicle. In this simple approach, we do have to take the maneuverability of the vehicle into account and make some choices about how many paths should be considered and their length. Increasing the number of paths gives more options to optimize, but at the cost of computation; however, that is usually not a big problem in the myopic approach. The choice of how long the paths should be can affect how reactive the agent is; with smaller steps, the agent can quickly react to changes, but it does

not look very far ahead. An illustration of different path planning methods can be found in Fig. 4.10; in the top left, we can see the myopic path planning. To improve myopic path planning, we can employ a strategy by linearly extending the paths evaluated, but only traveling a part of it before re-evaluating the path. This can be called a look-ahead strategy. This is illustrated in Fig. 4.10 b). A look-ahead enables the agent to look a bit further ahead without many of the computational challenges that come from non-myopic path planning.

The topic of non-myopic planning involves planning and evaluating multiple steps into the future and constantly reevaluating the plan. There are several reasons why one might want to implement such a strategy. Sometimes, focusing on short term goals comes at the cost of long term rewards, resulting in unbalanced sampling or even missing important locations entirely. There are also several non-informative reasons to use long term path planning. When navigating obstacles or more complex environments while gathering valuable information, non-myopic path-planning can be valuable. Lastly, the survey might include several prescriptive goals that need to be accomplished. This could include a specific location and time (Ge et al., 2025), surveying several intermittent goal locations (Janoš et al., 2021), or following a moving target (Pinto et al., 2013). Furthermore, we will look at some non-myopic algorithms and approaches.

Firstly, a non-myopic strategy can be the same as running a myopic strategy but with a multi-step search. We write a multi-step path as $\mathbf{P}_j = (\mathcal{P}_{j_1}, \mathcal{P}_{j_2} \dots, \mathcal{P}_{j_k})$, then we can choose the best multi-step path according to the acquisition function $f(\cdot)$ as

$$\mathbf{P}^* = \arg \max_{\mathbf{P}_j \in \mathcal{P}} f(\mathbf{P}_j). \quad (4.23)$$

When the best k -step path is chosen, the agent moves along the first part of the path $\mathcal{P}_{j_1}^*$. We can see that the number of paths can easily explode. Imagine that at each step, an agent can choose 8 lateral directions and 3 in depth, with a total of 24 paths. If we plan 3 steps deep, this results in 13824 different paths and can be infeasible to compute. This is also made harder because the information value from path segments cannot be considered independent. Tricks might have to be employed to reduce the number of paths evaluated. However, with a low k , this approach is still very local and can, in some cases, struggle to beat the myopic strategy (Fossum et al., 2021b).

One of the supposed issues with the myopic strategy is that it does not necessarily move towards the global best point but focuses on local goals. To do this, instead of evaluating the paths, we can find the best location to sample

$$\mathbf{s}^* = \arg \max_{\mathbf{s}_j \in \mathcal{S}_G} f(\mathbf{s}_j),$$

and then find a path from the current location \mathbf{a} to this goal. We write a feasible path from \mathbf{a} to \mathbf{s}^* as a path consisting of several segments $\mathbf{P}(\mathbf{a}, \mathbf{s}^*) = (\mathcal{P}_1, \dots, \mathcal{P}_k)$, then find the shortest feasible path

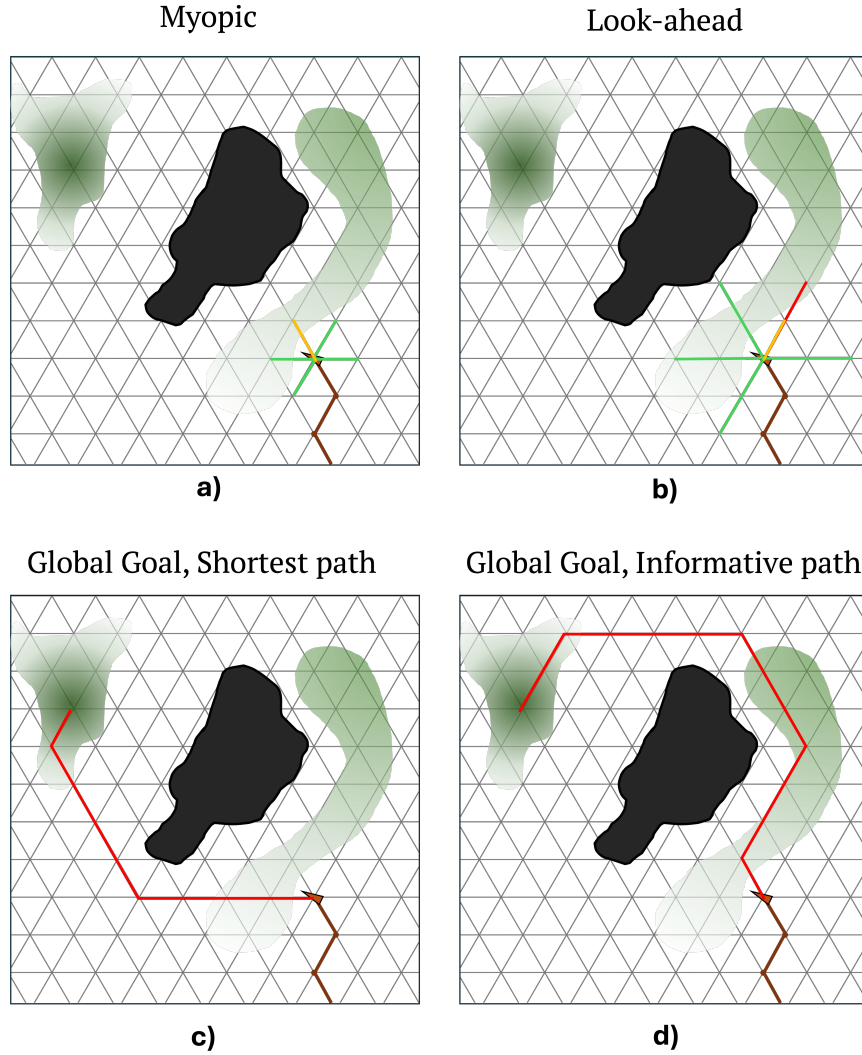


Fig. 4.10: Different scenarios for informative path planning. The grid is a waypoint graph for the agent to move on, the green is the information field where darker is more informative, the dark red path is the previous path, the bright red best path found, bright green are suggested paths and the black blob is an obstacle. **a)** myopic planning, **b)** look-ahead planning, **c)** global goal with shortest path planning, **d)** global goal with an informative path. All agents move to the first waypoint in the best path, and then reevaluate.

$$P_{\text{short}}^*(\mathbf{a}, \mathbf{s}^*) = \arg \min_{P(\mathbf{a}, \mathbf{s}^*)} \sum_{P_i \in P_j} C_d(P_i), \quad (4.24)$$

where C_d is just the distance of a path segment. This method can be illustrated in Fig. 4.10 c); here, the agent finds the shortest path to the end goal, but we can see that this is not necessarily the most informative path. Finding the shortest path in Eq. (4.24) is a well studied problem, and several algorithms can be used, like A^* (Kumawat et al., 2021). or a goal oriented rapidly exploring random tree (RRT*) (Noreen et al., 2016). Similar to the other approaches, we need not go the full distance of the path, but rather just a part and then reevaluate. An issue with this approach is that we do not compute the most informative path. A short path is good for reaching the goal, but the agent might lose out on paths that are a bit longer but much more informative.

To reduce the chance of missing important locations on the way, we can introduce a path cost for going to uninformative locations. In Eq. (4.24) we minimize the distance; however, we could include more terms in the cost of the path instead. Therefore, we introduce an informative cost function $C_I(\mathcal{P})$. This information cost can be based on one or several of the criteria discussed in Sec. 4.2 with some modifications so that informative locations correspond with a low path cost. Thus, the full cost of the path will be

$$C(\mathcal{P}) = C_d(\mathcal{P}) + C_I(\mathcal{P}). \quad (4.25)$$

We can further break the informative cost into several functions with associated weights

$$C_I(\mathcal{P}) = \sum C_{I_i}(\mathcal{P}) \cdot w_i, \quad (4.26)$$

where w_i is the weight for the corresponding informative cost. This function can be viewed as a cost valley, where it is cheaper to walk down in the informative parts of the topography. Using this informative function, we can define the best path to the goal location as

$$\mathbf{P}^*(\mathbf{a}, \mathbf{s}^*) = \arg \min_{\mathbf{P}(\mathbf{a}, \mathbf{s}^*)} \sum_{\mathcal{P}_i \in \mathcal{P}_j} C(\mathcal{P}_i). \quad (4.27)$$

An illustration of this is shown in Fig. 4.10 d); here we see that the endpoint of the path is at the global best location, but the path is not the shortest; rather, it is a more informative path. This is clearly an idealized case, but it illustrates how the agent can maneuver around obstacles towards a global goal while gathering valuable samples.

Finding the optimal path can be infeasible, especially within a reasonable time frame. However, a feasible path can be found. Effective algorithms will therefore not find the best path but rather a good enough path given the available time. An implementation of this cost valley with a RRT* path planner can be found in Ge et al. (2025). In this paper, obstacles are defined, an end point is specified, and a fixed time to complete the sampling is established.

Lastly, there is the much more difficult problem of finding the most informative path. Here we consider an agent who has a certain budget B for a survey. This involves solving the following equation

$$\mathbf{P}^* = \arg \max_{\mathbf{P}} f(\mathbf{P}), \quad C_d(\mathbf{P}) \leq B. \quad (4.28)$$

Note that a path consists of several sub paths, but the information value of each sub path is dependent on the others. With a dynamic field, this has to be recomputed often and will quickly become very expensive.

In this section, several ways of suggesting paths for sampling are shown. It seems that the main drawback of the non-myopic methods is the associated computational cost, but this is not the full picture. The effectiveness of the non-myopic path planners greatly depends on how well the model can predict regions that are far away in space and time. If it cannot predict faraway regions well, then planning toward those locations may not make a lot of sense and may fail to outperform myopic methods; it could even perform worse. In addition, some of the algorithms may not converge in a reasonable time, and the resulting paths may be suboptimal.

Chapter 5

Experiments

‘I forget, did you ever get any essays done at all?’
‘Well, not as such. No actual essays, but the reasons why not were absolutely fascinating. For instance, I discovered that...’

Douglas Adams
Dirk Gently’s Holistic Detective Agency

In the previous chapter, we discussed what could be considered valuable to sample, how to generate paths, and how to evaluate what the best path is. For each problem, there are several parameters, acquisition functions, and algorithms to test. So, this section will include some discussion on how these algorithms are evaluated and then implemented for a field experiment.

5.1 Simulation study

Before we even go into the ocean, we need to decide which algorithms to fly. Usually, there is a limited number of hours allocated in the field, but a nearly unlimited number of hours of simulation time. In addition, we will never have access to the ground truth in a real experiment, but in simulation, we have access to a perfect ground truth. Therefore, most of the algorithm selection and tuning takes place in simulation. The general setup for the simulation study is to simulate a diverse set of fields. Then, for each field, test a set of algorithms with different parameters and run some evaluations to see which performs the best.

Evaluation Metrics

Here we will list a number of metrics that can be valuable for different applications. For each metric, one can find the mean value and the variance. It can also be valuable to investigate the worst case scenarios or the agents in the fifth percentile. For some algorithms, the average performance might be good, but there is a significant chance of poor performance. These metrics will usually also be computed for different time steps in the simulation, thus evaluating how they evolve with time. To properly evaluate different algorithms, a range of evaluation metrics should be used and contrasted.

Here we denote $x_{\text{sim}}(\mathbf{s})$ as the true value in the field, $m(\mathbf{s}|\mathbf{y}_S)$ is the estimate of x conditional on the observations \mathbf{y}_S . The first metric is just the mean squared error (MSE) over the field \mathcal{S}_G

$$\text{MSE} = \frac{1}{N_G} \sum_{\mathbf{s} \in \mathcal{S}_G} (x_{\text{sim}}(\mathbf{s}) - m(\mathbf{s}|\mathbf{y}_S))^2. \quad (5.1)$$

The goal of this metric is to evaluate how well the agents learn the field. In many cases, we are only interested in specific parts of the field. Then, one can consider a targeted version of the MSE; for example, if we care about finding the largest value in the field, we might only compute the MSE for field values above a threshold $x_{\text{sim}}(\mathbf{s}) > t$. Other similar targeted MSEs can be calculated for the threshold and large gradients. The second is related to the uncertainty of the field. The mean variance (MV) can be written as

$$\text{MV} = \frac{1}{N_G} \sum_{\mathbf{s} \in \mathcal{S}_G} v(\mathbf{s}|\mathbf{y}_S). \quad (5.2)$$

Similarly to MSE, targeted versions can be applied for high values or values close to a threshold.

These are some of the more general metrics, but for each application, more specific metrics can be used. For example, when finding the boundary, a misclassification error can be computed, this can be seen in Fig. 4.8. When looking for the highest value, it can be shown at each time, similar to what is shown in Fig. 4.5. Each evaluation metric should evaluate whether the agent achieves an aspect of the overall goal.

Although simulation is extremely useful and, in practice, the only way to perform a sufficient amount of testing, there will always be a substantial gap between simulated conditions and the real world. One key discrepancy concerns how the agent moves in relation to the ocean. Currents, trim, and other environmental factors can cause the agent's motion to deviate from the desired path. Furthermore, underwater vehicles rely on an estimated position, and after spending an extended period submerged, this estimate can become inaccurate, potentially causing problems for the model. The simulated field will also differ greatly from the true field. This difference can, in some cases, lead to incorrect conclusions about the best algorithms. Lastly, and perhaps the most difficult hurdle, is that several of the adaptive algorithms react to the field, sometimes

with some randomness. This means that there is usually no assurance that an adaptive algorithm will perform better than a prescriptive algorithm under every circumstance. Furthermore, with unforeseen conditions and incorrect setups, the agent might not behave as expected based on the simulation study.

5.2 Field experiment

A simulation can never replace a field experiment; after all, the field experiment is the whole point. This section goes through the field experiments that were conducted for this thesis. It should be noted that not every experiment was a publishable success, but each trial was a learning experience. Many of the lessons were probably obvious to people who have spent some time in the field, but for someone on the more theoretical side, they were very valuable. In addition, they can be motivating for the work and foster collaboration.

There is a considerable gap between a sampling strategy that works in simulation and one that works in the field. Building up from a simulation to a field experiment needs to be done carefully. In addition to the informative parts of the agent, we also need safety measures, logic for communication, how robust the system is, and how easy it is to debug. These need to be implemented; often, things go wrong in unforeseen ways, and then the logs need to be detailed enough to retrace the issue and solve it. It is also important to test as closely to real conditions as possible, and often a simulation mode is added to vehicles that simulates the dynamics of the vehicle. In this mode, communication between systems and algorithm logic can be tested without leaving the lab. This does not reveal every issue.

Early on in our work, we found that when we launched a process from a laptop and then the vehicle lost contact with the laptop, it would eventually interrupt the process and stop it. This happened once, of course, on the last possible day of testing in Porto, where previous days had been lost to bad weather, so we returned home with no publishable results. The fix was very easy, but the issue was not revealed in the simulation and was tricky to figure out with the limited time in the field. We were able to replicate the issue in the lab. and then disconnecting the vehicle from the laptops (and WiFi) became part of the testing procedure. This is just one such story, but there are many similar ones. Some of the issues encountered were very familiar to people who have worked on such systems, such as the one mentioned, and some were very specific quirks associated with the vehicles with which we worked. These are things you only learn when leaving the comforts of the office.

We cannot always determine the effectiveness of an algorithm through a field experiment. With the setup of running a simulation study to choose the best algorithm and then using it in the field, we still have issues evaluating the best algorithm for the field. A simulation study can show clear benefits and drawbacks to different sampling strategies for different scenarios; however, this is not available in the field, and we have to assume that properties from the simulation carry over to the field. The inherent randomness also means that

adaptive strategies might not be strictly better than non-adaptive strategies, but rather better on average. This does mean that it can be hard for researchers to rely on adaptive algorithms. In addition, the sampling paths of the agents can be more difficult to explain. When there is no ground truth, it is hard to conclude that the chosen paths were the optimal choice. Over time, one can run adaptive and pre-scripted plans in parallel and then contrast the results. Eventually, patterns will emerge, but usually, the resources to do this are not available. There are other reasons why one might prefer adaptive strategies; with larger and more complex operations, where the number of agents outnumbers the human operators, pre-scripted plans are infeasible. In addition, if the durations of the missions are long enough and without human intervention for a long time, then more autonomy needs to be utilized. Developing data driven sampling strategies is a vital part of this autonomy.

5.3 Experiments timeline

During the work on this thesis, several field experiments were conducted. The timeline for this is shown in Tab. 5.1. This table shows the location, time, a small description, and the associated paper. The experiments included in this thesis are usually part of larger field tests, but only the relevant parts are described here. The experiments for Papers A and C are associated with the MASCOT project, while Papers B and D are associated with the HARVEST project.

Paper	Location	Year	Month	Description
-	Porto	2021	Oct.	Deployment in the Douro river. Using a 3D adaptive model to investigate the river plume
C*	Trondheim	2022	May	Using AUV to explore Nidelva river plume. Testing 2D Long-Horizon informative path planning for finding and classifying river water masses. Due to a bug, the AUV stope every 15 min, and had to be found and restarted. The mission finished, but due to the constant stopping and subsequent drifting the results were not usable.
-	Porto	2022	Oct.	Long-Horizon path planning in Douro river, similar to Nidelva a few months before. Due to bad weather (fog, large waves and wind), something called "fish attack" and short time in the field, the adaptive missions were very short and not usable. These experiments were re-done remotely at a later date.
C	Trondheim	2022	Sep.	Successfully running the Long-Horizon informative path planning from A to B.
A	Trondheim	2023	Jun.	Using AUV to find large salinity gradients in the Nidelva river. The mission was successful, but during another mission the AUV was lost. It was found 4 days later and 40 km away from where it was lost.
B	Mausund	2024	Jun.	Beyond-the-horizon experiments with several autonomous platforms, the majority of the scientific staff was in a control room 100 km away remotely controlling the vehicles. As a part of this campaign an adaptive algorithms for finding large chl <i>a</i> values was tested.
A	Trondheim	2024	Sep.	Replicate experiment for Paper A .
D	Mausund	2025	May	Second run of the beyond-the-horizon experiments in Mausund. As a part of this the goal was to fuse and compare sensor data from different sensors and platforms. In addition adaptive sampling using the particle camera was tested, but not successful.

Table 5.1: Overview of the field deployments during. The star marks means that we only got a partial result and could not be published.

Chapter 6

Paper overview

6.1 Paper A

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

André Julius Hovd Olaisen, Jo Eidsvik, Yaolin Ge

Published in *Environmetrics*, Volume 36, issue 6, 14pp , 2025 (Olaisen et al., 2025)

This work has been presented in the *22nd Annual Conference of the International Association for Mathematical Geosciences, 2023*, and in the *21st Norwegian Statisticians Meeting, 2024*.

Summary: The aim of this study is to detect and sample large lateral salinity gradients in a river plume using an AUV. The salinity field in the river plume is modeled as a spatio-temporal GRF, which implies that the directional gradients are Gaussian distributed as well. To enable real-time data assimilation and fast onboard predictions, we use a grid-free model, efficient matrix factorization techniques, and spatio-temporal thinning of observations that are far away from the current location. The sequential sampling design is based on proposing new sampling paths with a spider-leg shape. EI for the salinity gradients is used as an acquisition function to select the most informative paths. The design is myopic, but each proposed path includes a look-ahead extension beyond the portion that is actually sampled. Parameters in the sampling design and acquisition functions are explored in a simulation study to compare performance. In the simulated scenario, we test the ability to sample large gradients and to identify different regions with large salinity changes. Here, EI works well for finding large gradients and exploring potential new areas. Field experiments were conducted in the river plume of the Nidelva River in the Trondheimsfjord (June 2023 and September 2024). Each field experiment lasted approximately 2 hours and was successfully conducted without human intervention.

Key contributions:

- EI for selecting paths that target large lateral salinity gradients.
- Efficient onboard data assimilation and prediction (matrix factorization + spatio-temporal thinning).
- A low-cost look-ahead extension that improves on a purely myopic planner.

6.2 Paper B

Autonomous Underwater Vehicle Sampling for Hotspots in Chlorophyll A

André Julius Hovd Olaisen, Jo Eidsvik

Under revision in *Autonomous Robots*

Preliminary results were presented at the *21st Norwegian Statisticians Meeting in 2024*

Summary: The goal of this project is to effectively find and map large concentrations of phytoplankton in space and time. chl a is used as a proxy for phytoplankton abundance throughout this work. The chl a field is modeled using a spatio-temporal log-GRF, with a prior mean for the vertical distribution of chl a . The model is implemented using a grid-free sparse representation and thinning of locations far away in space and time. The sampling algorithm operates in two modes: when the AUV is at the surface, it suggests lateral directions with endpoint goal locations; in the second mode, it moves toward the end location and sequentially chooses depths to sample along the way. When the AUV returns to the surface at the endpoint, it starts the process again. The motivation for this design is to avoid unnecessary underwater turning. The proposed paths are assessed using an acquisition function. The algorithm, together with several alternative acquisition functions, is evaluated in a simulation study. Here, we test the acquisition functions' ability to find and explore different hotspots, as well as their ability to identify the largest concentrations of chl a . In simulation, EI scores well on both metrics compared with the other acquisition functions tested. The sampling design was tested as part of a larger campaign in Frohavet in 2024. The AUV ran the adaptive mission twice during the campaign and found depths with the highest chl a . During these missions, the AUV carried a particle camera enabling detection of zooplankton, along with sensors such as salinity and temperature. Results from these other sensors are also discussed in connection with each other.

Key contributions:

- EI-driven hotspot sampling for chl a with an AUV.
- A two-mode strategy (surface lateral planning + underwater depth selection) to reduce turning and onboard computation.
- Cross-sensor comparison of zooplankton counts (particle camera) with other in situ sensors.

6.3 Paper C

RRT*-enhanced long-horizon path planning for AUV adaptive sampling using a cost valley

Yaolin Ge, Jo Eidsvik, André Julius Hovd Olaisen

Published in *Knowledge-Based Systems, Volume 315:113261, 2025* (Ge et al., 2025)

An earlier simulation based version of this paper was presented at the *14th IFAC Conference on Control Applications in Marine Systems, Robotics, and Vehicles CAMS 2022* (Ge et al., 2022)

Summary: We investigate long-horizon informative path planning for sampling a river plume. Moving beyond a myopic planner is computationally demanding because it must anticipate future observations rather than react greedily. Our approach first identifies the globally optimal sampling region in the field and then constructs an informative trajectory to that region. This is achieved by building a cost-valley, a weighted cost surface that steers the vehicle toward highly informative locations. The cost surface combines two information criteria: integrated variance reduction (IVR), which reduces uncertainty over the entire field, and expected integrated Bernoulli variance (EIBV), which targets uncertainty along the plume boundary. We evaluate several weightings of these criteria. The cost-valley is explored using the RRT* algorithm. This is still a relatively computationally expensive approach, and several tricks had to be implemented to make it run in real-time. We assess the planner and various weightings for the information criteria in the cost-valley in a simulation study and compare the results with a greedy (myopic) planner. Performance metrics include mean-squared error, IBV, and traffic-flow analysis. The long-horizon planner yields marginally better results than the myopic baseline. Importantly, the long-horizon planner can also satisfy additional mission constraints, such as reaching a prescribed endpoint within a fixed time budget. The method was validated in field experiments conducted in the Nidelva river plume (Trondheim fjord). The autonomous underwater vehicle successfully navigated from the start location to the designated end location within the allotted 2.5 h while sampling the plume.

Key contributions:

- The informative cost-valley formulation (multi-criterion weighting of IVR and EIBV).
- An RRT*-based implementation enabling non-myopic planning with real-time constraints.
- Guaranteed satisfaction of mission constraints (arriving at a prescribed endpoint at a prescribed time).

6.4 Paper D

Spatio-temporal Calanus Mapping Using Different Vehicles and Sensors in a Campaign off the Norwegian Coast

André Julius Hovd Olaisen, Ahmed Abdelgayed, Jo Eidsvik, Erik Liu,
Karoline Barstein, Martin Ludvigsen

This work has been presented internally in the research group and is not yet submitted to any journals.

Summary: In this work, we investigate the use of heterogeneous autonomous marine vehicles and different sensors to map the distribution of phytoplankton and zooplankton. The measurements were obtained during a field campaign conducted during the spring bloom of zooplankton in late May 2025 in Frohavet. The campaign involved two different types of ASVs, two AUVs, and data collected from a boat. The different vehicles rarely sampled the same location at the same time, or, in the case of the AUVs, even at the same depth. This makes it challenging to compare measurements from different vehicles and sensors. We use different spatio-temporal models for the various sensor measurements; in particular, we are interested in the upper and lower biomass layers in the water column and the integrated biomass in these layers. These water-column statistics are compared across platforms and sensor types. We compare the zooplankton biomass estimates from the optical particle camera on the AUVs, echo-sounder on the ASVs, and the phytoplankton biomass estimates from the chl *a* sensor on the AUVs. We find a correlation between the zooplankton biomass estimates from the acoustic sensors on the ASV and the biomass estimates derived from the optical particle sensors on the AUV. This suggests that both types of sensors could be used to investigate zooplankton in a similar way, despite the fact that the sensors operate very differently. We also compare the phytoplankton biomass estimates from the chl *a* sensor on the AUV with the different zooplankton biomass estimates from the AUV, ASV, and the samples from the boat. In this case, we find no correlation or a negative correlation across all platforms.

Key contributions:

- Cross-platform comparison of heterogeneous sensors/vehicles despite sparse and asynchronous sampling.
- Sensor-specific spatio-temporal models, including a Poisson latent GRF for optical zooplankton observations to generate depth profiles.
- Evidence of correlation between optical particle-camera biomass estimates (AUV) and acoustic echo-sounder biomass estimates (ASV).

Chapter 7

Conclusion and Further Work

“What! you have solved it already?” “Well, that would be too much to say. I have discovered a suggestive fact, that is all.”

Sir Arthur Conan Doyle
The Sign of Four (Sherlock
Holmes)

7.1 Closing remarks

The main goal of the work has been to use statistical models together with data driven path planning algorithms to improve the sampling design for autonomous marine vehicles. The two primary phenomena investigated have been observing ocean fronts in the form of river plumes and finding large biomasses of plankton. The algorithms were tested in several simulation studies and, crucially, in field experiments. This means that for each phenomenon, we have worked from the demands of the problem and the limitations and strengths of the vehicle to build adaptive agents that were successfully tested in the field. The key contributions can be summed up as:

- Investigation of two key problems using autonomous marine vehicles; detecting ocean fronts in the form of river plumes and finding high concentrations of both phytoplankton and zooplankton.
- Using GRFs to model different spatio-temporal fields, such as salinity in river plumes, chlorophyll *a* using a log-normal GRF, and zooplankton counts using a Poisson latent GRF. This provides practical ways to reduce

computational costs while maintaining accurate predictions. Significant effort has also been made to reduce the computational cost of the models, such as efficient matrix factorization and spatio-temporal thinning.

- Exploring different acquisition functions for sequential sampling designs for an autonomous agent in spatio-temporal fields. This includes EI for directional salinity gradients and high densities of chl a , as well as EIBV for salinity thresholds.
- Incorporating the acquisition functions into informative path planning. This includes myopic path planning, myopic planning with a look-ahead, and non-myopic planning using RRT*.
- Testing the intelligent agents through different simulation studies and using different metrics to select the agents with the best characteristics.
- Implemented and tested algorithms on AUVs in several successful field experiments.
- Comparing different types of sensors from different vehicle types for the biomass of phytoplankton and zooplankton.

7.2 Further work

There are some direct improvements that could strengthen the field experiments demonstrated here. Due to different circumstances, several of the field experiments are on the shorter side and have few repeats. To demonstrate the improvements in autonomy, field experiments should last for many hours without human intervention, which would increase the confidence for truly autonomous agents. Increasing the number of repeats would also help validate the algorithms.

Another concrete improvement is to better account for the positional error accumulated by agents while submerged. When vehicles are submerged and lose GPS connection, they must rely on estimates for their current position. Several practical approaches can be used to better approximate the position. Acoustic sensors can be used for self-localization (Miller et al., 2018), or they can be used to range platforms that have a known location (Kussat et al., 2005). After the AUV surfaces and acquires a GPS fix, the estimated positions can be re-estimated using Kalman filters. This can be further improved by utilizing the ocean current field. These methods will only reduce the uncertainty in positioning, so a further improvement would be to properly utilize uncertain locations in the models. A computationally intensive approach is to use Monte Carlo samples for the observed path to estimate the conditional field.

Statistical models have been well studied and applied to spatial fields, some strengths and weaknesses have been discussed here. There are several data driven machine learning models that can be applied (Sonnewald et al., 2021; Lou et al., 2023); these include support vector machines (SVM) (Lins et al., 2013),

convolutional neural networks (CNNs) , LSTMs (Long-short-term memory) (Liu et al., 2018), and spatio-temporal decision trees (Wu et al., 2024). Another possibility is to use hybrid approaches that combine physics informed models and data driven models (Willard et al., 2020). In general, a more thorough comparison of different possible models could be valuable. Several promising models might not work in cases where the data is too limited.

There is still the issue of validating the algorithms with field experiments. This comes down to two issues; there is no ground truth available in the field, and running comparisons between algorithms requires several agents to be run at the same time over several replicates, this is usually not feasible. Efforts can be made to establish a ground truth; for example, UAVs can be driven in a lawnmower pattern, and because they use an optical sensor and move faster than ocean going platforms, they can collect data for a larger field more quickly. This will not be a true comparison because the optical sensor will not easily observe the field in the same way as the in-situ sensors. There are several other possibilities for establishing a ground truth, but it will probably come down to doubling up the sampling efforts. This might be the wrong direction because it is an inefficient use of resources. Another approach is to not use more vehicles but rather to employ switch behavior between sampling strategies. Some work is required to determine if this is a feasible way to compare algorithms, but it is an avenue of research.

There are complex interactions between the physical and biochemical properties of the ocean, and vehicles are often equipped with several different sensors. With a limited amount of resources, the sampling design could benefit from a multi-variate approach. This comes in the form of multivariate modeling and/or multi-objective surveys. Multivariate modeling can still be done with GRFs (Ver Hoef and Cressie, 1993), but a range of other models can be used, such as GMRFs (Kreuzer and Solowjow, 2018) or an array of machine learning techniques (Nikparvar and Thill, 2021). There are many opportunities to improve onboard modeling, considering that computers are only getting faster. There have been efforts made using variants of multi-objective acquisition functions; for example, Fossum et al. (2021b) focused on joint excursion sets for salinity and temperature. Another approach exemplified by Ge et al. (2025) (C) combines two objective functions for a univariate model through a linear combination with different weights, this method is often called weighted-sum scalarization. Similar approaches can be tested for different problems, but there are other methods that would be interesting. One promising method is to use Pareto fronts between the different objective functions; the volume inside the Pareto front is often called the hypervolume, and one can create sampling designs that increase this hypervolume (Kristensen et al., 2016).

One need not be limited to running adaptive algorithms for just one vehicle at a time. One can increase sampling efficiency by either increasing the number of similar vehicles in a *homogeneous multi-robot system*, which is often also called swarm robotics (Brambilla et al., 2013; Schranz et al., 2020), or by using different vehicles with different sensors in a *heterogeneous multi-robot system* (Rizk et al., 2019). In a homogeneous system, all the robots are more or less the

same; therefore, the issue often revolves around information sharing, which is especially difficult for AUVs, as well as task sharing and division, or running in joint formations. For a heterogeneous fleet, many of the same problems apply, but there are some additional aspects to consider. The sensors are different; they can measure different things or the same things but with different methods. The different robots might also have varying operating speeds, ranges, and possible modes of communication. Thus, each such case is unique, with many different options.

Bibliography

- Haakon Bakka, Håvard Rue, Geir-Arne Fuglstad, Andrea Riebler, David Bolin, Janine Illian, Elias Krainski, Daniel Simpson, and Finn Lindgren. Spatial modeling with r-inla: A review. *Wiley Interdisciplinary Reviews: Computational Statistics*, 10(6):e1443, 2018.
- Sünnje L Basedow, David McKee, Ina Lefering, Astthor Gislason, Malin Daase, Emilia Trudnowska, Einar Skarstad Egeland, Marvin Choquet, and Stig Falk-Petersen. Remote sensing of zooplankton swarms. *Scientific reports*, 9(1):686, 2019.
- Igor M Belkin, Peter C Cornillon, and Kenneth Sherman. Fronts in large marine ecosystems. *Progress in Oceanography*, 81(1-4):223–236, 2009.
- Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.
- Jens Einar Bremnes, Ingrid Bouwer Utne, Thomas Røbekk Krogstad, and Asgeir Johan Sørensen. Holistic risk modeling and path planning for marine robotics. *IEEE Journal of Oceanic Engineering*, 50(1):252–275, 2024.
- Cecilie Broms and Webjørn Melle. Seasonal development of calanus finmarchicus in relation to phytoplankton bloom dynamics in the norwegian sea. *Deep Sea Research Part II: Topical Studies in Oceanography*, 54(23-26):2760–2775, 2007.
- Charles G Broyden. The convergence of a class of double-rank minimization algorithms: 2. the new algorithm. *IMA journal of applied mathematics*, 6(3): 222–231, 1970.
- Clément Chevalier, David Ginsbourger, and Xavier Emery. Corrected kriging update formulae for batch-sequential data assimilation. In *Mathematics of Planet Earth: Proceedings of the 15th Annual Conference of the International Association for Mathematical Geosciences*, pages 119–122. Springer, 2013.
- Clément Chevalier, Julien Bect, David Ginsbourger, Emmanuel Vazquez, Victor Picheny, and Yann Richet. Fast parallel kriging-based stepwise uncertainty

- reduction with application to the identification of an excursion set. *Technometrics*, 56(4):455–465, 2014.
- Noel Cressie and Christopher K Wikle. *Statistics for spatio-temporal data*. John Wiley & Sons, 2015.
- Trine Dale, Francisco Rey, and Berit Riddervold Heimdal. Seasonal development of phytoplankton at a high latitude oceanic site. *Sarsia*, 84(5-6):419–435, 1999.
- Emlyn John Davies. SilCam - SINTEF — sintef.no. <https://www.sintef.no/sintef-ocean/silcam/>, 2026. [Accessed 19-05-2026].
- Robert S Dietz. The sea’s deep scattering layers. *Scientific American*, 207(2):44–51, 1962.
- Raffaele Ferrari. A frontal challenge for climate models. *Science*, 332(6027):316–317, 2011. doi: 10.1126/science.1203632. URL <https://www.science.org/doi/abs/10.1126/science.1203632>.
- Trygve Olav Fossum, Petter Norgren, Ilker Fer, Frank Nilsen, Zoe Charlotte Koenig, and Martin Ludvigsen. Adaptive sampling of surface fronts in the arctic using an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 46(4):1155–1164, 2021a.
- Trygve Olav Fossum, Cédric Travelletti, Jo Eidsvik, David Ginsbourger, and Kanna Rajan. Learning excursion sets of vector-valued gaussian random fields for autonomous ocean sampling. *The annals of applied statistics*, 15(2):597–618, 2021b.
- Pilar García-Soidán, Raquel Menezes, and Óscar Rubiños. Bootstrap approaches for spatial data. *Stochastic environmental research and risk assessment*, 28(5):1207–1219, 2014.
- Roman Garnett. *Bayesian optimization*. Cambridge University Press, 2023.
- Yaolin Ge, André Julius Hovd Olaisen, Jo Eidsvik, R Praveen Jain, and Tor Arne Johansen. Long-horizon informative path planning with obstacles and time constraints. *IFAC-PapersOnLine*, 55(31):124–129, 2022.
- Yaolin Ge, Jo Eidsvik, and Tore Mo-Bjørkelund. 3-d adaptive auv sampling for classification of water masses. *IEEE Journal of Oceanic Engineering*, 48(3):626–639, 2023. doi: 10.1109/JOE.2023.3252641.
- Yaolin Ge, Jo Eidsvik, and André Julius Hovd Olaisen. Rrt*-enhanced long-horizon path planning for auv adaptive sampling using a cost valley. *Knowledge-Based Systems*, 315:113261, 2025. ISSN 0950-7051. doi: <https://doi.org/10.1016/j.knsys.2025.113261>. URL <https://www.sciencedirect.com/science/article/pii/S0950705125003089>.

- Aiqin Han, Minhan Dai, Shuh-Ji Kao, Jianping Gan, Qing Li, Lifang Wang, Weidong Zhai, and Lei Wang. Nutrient dynamics and biological consumption in a large continental shelf system under the influence of both a river plume and coastal upwelling. *Limnology and Oceanography*, 57(2):486–502, 2012.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *Model Assessment and Selection*, pages 219–259. Springer New York, New York, NY, 2009. ISBN 978-0-387-84858-7. doi: 10.1007/978-0-387-84858-7.7. URL https://doi.org/10.1007/978-0-387-84858-7_7.
- Aatira Hilal, Shabir Ahmed Bangroo, Nayar Afaq Kirmani, Javaid Ahmed Wani, Asim Biswas, Mohammad Iqbal Bhat, Khushboo Farooq, Owais Bashir, and Tajamul Islam Shah. Geostatistical modeling—a tool for predictive soil mapping. In *Remote Sensing in Precision Agriculture*, pages 389–418. Elsevier, 2024.
- Alexander R Horner-Devine, Robert D Hetland, and Daniel G MacDonald. Mixing and transport in coastal river plumes. *Annual Review of Fluid Mechanics*, 47(1):569–594, 2015.
- Y Huot, M Babin, F Bruyant, C Grob, MS Twardowski, and Hervé Claustre. Does chlorophyll a provide the best index of phytoplankton biomass for primary productivity studies? *Biogeosciences discussions*, 4(2):707–745, 2007.
- IPCC. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2013. ISBN ISBN 978-1-107-66182-0. doi: 10.1017/CBO9781107415324. URL <http://www.climatechange2013.org>.
- Jaroslav Janoš, Vojtěch Vonásek, and Robert Pěnička. Multi-goal path planning using multiple random trees. *IEEE Robotics and Automation Letters*, 6(2):4201–4208, 2021.
- Yan Kang, Delu Pan, Yan Bai, Xianqiang He, Xiaoyan Chen, Chen-Tung Arthur Chen, and Difeng Wang. Areas of the global major river plumes. *Acta Oceanologica Sinica*, 32(1):79–88, 2013.
- Bengt Karlson, Anna Godhe, Caroline Cusack, and Eileen Bresnan. Introduction to methods for quantitative phytoplankton analysis. *B, Karlson, C. Cusack and E. Bresnan (Eds.). Microscopic and molecular methods for quantitative phytoplankton analysis. UNESCO*, 2010.
- Victor Klemas and Xiao-Hai Yan. Subsurface and deeper ocean remote sensing from satellites: An overview and new results. *Progress in oceanography*, 122: 1–9, 2014.
- Edwin Kreuzer and Eugen Solowjow. Learning environmental fields with micro underwater vehicles: a path integral—gaussian markov random field approach. *Autonomous robots*, 42(4):761–780, 2018.

- Jesper Kristensen, You Ling, Isaac Asher, and Liping Wang. Expected-improvement-based methods for adaptive sampling in multi-objective optimization problems. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 50114, page V02BT03A024. American Society of Mechanical Engineers, 2016.
- Sunita Kumawat, Chanchal Dudeja, and Pawan Kumar. An extensive review of shortest path problem solving algorithms. In *2021 5th International Conference on Intelligent Computing and Control Systems (ICICCS)*, pages 176–184. IEEE, 2021.
- Neil Harvey Kussat, C David Chadwell, and Richard Zimmerman. Absolute positioning of an autonomous underwater vehicle using gps and acoustic measurements. *IEEE Journal of Oceanic Engineering*, 30(1):153–164, 2005.
- Carol Lalli and Timothy R Parsons. *Biological oceanography: an introduction*. Elsevier, 1997.
- Winfried Lampert. Ultimate causes of diel vertical migration of zooplankton: new evidence for the predator-avoidance hypothesis. In *Diel vertical migration of zooplankton*, pages 79–88, 1993.
- Andone C Lavery, Dezhang Chu, and James N Moum. Measurements of acoustic scattering from zooplankton and oceanic microstructure using a broadband echosounder. *ICES journal of marine science*, 67(2):379–394, 2010.
- Tukimat Lihan, Sei-Ichi Saitoh, Takahiro Iida, Toru Hirawake, and Kohji Iida. Satellite-measured temporal and spatial variability of the tokachi river plume. *Estuarine, Coastal and Shelf Science*, 78(2):237–249, 2008. ISSN 0272-7714. doi: <https://doi.org/10.1016/j.ecss.2007.12.001>. URL <https://www.sciencedirect.com/science/article/pii/S0272771407005392>.
- Isis Didier Lins, Moacyr Araujo, Márcio das Chagas Moura, Marcus André Silva, and Enrique López Droguett. Prediction of sea surface temperature in the tropical atlantic by support vector machines. *Computational Statistics & Data Analysis*, 61:187–198, 2013.
- Jun Liu, Tong Zhang, Guangjie Han, and Yu Gou. Td-lstm: Temporal dependence-based lstm networks for marine temperature prediction. *Sensors*, 18(11):3797, 2018.
- Ranran Lou, Zhihan Lv, Shuping Dang, Tianyun Su, and Xinfang Li. Application of machine learning in ocean data. *Multimedia Systems*, 29(3):1815–1824, 2023.
- Esam Mahdi, Ali H Abuzaid, and Abdu MA Atta. Empirical variogram for achieving the best valid variogram. *Communications for Statistical Applications and Methods*, 27(5):547–568, 2020.

- Abigail McQuatters-Gollop, Philip C Reid, Martin Edwards, Peter H Burkill, Claudia Castellani, Sonia Batten, Winfried Gieskes, Doug Beare, Robert R Bidigare, Erica Head, et al. Is there a decline in marine phytoplankton? *Nature*, 472(7342):E6–E7, 2011.
- Renato Mendes, Magda Catarina Sousa, Maite deCastro, Moncho Gómez-Gesteira, and João Miguel Dias. 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, 2016. ISSN 0079-6611. doi: <https://doi.org/10.1016/j.pocean.2015.11.006>. URL <https://www.sciencedirect.com/science/article/pii/S0079661115002414>.
- Alexander Miller, Boris Miller, and Gregory Miller. Auv navigation with seabed acoustic sensing. In *2018 Australian & New Zealand Control Conference (ANZCC)*, pages 166–171. IEEE, 2018.
- Werner G Müller and Dale L Zimmerman. Optimal designs for variogram estimation. *Environmetrics: The official journal of the International Environmetrics Society*, 10(1):23–37, 1999.
- Ahmad Mustapha. Influence of river plume on variability of chlorophyll a concentration using satellite images. *Journal of Applied Sciences*, 2011.
- Behnam Nikparvar and Jean-Claude Thill. Machine learning of spatial data. *ISPRS International Journal of Geo-Information*, 10(9):600, 2021.
- NOAA. Noaa. how much oxygen comes from the ocean? <https://oceanservice.noaa.gov/facts/ocean-oxygen.html>, 2026a. Accessed: 2026.06.02.
- NOAA. What are plankton? — oceanservice.noaa.gov. <https://oceanservice.noaa.gov/facts/plankton.html>, 2026b. [Accessed 16-05-2026].
- Iram Noreen, Amna Khan, Zulfiqar Habib, et al. Optimal path planning using rrt* based approaches: a survey and future directions. *International Journal of Advanced Computer Science and Applications*, 7(11):97–107, 2016.
- W John O’Brien. The predator-prey interaction of planktivorous fish and zooplankton: recent research with planktivorous fish and their zooplankton prey shows the evolutionary thrust and parry of the predator-prey relationship. *American scientist*, 67(5):572–581, 1979.
- Anouck Ody, David Doxaran, Romaric Verney, François Bourrin, Guillaume P. Morin, Ivane Pairaud, and Aurélien Gangloff. Ocean color remote sensing of suspended sediments along a continuum from rivers to river plumes: Concentration, transport, fluxes and dynamics. *Remote Sensing*, 14(9), 2022. ISSN 2072-4292. doi: 10.3390/rs14092026. URL <https://www.mdpi.com/2072-4292/14/9/2026>.

- André Julius Hovd Olaisen, Yaolin Ge, and Jo Eidsvik. Using expected improvement of gradients for robotic exploration of ocean salinity fronts. *Environmetrics*, 36(6):e70037, 2025.
- Donald B Olson, Gary L Hitchcock, Arthur J Mariano, Carin J Ashjian, Ge Peng, Redwood W Nero, and Guillermo P Podestá. Life on the edge: marine life and fronts. *Oceanography*, 7(2):52–60, 1994.
- Kaare Brandt Petersen, Michael Syskind Pedersen, et al. The matrix cookbook. *Technical University of Denmark*, 7(15):510, 2008.
- José Pinto, Margarida Faria, Joao Fortuna, Ricardo Martins, Joao Sousa, Nuno Queiroz, Frederic Py, and Kanna Rajan. Chasing fish: Tracking and control in a autonomous multi-vehicle real-world experiment. In *2013 OCEANS-San Diego*, pages 1–6. IEEE, 2013.
- Luca Polimene, Sevrine Sailley, Darren Clark, Aditee Mitra, and J Icarus Allen. Biological or microbial carbon pump? the role of phytoplankton stoichiometry in ocean carbon sequestration. *Journal of Plankton Research*, 39(2):180–186, 2017.
- Yara Rizk, Mariette Awad, and Edward W Tunstel. Cooperative heterogeneous multi-robot systems: A survey. *ACM Computing Surveys (CSUR)*, 52(2): 1–31, 2019.
- Melanie Schranz, Martina Umlauf, Micha Sende, and Wilfried Elmenreich. Swarm robotic behaviors and current applications. *Frontiers in Robotics and AI*, 7:36, 2020.
- Mel Siegel. The sense-think-act paradigm revisited. In *1st International Workshop on Robotic Sensing, 2003. ROSE'03.*, pages 5–pp. IEEE, 2003.
- Dag Slagstad and Thomas A McClimans. 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, 2005.
- NTNU SmallSat Lab. Observational Pyramid - NTNU — ntnu.edu. <https://www.ntnu.edu/smallsat/observational-pyramid>, 2026. [Accessed 19-05-2026].
- Theodore J Smayda. Harmful algal blooms: their ecophysiology and general relevance to phytoplankton blooms in the sea. *Limnology and oceanography*, 42(5part2):1137–1153, 1997.
- Maike Sonnewald, Redouane Lguensat, Daniel C Jones, Peter D Dueben, Julien Brajard, and Venkatramani Balaji. Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, 16(7):073008, 2021.

- Michael L Stein. The screening effect in kriging. *The Annals of Statistics*, 30(1):298–323, 2002.
- Andy Stock. Choosing blocks for spatial cross-validation: lessons from a marine remote sensing case study. *Frontiers in Remote Sensing*, 6:1531097, 2025.
- Lynne D Talley. *Descriptive physical oceanography: an introduction*. Academic press, 2011.
- Andrew C. Thomas and Ryan A. Weatherbee. Satellite-measured temporal variability of the columbia river plume. *Remote Sensing of Environment*, 100(2):167–178, 2006. ISSN 0034-4257. doi: <https://doi.org/10.1016/j.rse.2005.10.018>. URL <https://www.sciencedirect.com/science/article/pii/S0034425705003482>.
- Emmanouil G Tsardoulias, A Iliakopoulou, Andreas Kargakos, and Loukas Petrou. A review of global path planning methods for occupancy grid maps regardless of obstacle density. *Journal of intelligent & robotic systems*, 84(1): 829–858, 2016.
- Anete Vagale, Robin T Bye, Rachid Oucheikh, Ottar L Osen, and Thor I Fossen. Path planning and collision avoidance for autonomous surface vehicles ii: a comparative study of algorithms. *Journal of Marine Science and Technology*, 26(4):1307–1323, 2021.
- Jay M Ver Hoef and Noel Cressie. Multivariable spatial prediction. *Mathematical Geology*, 25(2):219–240, 1993.
- Keith Vickery. Acoustic positioning systems. a practical overview of current systems. In *Proceedings of the 1998 Workshop on Autonomous Underwater Vehicles (Cat. No. 98CH36290)*, pages 5–17. IEEE, 1998.
- Yihe Wang, Hui Wu, Lei Gao, Fang Shen, and X San Liang. Spatial distribution and physical controls of the spring algal blooming off the changjiang river estuary. *Estuaries and Coasts*, 42(4):1066–1083, 2019.
- Jared Willard, Xiaowei Jia, Shaoming Xu, Michael Steinbach, and Vipin Kumar. Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919*, 1(1):1–34, 2020.
- Huisheng Wu, Lejie Wang, Xiaochun Ling, Long Cui, Ruixue Sun, and Na Jiang. Spatiotemporal reconstruction of global ocean surface pco2 based on optimized random forest. *Science of The Total Environment*, 912:169209, 2024.
- Jianhao Wu, Chang Liu, Vladimir Filaretov, Dmitry Yukhimets, Rongjie Cai, Ao Zheng, and Alexander Zuev. Review of research on cooperative path planning algorithms for auv clusters. *Drones*, 9(11):790, 2025.
- Maurice Yonge. 1.—the inception and significance of the challenger expedition. *Proceedings of the Royal Society of Edinburgh, Section B: Biological Sciences*, 72(1):1–13, 1972.

- Zheng Zeng, Lian Lian, Karl Sammut, Fangpo He, Youhong Tang, and Andrew Lammas. A survey on path planning for persistent autonomy of autonomous underwater vehicles. *Ocean Engineering*, 110:303–313, 2015.
- Jing Zhou, Yulin Si, and Ying Chen. A review of subsea auv technology. *Journal of Marine Science and Engineering*, 11(6):1119, 2023.
- Weining Zhu, Qian Yu, Yong Q. Tian, Robert F. Chen, and G. Bernard Gardner. Estimation of chromophoric dissolved organic matter in the mississippi and atchafalaya river plume regions using above-surface hyperspectral remote sensing. *Journal of Geophysical Research: Oceans*, 116(C2), 2011. doi: <https://doi.org/10.1029/2010JC006523>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2010JC006523>.

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

June 22, 2026

Abstract

We develop, test, and deploy a sampling design strategy that enables an autonomous underwater vehicle (AUV) to explore and detect large gradients in spatio-temporal random fields. Our approach models the field using a Gaussian random field, which means that the directional derivatives of the field are Gaussian distributed. Leveraging fast matrix factorization and data thinning techniques, we obtain real-time data assimilation and design evaluation onboard the AUV. At each stage in the dynamic framework, possible design transects are formed based on a spider-leg search space pattern, and the agent chooses the optimal design for the next stage. The design criterion used is based on expected improvement (EI) in directional derivatives. This means that we compute the expected value of observing a larger derivative than what has been seen already. EI is among the most popular acquisition functions in Bayesian optimization. To evaluate the effectiveness of this approach, we conduct a simulation study comparing EI with alternative selection criteria. Our algorithm was embedded on an AUV which was deployed for characterizing a river plume frontal system in a Norwegian fjord. Using EI in the salinity field derivatives, the vehicle successfully sampled the fjord for approximately two hours without human intervention in two separate field experiments.

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

1 Introduction

With advances in sensor technology and embedded computing, there is a drive to develop intelligent systems for monitoring dynamic environmental processes. Recent examples include internet-of-things for smart sensor networks measuring air pollution (Dhingra et al., 2019), embedded systems and AI for agriculture (Shadrin et al., 2019), robotic systems for understanding environmental variability (Dunbabin and Marques, 2012) and cyber-physical systems for ecological monitoring (Schranz et al., 2021).

Motivated by challenges in underwater robotics, our focus is on developing algorithms for an autonomous underwater vehicle (AUV) to dynamically explore and detect ocean fronts. Frontal systems are important in meteorology and oceanography as they tend to be key drivers of the physical systems, 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). These fronts are by nature always changing. Hence, for a robotic agent to be able to monitor them, they need to be able to make intelligent decisions based on real-time measurements. In this paper, we explore ocean fronts arising from fresh river water meeting the saline ocean water, this is called a river plume front. The spatio-temporal variable of interest is then ocean salinity and its derivatives which capture the characteristics of the river plume.

Ocean fronts have been described by numerical ocean models (Lermusiaux, 2006) and satellite data (Hopkins et al., 2010). However, these have limitations: Numerical models are often biased from initial values or boundary conditions in the partial differential equations. Satellite data are sparse in time, only on the ocean surface and not available on a cloudy day. AUVs have become an important tool for oceanographic in-situ sampling, and they hence provide new opportunities for detecting and understanding ocean fronts. AUVs can navigate underwater and carry a range of sensors such as a salinity sensor. With capabilities for onboard computing, the AUV can react to data and even assimilate data in a model. This again facilitates real-time intelligent operations, with optimal sampling design decisions being made on-the-fly. Only then can one get the optimal benefits of AUV deployments, which are otherwise hampered by the limited underwater communication opportunities.

The capabilities of such embedded systems can be improved by leveraging knowledge from spatio-temporal statistics and design of experiments Wang et al. (2012, 2020); 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. Recent studies in spatial design for environmental applications include that of Low et al. (2024) who used an integrated prediction variance design criterion for groundwater monitoring, Koski and Eidsvik (2024) used value of information analysis to select lake monitoring sites, and Prentius and Grafström (2024) who considered new spatial balancing measures in designs. The optimal design in such situations is often only available via heuristic search solutions. Also, the design criteria and suggested solutions tend to be case-specific.

Adaptive sampling design is often more optimal than conventional pre-scripted designs, see e.g. the examples in Chao and Thompson (2001) and Thilan et al. (2023). Adaptive algorithms allow the sampling design to be modified based on the data that have become available. With an AUV exploring the ocean, it is also natural to focus on such adaptive or sequential designs because the search is limited to the path chosen by the AUV, without opportunities for quickly moving the AUV to a completely different location. Again, the criterion and method of choice for adaptive designs often depend on the situation: Krause et al. (2008) studied sensor placement optimizing the mutual information. Di Biase et al. (2024) used inclusion probabilities in an animal populations characterization problem and Bonneau et al. (2014) used reinforcement learning to select sampling

sites under cost constraints in an agricultural application. Other active learning-type criteria for adaptive sampling include that of level sets characterization (Hitz et al., 2017; Bottarelli et al., 2019) or minimizing the variance of excursion sets (Fossum et al., 2021; Foss et al., 2022).

Related sequential search problems are often employed in Bayesian optimization problems, e.g. Garnett (2023). Expected improvement (EI) is then commonly used as a criterion for function evaluation here.

The goal of this paper here is to develop, test and deploy design algorithms for efficient AUV sampling of spatial derivatives. The AUV has a relatively small on-board computer, and for real-time operation one must take care in the implementation. We embed a Gaussian random field (GRF) model on-board the AUV. Leveraging effective matrix computations at the size of the data collected during the current stage of the sampling scheme, we enable efficient data assimilation and design computation. Our suggested scheme is based on maximizing the EI in directional derivatives over a set of transects that form locally at each sampling stage.

The main contributions of this paper are

- A spatio-temporal sampling approach with a dynamic agent effectively exploring large derivatives in the field.
- A fast scalable algorithm for updating and design choices based on GRFs and sparse observation points in the vicinity of the current position.
- Field deployments with an AUV adaptively sampling a river plume front in a Norwegian fjord for more approximately 2 hours each.

The novelty of our approach is mainly to bridge ideas from statistical sampling designs and robotic sampling, for efficient sampling for spatial derivatives. Such developments are important for the exploration of ocean fronts, which are currently studied intensely with the challenges of climate change.

There are other recent approaches exploring the properties of ocean fronts with AUVs. Zhang et al. (2019) used in-situ AUV measurements to follow the front of mixing cold and warm water masses in Monterey Bay, California. Fossum et al. (2021) deployed an AUV in Arctic waters to understand the frontal zone near the ice shelf. Fonseca et al. (2023) used AUV samples to map the front of chlorophyll 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 these articles use spatio-temporal statistical models or approaches from spatial design, which would likely have improved the mapping performance.

In Section 2, we describe the problem motivating this study and define the necessary notation. In Section 3, we set up the required building blocks on spatio-temporal GRFs and the properties of their derivatives, along with methods for Gaussian model updating. In Section 4, we present the method and algorithm for effective adaptive spatial sampling for characterizing large directional derivatives in the field. In Section 5, we demonstrate the properties of the suggested design algorithm in a simulation study. In Section 6, we show the results of AUV deployments in the Trondheim fjord in Norway. In Section 7, we provide conclusions and point to future work.

2 Problem description and notation

Let $\mathbf{s} = (s_e, s_n, s_d, t)$ be a point in space and time. Here, $(s_e, s_n, s_d) \in \mathcal{D} \subset \mathcal{R}^3$ represent east, north and depth coordinates in a domain \mathcal{D} while $t > 0$ is a temporal index. The spatio-temporal quantity of interest, denoted $x(\mathbf{s}) \in \mathcal{R}$, is ocean salinity in our application.

We are primarily interested in detecting large changes in the spatio-temporal quantity. In our application, this would indicate ocean frontal zones. We define the directional gradient from location \mathbf{s} to \mathbf{s}' by

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

With decreasing distance $d(\mathbf{s}, \mathbf{s}')$, we obtain the field derivative at \mathbf{s} in the direction towards \mathbf{s}' . In practice, we instead consider distance $d = d(\mathbf{s}, \mathbf{s}')$ as a tuning parameter that can be specified in the context of the application and the operational design constraints.

An observation made at space-time location \mathbf{s} is denoted $y(\mathbf{s})$, and $\mathbf{y}(\mathcal{S}) = (y(\mathbf{s}_1), y(\mathbf{s}_2), \dots, y(\mathbf{s}_N))$ are observations from a set of sampling points $\mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_N\}$. Considering sequential data gathering with stages $k = 1, \dots$, size- N_k batch of data $\mathbf{y}_k = \{y(\mathbf{s}_1^{(k+1)}), \dots, y(\mathbf{s}_{N_k}^{(k+1)})\}$ are collected at locations $\mathcal{S}_k = \{\mathbf{s}_1^{(k+1)}, \dots, \mathbf{s}_{N_k}^{(k+1)}\}$. Hence at stage k , $N_{1:k} = \sum_{l=1}^k N_l$ data points are available. We denote the set of all sampling locations up to this stage 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 observations

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

We represent the quantity of interest by a GRF model. The differences in equation (1) are then also Gaussian distributed. We further assume a Gaussian model for the observations in equation (3). With these assumptions, one can evaluate proposed sampling designs and conduct data assimilation in near real-time, without resorting to large computing resources.

Our setup, where an agent moves at a limited speed along its trajectory, presents a unique data collection challenge. Traditional space-covering designs would require extensive time for exploration due to the agent's inability to quickly reposition. Instead, we develop, test, and deploy an algorithm that sequentially evaluates design adjustments from the agent's current position. This grid-free approach allows the agent to dynamically select interesting trajectories as they emerge, avoiding the need for a pre-specified waypoint grid, which is a common constraint in robotics applications.

At each stage k , the agent computes expected rewards for staying on the same trajectory and for changing to another direction. Design directions with large expected derivatives are preferred in our case. Each of J possible designs consists of a transect spanning out from the current location, having 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 $\mathbf{p}_{k,l}^j \in \mathcal{D}$. 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.

Specifically, in our application, we propose a spider-leg design template for the agent. This is constructed at constant depth coordinate, focusing on directional gradients in the lateral domain. Extensions to depth or time variation are possible without much additional computing requirements, but salinity changes are often most difficult to characterize in the lateral domain. Depth variations are often more stable, and we assume that the agent moves significantly faster than the phenomenon develops over time. With this spider-leg template in the lateral parts of the domain, we limit the scope to having $J = 7$ possible new directions in the lateral domain.

The agent is able to maintain a nearly constant velocity and with sensors sampling at a constant frequency. Two locations \mathbf{s} and \mathbf{s}' which are used when computing the differences are hence separated by the same distance during the deployment.

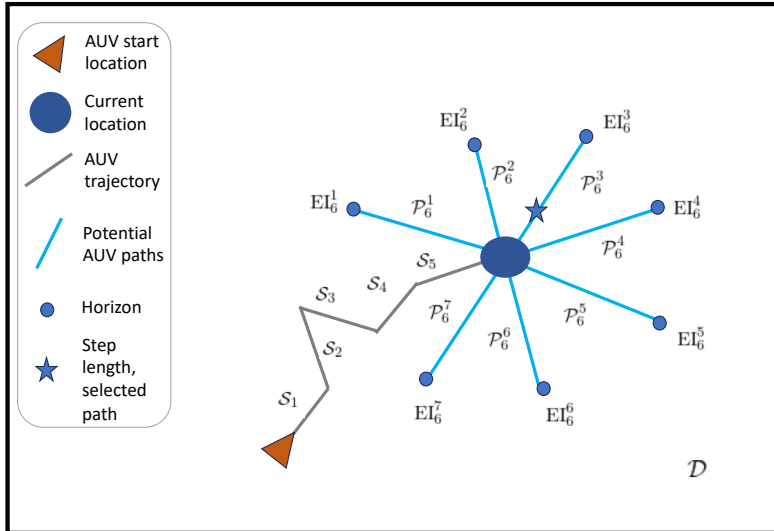


Figure 1: Illustration of an agent’s path made up of transects at 5 earlier stages. At the current location, the agent chooses one of the 7 possible designs (not possible to turn straight back). The design selection criteria is here expected improvement (EI) for the directional salinity gradients.

Figure 1 shows the situation with an agent at the current location (blue circle), having passed five stages. It must now make a decision about where to go next, for stage six. The agent is not constrained to travel the entire transect of length M_j . Instead, it conducts new design evaluations after moving one step-length down the best transect. In this illustration, design number 3 is selected as the most valuable, and the agent moves a step-length in this direction (marked by a star).

3 Spatio-temporal Gaussian random fields

Robotic units and sensor systems have limited computing, storage and communication capabilities. Ideas from surrogate modeling (Gramacy, 2020; Fuhg et al., 2021) are hence highly relevant in this context because they balance realistic representations of the spatio-temporal process and ease of computations. Here, a GRF is used as a surrogate model onboard the agent.

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. The covariance function is specified from multiple ocean models as well as previously acquired data from the domain of interest. We assume a separable space-time covariance model. A similar separable covariance kernel has been used to model ocean salinity and temperature by Sahu and Challenor (2008). For a

Gaussian covariance function, we have

$$C(\mathbf{s}, \mathbf{s}') = \sigma^2 \exp\left(-\left(\frac{|s_n - s'_n|^2 + |s_e - s'_e|^2}{\phi_s^2}\right)\right) \exp\left(-\left(\frac{|t - t'|}{\phi_t}\right)^2\right), \quad (4)$$

where we assume constant sampling depth s_d and with ϕ_s and ϕ_t being the correlation decay parameters of space and time, respectively. We will go into how these parameters are estimated for the field experiments in section 6.1

For any set of N space-time locations $\mathcal{S} = \{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_N\}$, the length- N random vector $\mathbf{x}_{\mathcal{S}}$ 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

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

3.2 Directional differences and GRFs

The variable difference $g(\mathbf{s}, \mathbf{s})$ in Equation (1) is a linear combination of two Gaussian distributed variables, and hence it is also Gaussian distributed. 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

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

One can consider such differences between any pairs of variables at equal distances along a transect line in the spatial domain (Figure 1). The resulting random vector of directional differences is multivariate Gaussian distributed.

When studying properties of such derivatives we notice one of the main benefits of using a Gaussian covariance function. Figure 2 shows five different 1D GRF realizations, one using an exponential covariance function ($\nu = 1/2$), Matérn covariance function with $\nu = (3/2, 5/2, 7/2)$ and Gaussian ($\nu \rightarrow \infty$) covariance function. For each of the covariance functions the correlation between points with a distance of 300 m is 0.05. The realizations are also conditioned on 10 points, the points are drawn from a GRF with a Matérn covariance function using $\nu = 3/2$ with independent Gaussian measurement noise. Among the five realizations using the Gaussian covariance function is the smoothest, and the corresponding gradient is also smooth. The realization with Matérn $\nu = 7/2$ covariance is also quite smooth. The derivatives for $\nu = 1/2$ and $\nu = 3/2$ the Matérn $\nu = 3/2$ results in large and not very smooth derivatives. This means that it is likely more robust to impose a smooth correlation function when the goal is to search for hot-spots in gradients.

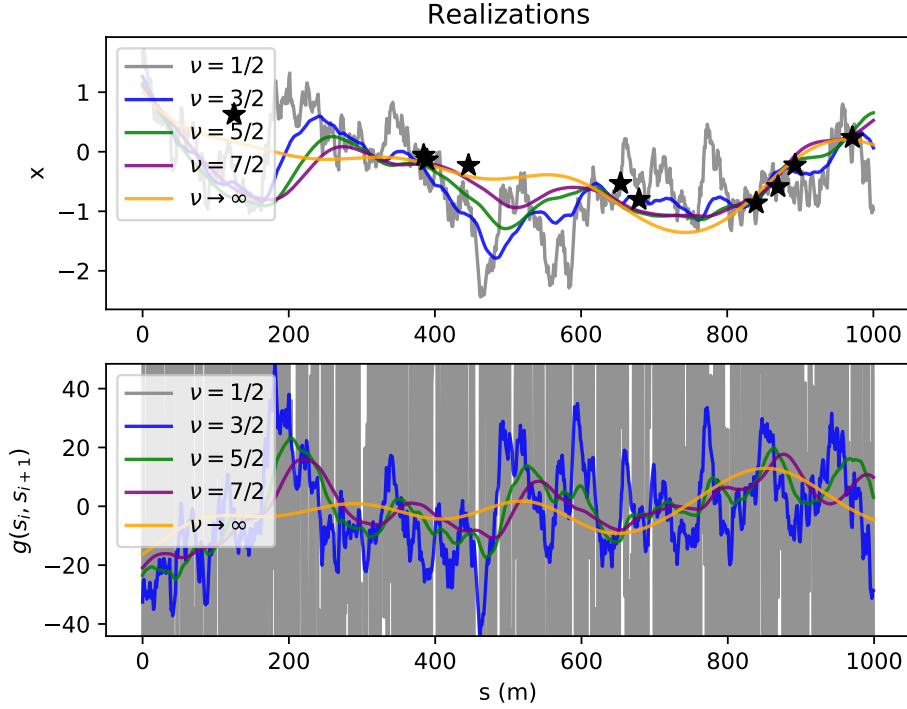


Figure 2: One realization using 5 different values for ν . Exponential ($\nu = 1/2$), Matérn with $\nu = (3/2, 5/2, 7/2)$ and Gaussian ($\nu \rightarrow \infty$) covariance function. These realizations are conditioned on the 10 star points in the top plot. The Gaussian covariance function gives the smoothest realizations and gradients. The plotting scale for the gradients is truncated, as the gradients for the exponential covariance function become extremely large.

3.3 Conditioning to in-situ observations

The agent updates the on-board model based on in-situ measurements. The observation 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 τ^2 is the error variance. Observation errors at different locations are assumed to be independent.

The additive Gaussian noise assumptions in Equation (10) entail that the conditional model remains Gaussian. For design location set $\mathcal{P} \in \mathcal{P}_{k+1}^j$ at stage $k+1$, given the currently available data $\mathbf{y}_{1:k}$, mean vector and covariance matrix are

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

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

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

In particular, at any two points $\mathbf{s} \in \mathcal{P}$ and $\mathbf{s}' \in \mathcal{P}$, the variable $(x(\mathbf{s}), x(\mathbf{s}'))$ is bivariate Gaussian distributed. Their scaled difference is then also Gaussian, see Equation (8)-(9).

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 $N_{1:k} \times N_{1:k}$. This is of order $O(N_{1:k}^3)$ calculations, and it 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). 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 available from the previous stage, and

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

The required inversion in this expression is for matrix \mathbf{C} , which is of moderate size $N_{k+1} \times N_{k+1}$ because it only involves the new batch. Equation (14) 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.

Even though this matrix factorization trick allows efficient computing onboard the agent, the mere size of the covariance matrix grows over stages and this leads to evaluation times that hinder real-time operation. To address this challenge further, we implement an approach for reducing the data size over space-time by thinning data from far away/long ago. This notion of nearest observations having the largest impact on the Kriging prediction is sometimes referred to as spatial screening (Stein, 2002). In our setting, we set a threshold computation time for the data assimilation and prediction steps. If the agent's on-board computation time gets larger than this threshold time, we thin the data points in memory by half. Thinned data locations are far away from the agent's current position, and the variables at thinned locations have a low correlation with the prediction points. After this thinning step, the matrix inverse \mathbf{A} in Equation (13) must be re-computed before re-commencing with the recursive factorization formula for the next stages.

Note that our procedure treats the model parameters of the GRF model as fixed. Hence, they are not re-estimated from in-situ data. One could in principle update model parameters as well, say, using maximum likelihood estimation, but this would not scale sufficiently well for real-time operation because the agent's computer is relatively small. If such parameter updating is desired, we suggest pausing the operation. The agent then takes extra time to surface and can communicate data to a larger computer using 4G. Once this faster computer has re-estimated model parameters, they are communicated back to the agent, which can continue the operation with a revised model description and new parameter values.

4 Adaptive sampling design

For the adaptive selection of designs as illustrated by the spider leg design in Figure 1, the agent first evaluates the expected rewards of the 7 possible transects. Next, an optimal decision (direction) is made based on the highest expected reward in terms of directional differences. This procedure is repeated over stages as the agent traverses the domain of interest. The presentation in Section 4.1 describes the suggested sampling approach using a single stage (with suppressed notation for stage k), while Section 4.2 outlines the algorithm for multiple stages in detail.

4.1 Expected improvement for spider legs transects

The design transects \mathcal{P}^j , $j = 1, \dots, 7$, in Figure 1 are evaluated based on the updated conditional distributions in Equation (11)-(12), which also defines the Gaussian distribution for directional differences, see Equation (8)-(9). The transect design choice is made according to the optimal expected gain in directional differences. For this purpose, we let g_{\max} denote the largest absolute directional difference observed thus far in the sampling. We first present the probability of improvement (PoI) in this directional difference. Then the expected improvement (EI) is described. (In Section 5 we compare PoI and EI in a simulation study.)

Probability of improvement

The PoI for a design is the probability of having a difference larger than g_{\max} for the selected observation points. For a transect $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_N\}$, variables at two neighboring locations $\mathbf{p}_i, \mathbf{p}_{i+1}$, $i = 1, \dots, N - 1$ define the difference. To simplify notation, the conditional distribution of the difference $g_i = g(\mathbf{p}_i, \mathbf{p}_{i+1})$ is denoted $\mathcal{N}(\zeta_i, \eta_i^2)$. Then the probability that $|g(\mathbf{p}_i, \mathbf{p}_{i+1})|$ is larger than g_{\max} is

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

where $\Phi(z)$ is the standard normal cumulative distribution function evaluated at z .

The PoI in Equation (14) is computed for all two-neighbor locations along a transect. Next, the optimal transect according to the largest PoI is defined by

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

There exist alternative forms of PoI. For instance, Jones (2001) suggested using a shifted mean to encourage more exploration in the designs. It is however not easy to tune the fudge shift parameter.

Expected improvement

EI represents the expected gain in difference over g_{\max} . Here, the gain or improvement is

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

For Gaussian distributions, there is a closed-form solution for the EI, see e.g. Jones (2001). We have

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

where $\phi(z)$ is the standard normal probability density function evaluated at z .

EI has been used much in for instance Bayesian optimization (Garnett, 2023). Compared to PI, it tends to more naturally balance exploitation and exploration.

The EI is computed for all two-neighbor locations along a transect design. Next, the optimal design is selected based on the largest EI:

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

4.2 Algorithm

At each stage k one conducts i) data assimilation (Section 3), ii) prediction of directional derivatives along all possible design transects, iii) evaluation of the design criteria to select the optimal transect (Section 4.1), iii) sampling of data (assumed to be represented by the model in Equation (3)). Then returns to step i) for the next stage. The algorithm using EI as a design criterion is summarized in Algorithm 1.

Along with the statistical model, data assimilation and the statistical design criteria, some further details are needed to fully describe this algorithm. For practical implementation, one must define the length or horizon for the $J = 7$ possible transects. This defines N_{k+1} and the locations in designs P_{k+1}^j , $j = 1, \dots, J$. Further, for proactive deployment, the agent does not necessarily move until the end of this transect. Instead, it sets a waypoint (WP) at a shorter step length. (This is illustrated by the star in Figure 1.) The reason for predicting far into the future and moving only a short path is that one can react quickly to the new measurements while looking for promising far-away large derivatives in the design criterion.

Algorithm 1 shows the steps of the sequential procedure. At each stage k , the agent computes the optimal design, and based on this it sets the next target waypoint (WP_k). In the real-world deployments, the initial point (WP_0) is a location that is expected to be near the front. Once, the next waypoint is known, the agent collects data according to this selected design direction heading towards this waypoint, updates the model, and then evaluates all designs for the next stage when it reaches the waypoint.

Figure 3 illustrates parts of this procedure as it folds out in the simulation study. In the displays, the agent has sampled for 20 steps. It is now at the white triangle, about to make a decision on where to move next. The 7 candidate transects are shown in dark green. The next waypoint (white star) is set a step length along the transect having the largest EI in directional differences.

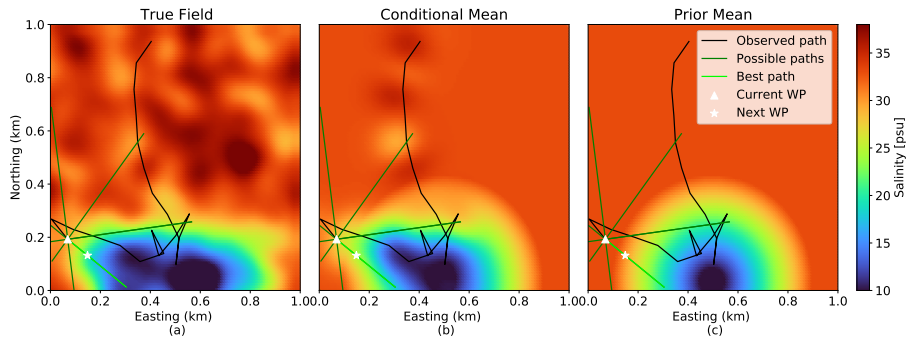


Figure 3: Illustration of spider leg designs for (a) one realization of a true salinity field, (b) the conditional mean salinity field, and (c) the prior mean field. All three plots show the observed path of the AUV in black, and the possible design transects in dark green. The optimal design is in light green and the next waypoint is one step-length down this optimal design transect.

Recall that the computation time does not depend on the size of the field domain in our case. It depends on i) the number of data points, ii) the frequency of design evaluations, and iii) the number of designs evaluated. For i), we implement efficient matrix factorization and thinning of data (Section 3.4). For ii) and iii), we have algorithmic tuning parameters defining the length of the 7 planning transects and the step length along a transect forming the next waypoint. We describe these further in the context of the simulation study and then use the same for the real-world deployments.

Algorithm 1 Adaptive sampling for large derivatives.

Require:

$\mu(\mathbf{s}), C(\mathbf{s}, \mathbf{s}')$, operational domain \mathcal{D} , distance d , step-length along transect, transect horizon.

$\mathcal{S}_0 = \emptyset, \mathbf{y}_0 = \emptyset, g_{\max,0} = 0$, waypoint WP_0 near the expected front.

repeat For each time $k = 0, 1, \dots$ ▷ Loop over stages

Define spider leg designs \mathcal{P}_{k+1}^j for $j = 1, \dots, J$ transects.

Set $\text{EI}_{\max} = 0$

repeat For each spider-leg $\mathcal{P}_{k+1}^j, j = 1, \dots, J$ ▷ Design evaluation

Predict $\mathbf{m}_{\mathcal{P}_{k+1}^j}$ and $\Psi_{\mathcal{P}_{k+1}^j}$ conditional on data $\mathbf{y}_{1:k}$ at $\mathcal{S}_{1:k}$. ▷ Eq. (11)-(12).

Compute $\text{EI}_{k+1}^j = \max_{\mathbf{p}_i \in \mathcal{P}_{k+1}^j} E[\max(|g_i| - g_{\max,k}, 0)]$ ▷ Eq. (17).

if $\text{EI}_{k+1}^j > \text{EI}_{\max}$ **then**

$\text{EI}_{\max} = \text{EI}_{k+1}^j$

Set $\mathcal{P}_{k+1}^* = \mathcal{P}_{k+1}^j$

end if

until

Along transect \mathcal{P}_{k+1}^* , set size- N_{k+1} set of locations \mathcal{S}_{k+1} , with end location WP_{k+1}

The agent moves from $\text{WP}_k = \mathbf{s}_0^{(k+1)}$ to $\text{WP}_{k+1} = \mathbf{s}_{N_{k+1}}^{(k+1)}$,

Data \mathbf{y}_{k+1} is gathered at locations \mathcal{S}_{k+1}

Update $\mathbf{y}_{1:k+1} = (\mathbf{y}_{1:k}, \mathbf{y}_{k+1})$ and $\mathcal{S}_{1:k+1} = \{\mathcal{S}_{1:k}, \mathcal{S}_{k+1}\}$ ▷ Augment dataset

Find the largest absolute derivative $g_{k+1} = \max_{i=1, \dots, N_{k+1}} \left\{ \left| \frac{y(\mathbf{s}_i^{(k+1)}) - y(\mathbf{s}_{i-1}^{(k+1)})}{d} \right| \right\}$

Update maximum derivative $g_{\max,k} = \max(g_{\max,k}, g_{k+1})$

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

Points in memory are thinned.

end if

until

5 Simulation study

We test and compare different sampling design strategies on a replicate study. The results of this study provide recommendations for the real-world setup in the next section.

5.1 Simulation setup

We consider an AUV that can operate in a square 2D domain of size $(0, 1) \times (0, 1)$ km² as shown in Figure 3. In this domain of interest, we generate 100 replicate true fields from a 2D GRF model. No time variation is used here. The mean (see Figure 3 right) is specified to mimic a river outlet in the southern part. There are hence low salinity values in the south and higher salinity in the north. We use a Gaussian covariance function:

$$C(\mathbf{s}, \mathbf{s}') = \sigma^2 \exp \left(- \left(\frac{|s_n - s'_n|^2 + |s_e - s'_e|^2}{\phi_s^2} \right) \right), \quad (19)$$

with $\phi_s = 100$ m and $\sigma = 4$. This initial GRF model is updated with data during the simulation run.

The starting location for the AUV is sampled uniformly, at least 700m from the modeled outlet. For all experiments, the AUV will run for a total of 5000 m. In the design algorithm, we adjust

transect alternatives crossing the boundaries of the domain. The speed of the AUV is set to be 1 m/s with a sampling frequency of 1 Hz. This means that the distance between sampled points is 1 m. The measured salinity noise variance is $\tau^2 = 2^2$. In the algorithm, we use a transect horizon of 500m in the design evaluation and step-length 100m for waypoint distances. This means that the model assimilates $N_k = 100$ data points at each stage.

5.2 Evaluation metrics

We study statistical properties of the suggested design methods using various metrics such as i) the largest directional difference observed during the mission, and ii) how many important high-gradient regions the AUV was able to visit during its operation.

We monitor the largest measured absolute directional gradient $g_{\max,k}$ over mission time and replicate runs. Ideally, we want this value to grow quickly over stages k . This metric is also a good indicator of whether the AUV has detected a large gradient during the mission. However, the metric does not indicate the AUV's ability to explore large gradients in various regions of the field.

For domain exploration, the AUV should visit several regions where there are large gradients in the salinity field. To gauge this kind of exploration, we split the domain into 400 equal-size regions ($50 \times 50 \text{ m}^2$ size). We then find the 20 % regions with the largest absolute gradients in the generated GRF. We refer to these as important regions. Notably, the important regions vary over the replicates because they are computed based on different true fields. However since there is no time variation, they remain the same during simulation time for each replicate. As a performance metric for exploring regions with large gradients, we count how many of these important regions the AUV visits during its mission in the field.

For illustration, Figure 4 shows gradients in the salinity field (left display) for one of the replicate runs. The important regions for this replicate realization of the field are marked with red colors in




Figure 4: One realization of absolute gradient (left plot) and the corresponding important regions (right plot). The important regions cover 20 % of the domain and they represent regions with the largest gradients.

the right display.

5.3 Simulations results

We use the two different design criteria for actively learning large gradients; PoI in Equation (14) and EI in Equation (17). In addition to these two, we compare performance with three other strategies; select the direction with the largest expected difference, select the direction with the largest variance in the differences, and select at random among the 7 possible transects. All strategies rely on the same 100 replicates, the same starting points, and they use the same statistical model for data assimilation. The only difference is in the sampling design objective function used.

The results of running replicate simulations with these five different strategies are shown in Figure 5. We first study the AUV’s ability to detect large directional derivatives. In Figure 5 left) we show the maximum derivative discovered by each of the five strategies over the distance traveled by the AUV (first axis). The thick line is the mean over the 100 replicates while the shaded region represents 2 standard deviations in this mean. Here, EI works better than the other strategies, followed by PoI and the design strategy based on the largest expected difference. It appears as if PoI gradually improves over time, compared with the strategy focused on large expected derivatives. The gap between PoI and EI closes towards the end of the simulation, but EI increases much faster in the beginning of the simulation. Compared with EI, the strategies of PoI and large expected differences are both more prone to get stuck in a local maxima, and it spends much sampling effort locally before moving away. The design strategy based on large variance in the differences is worse than the best three, but it is much better than the random strategy which is worse than the others here.

In Figure 5 right) we plot the number of important regions the AUV has visited after running a given distance. Again, the thick lines indicate the mean number of regions while the shaded region represents 2 standard deviations in this mean. The ranking of the different design criteria is the same as for the largest detected derivative (left display), but the relative performance of different strategies is somewhat different. Here, capturing more of the exploration capabilities of the various strategies, EI is the best and remains better than PI for all stages. The strategy focused on the largest expected differences is now much worse than PI. It is more likely to get stuck in promising regions, not exploring the domain so well.

To gain more insight into the properties of various strategies, we study their worst-case and best-case outcomes over the replicate runs. This relates to risk in our case because it is 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 operation. In terms of the largest derivative g_{\max} , EI has a narrower range in the 95 to 5 percentiles than PI. For the 95th and 50th percentile of important regions covered, EI and PoI have similar performance, but for the 5 percentile, the difference is large. Here, EI is able to sample in 39 (out of 80) important regions, while PoI is only able to sample in 19. The other three design strategies are much worse in the 5 percentile cases both for the largest derivative and important regions. This means that EI is more robust in its performance than PI and other criteria. In particular, its worst-case performance is much better at exploring the important regions.

Regarding algorithmic parameter tuning, we conducted other test runs to evaluate the impact of various inputs. For the transect horizon, a longer horizon tends to give a faster increase in the performance metrics. The main reason to limit the horizon is computational cost and the goal of near real-time operation. For the step length used to set a new waypoint, we learned that it should not be too long. A shorter step length enables faster AUV reaction to variabilities in the field. Using a long step length, the AUV can overshoot a region with large derivatives, and it wastes time traveling the long way back.

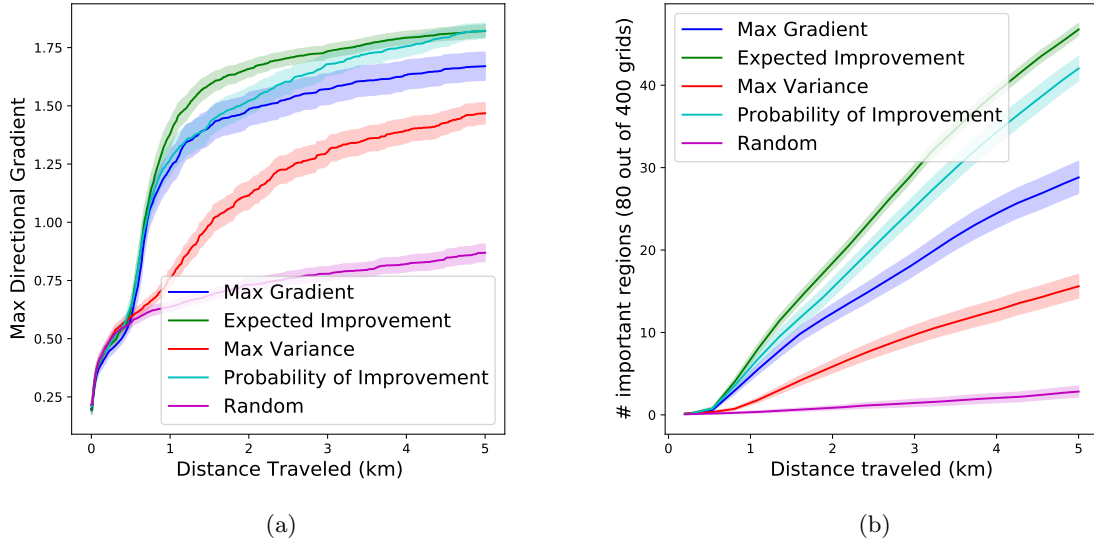


Figure 5: Comparison of different design criteria as a function of distance traveled (operation time). Results are averaged over 100 realizations (solid line) with ± 2 standard errors (shaded). (a) Increase in direction derivative. (b) Increase in important regions visited.

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	13

Table 1: Percentile table for the different objective functions with the two different evaluation metrics. The percentiles are computed from the 100 replicate runs and they are calculated at the end of the operation.

6 Case study

The proposed algorithm was tested in two field deployments in the Trondheim Fjord, Norway. The first took place on June 22, 2023 and the second on September 25, 2024.

We first outline the parameter specification in the spatio-temporal GRF using numerical ocean model data. Then we describe the AUV setup, and finally show the results of the deployment.




Figure 6: Left: SINMOD simulation results of the surface salinity of the Trondheim fjord near the Nidelva river outlet. Right: The absolute salinity gradient computed from the SINMOD simulation. The plots show river outlets and river plume fronts in dark red.

6.1 Prior model specification based on SINMOD

We have access to a numerical ocean model of the fjord. This model is called SINMOD and it has been developed by SINTEF Ocean (Slagstad and McClimans, 2005). From inputs of river flow, bathymetry and tides, SINMOD solves partial differential equations to mimic currents, temperature and salinity in the fjord. We are mainly interested in the salinity variable and how the salinity changes in space and time. Figure 6 illustrates a snapshot from a SINMOD simulation, where the left display shows the salinity and the right shows the absolute gradient. At this time in the simulation there is a transition from high to low tide¹. The darker red regions in Figure 6 right) indicate zones where the salinity variable changes rapidly. We notice large gradients near the multiple river outlets but also further out in the fjord highlighting river plume fronts. For the field deployment, the AUV will explore the north-east river outlet and its plume. This is the outlet with the highest discharge. It is also the part with less marine traffic, reducing the risk of collision. SINMOD simulations are computationally heavy to run and they cannot be updated onboard the AUV. Instead, we rely on SINMOD simulations to fit a GRF surrogate model for the salinity. We use simulations done sometime before the mission to specify the prior mean for the operation time. The spatial and temporal correlation decay parameters ϕ_t and ϕ_s , and the variance σ^2 in Equation (19) are estimated prior to deployment. For the September 2024 experiment, these parameters were estimated using maximum likelihood estimation on AUV and SINMOD data from June 2023. The parameters were specified to $\phi_t = 920$ s, $\phi_s = 120$ m and $\sigma = 0.36$. For the June 2023 experiment the parameters were estimated using the SINMOD simulations. The parameters were then set based on empirical and modeled variogram matching, giving $\phi_t = 5400$ s, $\phi_s = 530$ m and $\sigma = 2$. Maximum likelihood estimation gave very similar parameter estimates, but a somewhat shorter correlation range. Post-processing of the results shows that differences in estimates are caused by seasonal variety, river flow and tidal conditions rather than the method of parameter estimation. The salinity measurement noise standard deviation τ in Equation (10) was estimated to be $\tau = 0.27$ from previous AUV deployments in similar conditions.

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



Figure 7: LAUV-Thor maneuvering at the surface during pre-launch protocol testing.

6.2 AUV Setup

A Light Autonomous Underwater Vehicle (LAUV) from NTNU’s Applied Underwater Robotics Laboratory (AURLab) was used in the experiments. Pre-launch protocol tests consisted of standard remote control verification (Figure 7). 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 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 aims to maneuver at a depth of 1 m where we expect large variability in salinity (Berild et al., 2024). Also, the AUV is less prone to collisions with small boats when keeping this depth, so it reduces risk. Because of waves, AUV pitch and buoyancy controls, there is some variability in the true depth (0.7-1.3 m). The AUV is programmed 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. The sampling frequency is set at 0.2 Hz. It was configured to re-surface at 10-minute intervals for navigational adjustments. The starting location is close to the river mouth where the salinity changes are expected to be highest.

The suggested sampling design algorithm contains tuning parameters: the step-length was here set to 250 m, while the prediction horizon was 750 m. The AUV evaluated 7 transects at each stage, outlined in Algorithm 1. The maximum planning time is set at 5 s, when the AUV uses more time than that when making a decision the data points in memory will be thinned

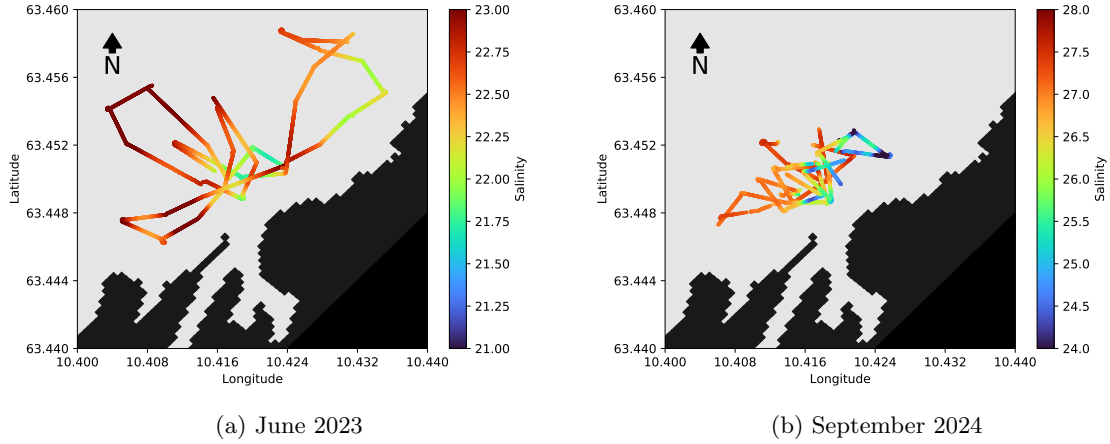


Figure 8: AUV paths for the two missions; June 2023 (left) and September 2024 (right). Note the different color scales for salinity. Rapid color changes indicate larger salinity gradients. The black parts is land.

6.3 Results

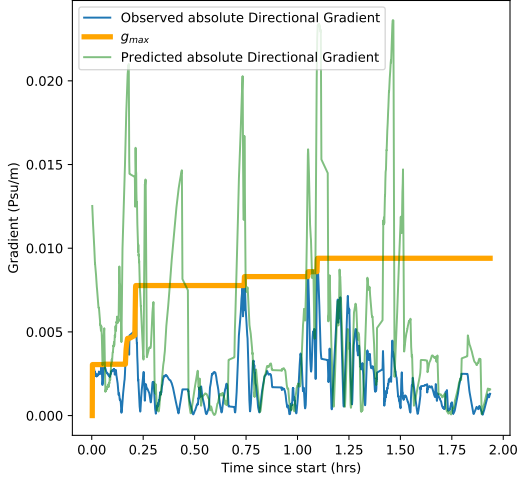
The experiment on 22 June 2023 ran for 2 h 10 min and the one on 25 September 2024 ran for 1 h 30 min. During the June 2023 the AUV took 39 stages and during September 2024 it took 27 stages. We next present the differences and similarities in the results of the two deployments.

The conditions of the river plume were quite different during the two experiments. First, the average river discharge measured on 22 June 2023 was $57.7 \text{ m}^3/\text{s}$ with an average of $160.6 \text{ m}^3/\text{s}$ the 30 days prior to the deployment. On September 25 2024, the discharge was measured at $41.7 \text{ m}^3/\text{s}$ with an average of $88.8 \text{ m}^3/\text{s}$ the 30 days prior to deployment ². This indicates that there was more freshwater flowing into the fjord during test 1, and that more fresh water had gone into the fjord prior to the deployment. Second, the tide was going from low to high tide during June 2023 experiment, while the September 2024 experiment 2 was conducted during low tide.


The paths made by the AUV during the two experiments are shown in Figure 8. The left display shows results for June 2023 and the right one is for September 2024. The color along the paths indicate the estimated salinity values, and changes in color along a transect indicate large derivatives in the salinity. Note that the color scales are different for the two plots. In June 2023 the levels ranged from 21.5-23, while in September 2024 the values ranged from 23-28. The differences in salinity are likely due to the difference in river discharge during and preceding the experiments. In the June 2023 deployment (left), the AUV explored a larger area because the salinity variations were not very high, while the September 2024 deployment discovered larger variations and the AUV stayed close to the river outlet.

Figure 9 shows the maximum salinity gradient increase during the experiment. The plot also indicates the observed absolute directional gradient and the predicted gradients. There are four distinct shifts in g_{\max} during the June 2023 experiment: in the start, 0.2 hours, 0.75 hours, and 1.1 hours into the mission. Around 1 hour to 1.5 hours into the mission the AUV sampled in a region

²River discharge data are gathered from sildre.nve.no using station 123.20.0. This station is positioned below Nedre Liarfoss hydropower station



(a) June 2023



(b) September 2024

Figure 9: Increase in maximum absolute derivative during the two missions (orange). The plots also show the measured absolute salinity change (blue) and the predicted salinity change derivative (green). Note the different scale on the second axes.

with plenty of salinity change, this is where the largest gradient is found. After this time, the AUV did not measure any very large salinity derivatives. For the September 2024 experiment, the AUV measured large derivatives during the first part and the second half of the mission, while it found and predicted only small derivatives in the middle part of the mission (0.6-0.9 hours). Note again the scale differences. There were much sharper salinity derivatives near the river outlet in September 2024, because of less flow in the river. Figure 9 (green colors) shows that when the AUV predicted a large derivative, it often observed (blue colors) a large derivative as well. This means that the GRF model and the updating approach have merits. It guided the AUV to locations of large derivatives and it was quite often right in its predictions. However, the predicted derivative is usually larger than what is observed. These inflated predictions can cause some issues in the exploration, say in the second experiment where the AUV stayed very close to the river plume, rather than exploring further out in the fjord.

Figure 10 shows the EI results over stages of the AUV sampling. The selected transects during the stages of the AUV deployments is in the red solid curve, while the dots indicate other EI values for other alternative design transects at this stage. The EI usually starts out with large values while the g_{\max} is low. Later in the operation, the EI drops closer to zero, and the AUV explores different regions where it does not necessarily predict that the derivative will be large, but rather that the variance is large. On the left plot (June 2023) we notice this behavior of EI decline. It measures the highest EI at step 4 and then it decreases over stages. This behavior is not so clear for the September 2024 deployment (plot b), where the EI remains high throughout the experiment. This can be explained by the fact that the AUV indeed measured two larger derivatives during the latter half of the experiment, and thus a high EI would be reasonable in the vicinity of these observations.

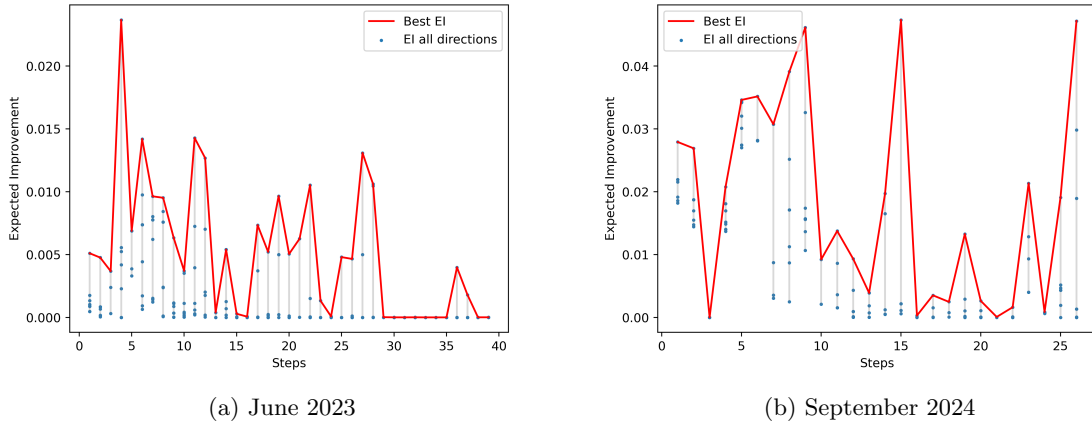


Figure 10: The EI during the stages of the two experiments. For each step, the dots show the EI for the different possible transects. The red line is then the EI for the transect with the highest EI which is chosen by the AUV.

7 Closing remarks

We have presented a new approach for spatio-temporal sampling designs, where an agent aims to locate regions exhibiting large derivatives. For our application in oceanography, such locations indicate frontal zones with much biochemical activity. The approach for adaptive sampling is based on a Gaussian random field model, and the directional differences in the field are then also Gaussian distributed. Leveraging matrix relations and data thinning, we then implement efficient data assimilation and design evaluations 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.

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. In the simulation study, we achieved the best performance using expected improvement as a design criterion. Particularly when studying the agent’s ability to explore several parts of the domain that hold large gradients. Two field experiments (June and September) demonstrated real-time operation in harsh ocean conditions. In June, there was much water flow and the autonomous underwater vehicle searched for gradients further out in the fjord compared to findings in the September mission where the water flow was low and the vehicle stayed close to the river outlet.

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 with reasonable statistical properties.

For future work, we plan to investigate more nuanced algorithms where one can tune the distance and design parameters to automatically capture the right scales on the fly. 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. The methods described in this paper can hence also be used for such applications. We used a relatively standard spatio-temporal covariance model here. It would be interesting to try more complex temporal dynamics such as advection-diffusion processes (Sigrist et al., 2015; Foss et al., 2022) or non-stationarity in space (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 Aurlab³ for making the deployment in the fjord possible, especially Karoline Barstein and Sebastian Sikora for operating supporting 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.
- Berild, M. O., Ge, Y., Eidsvik, J., Fuglstad, G.-A., and Ellingsen, I. (2024). Efficient 3d real-time adaptive auv sampling of a river plume front. *Frontiers in Marine Science*, 10:1319719.
- Bonneau, M., Gaba, S., Peyrard, N., and Sabbadin, R. (2014). Reinforcement learning-based design of sampling policies under cost constraints in markov random fields: Application to weed map reconstruction. *Computational Statistics & Data Analysis*, 72:30–44.
- Bottarelli, L., Bicego, M., Blum, J., and Farinelli, A. (2019). Orienteering-based informative path planning for environmental monitoring. *Engineering Applications of Artificial Intelligence*, 77:46–58.
- 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.
- Chao, C.-T. and Thompson, S. K. (2001). Optimal adaptive selection of sampling sites. *Environmetrics: The official journal of the International Environmetrics Society*, 12(6):517–538.
- 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.

³<https://www.ntnu.edu/aur-lab>

- Di Biase, R. M., Marcheselli, M., and Pisani, C. (2024). Achieving spatial balance in environmental surveys under constant inclusion probabilities or inclusion density functions. *Environmetrics*, page e2869.
- Dunbabin, M. and Marques, L. (2012). Robots for environmental monitoring: Significant advancements and applications. *IEEE Robotics & Automation Magazine*, 19(1):24–39.
- 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.
- Garnett, R. (2023). *Bayesian optimization*. Cambridge University Press.
- 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.
- Hitz, G., Galceran, E., Garneau, M.-È., Pomerleau, F., and Siegwart, R. (2017). Adaptive continuous-space informative path planning for online environmental monitoring. *Journal of Field Robotics*, 34(8):1427–1449.
- 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.
- Jones, D. R. (2001). A taxonomy of global optimization methods based on response surfaces. *Journal of global optimization*, 21:345–383.
- Koski, V. and Eidsvik, J. (2024). Sampling design methods for making improved lake management decisions. *Environmetrics*, page e2842.
- 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.

- Low, M. I., Bowman, A. W., Jones, W., and Bonte, M. (2024). Exact optimisation of spatiotemporal monitoring networks by p-splines with applications in groundwater assessment. *Environmetrics*, page e2874.
- LSTS (2022). Inter module communication protocol.
- 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 toolchain for networked vehicle systems. In *2013 MTS/IEEE OCEANS - Bergen*, pages 1–9.
- Prentius, W. and Grafström, A. (2024). How to find the best sampling design: A new measure of spatial balance. *Environmetrics*, page e2878.
- Quigley, M. (2009). Ros: an open-source robot operating system. In *ICRA 2009*.
- Sahu, S. K. and Challenor, P. (2008). A space-time model for joint modeling of ocean temperature and salinity levels as measured by argo floats. *Environmetrics*, 19(5):509–528.
- Schranz, M., Di Caro, G. A., Schmickl, T., Elmenreich, W., Arvin, F., Şekercioglu, 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.
- Stein, M. L. (2002). The screening effect in kriging. *The Annals of Statistics*, 30(1):298–323.
- Thilan, A., Menéndez, P., and McGree, J. (2023). Assessing the ability of adaptive designs to capture trends in hard coral cover. *Environmetrics*, page e2802.
- 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.

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.

Autonomous Underwater Vehicle Sampling for Hotspots in Chlorophyll A

André J. H. Olaisen^{1*} and Jo Eidsvik¹

^{1*}Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway.

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

Abstract

With the goal of effective chlorophyll mapping in space and time using autonomous underwater vehicles, we study algorithms for hotspot sampling. We develop, implement and test a sampling algorithm using the expected improvement in chlorophyll to evaluate potential sampling designs. Working in a depth and lateral mode, the vehicle intelligently chooses the survey path for each sampling stage. The onboard system for design evaluation and data assimilation builds on a Gaussian random field representation of the log chlorophyll variations in north, east, depth, and time coordinates. Using a grid-free approach with spatio-temporal thinning of data, the random field model is quickly updated with in situ measurements made by the autonomous vehicle. This approach further enables real-time calculations for the expected improvement criterion used for sampling decisions. Via simulation studies, we compare the suggested sampling approach with others. Results of trials in the Mausund area, Norway, are used to demonstrate field deployments of our implementation for adaptive sampling. Here, we also discuss and compare the chlorophyll samples with that of zooplankton data gathered at the same locations.

Keywords: autonomous underwater vehicle, adaptive sampling, path planning, chlorophyll, expected improvement

Acknowledgments

We acknowledge support from Norwegian Research Council (RCN) through the MAS-COT project 305445 and SFI Harvest project 309661. The AUV and engineers to help the operation were provided by AURLab NTNU.

1 Introduction

Half of the oxygen we breathe comes from phytoplankton in the ocean (Field et al. 1998). Crucial for mankind and the primary productivity in the sea, phytoplankton is also essential for the development of grazing zooplankton and the entire food chain in the oceans. It is hence important to understand the mechanisms that drive phytoplankton, its correlations with zooplankton and biomass production, particularly in times of climate change, where one anticipates increasing ocean temperatures, potentially leading to reduced oxygen production (Sekerci and Petrovskii 2015).

A key enabler for the improved characterization of phytoplankton is the development of new sensor system technologies and algorithms to sample plankton in the ocean, and to couple these with models to gain an understanding of their spatial-temporal distribution. Phytoplankton dynamics can be modeled through coupled numerical ocean simulations that incorporate both physical and biochemical oceanographic processes. While such numerical ocean models provide much skill and interpretation, it is, however, difficult to correctly specify the initial conditions and forcing parameters. This problem may lead to bias in the predictive spatial-temporal distribution of plankton. Another source of phytoplankton information is satellite imagery data. Ocean color products computed from satellite data correlate with Chlorophyll A (ChlA), which is again indicative of the phytoplankton types and biomass (Brewin et al. 2017). Satellite data are, however, sparsely sampled and not always available due to cloud coverage and other factors such as turbulence and dissolved organic matter (Li et al. 2021). Moreover, phytoplankton blooms can be initiated far below the surface (Cullen 1982; Coon et al. 1987), and they are difficult to detect at an early stage using satellite data which reflect the situation near the surface (Cornec et al. 2021).

In situ measurements of water samples are highly informative and provide accurate data on plankton. These data are useful for understanding vertical migration and for calibrating ocean models or satellite data. However, water samples are not particularly useful for plankton mapping because they are sparsely sampled, providing data only at select points. One gains much flexibility by using an autonomous underwater vehicle (AUV) for sampling. AUVs enable in situ measurements of plankton at various depths, with the opportunity to move the AUV with its sensors to desired sampling locations, rather than staying at one location or drifting with the currents. AUVs are further equipped with computing capabilities, leveraging on-board data assimilation of the in situ observations in a model formulation, and using that updated model to find promising sampling designs leading the AUV to sense at new locations of potentially high plankton intensity.

There have been several studies recently on informative robotic sampling in the ocean. Flaspohler et al. (2019) used Monte Carlo tree search, entropy-based criteria and fast Gaussian random fields (GRFs) calculations to localize and sample coral hotspots with an autonomous surface vessel. Stankiewicz et al. (2021) applied an upper confidence bound and GRF calculations to map dissolved oxygen in the Severn River, Maryland. Berget et al. (2023) conducted variance-guided AUV sampling for mine tailings concentrations, using an onboard GRF model fitting an advection-diffusion process in space and time. Chen et al. (2024) showed benefits of having a

non-stationary Gaussian model for information gathering, using an autonomous surface vessel to improve underwater terrain maps in a quarry lake. Related to plankton and ChlA sampling, [Das et al. \(2015\)](#) conducted mean and upper confidence bound strategies for AUV sampling of phytoplankton abundance in Monterey Bay. [Manjanna et al. \(2018\)](#) relied on prediction variance for informative surface vessels sampling of chlorophyll in a Colorado reservoir. [Zhang et al. \(2019\)](#) used an AUV with spiral-type up and down movement to detect and track the deep chlorophyll maximum layer from ChlA and temperature measurements. [Fossum et al. \(2019\)](#) performed mean and variance-based AUV sampling for the ChlA maxima layers off the coast of Norway. [Zhang et al. \(2021\)](#) used an AUV to detect the deep chlorophyll maximum layer and drift with it while other sensors measured other parts of the phenomenon. [Fonseca et al. \(2023\)](#) relied on a GRF model and its gradients to conduct adaptive AUV sampling of ChlA in an algal bloom in the Baltic Sea. [Karakose and Bal \(2024\)](#) used a GRF model and upper confidence bounds for adaptive ChlA sampling off the coast of Alaska.

The approach outlined in this paper differs from these in the way we set up a sampling design and enable fast calculations for real-time large-scale operations. A GRF model is updated onboard the AUV with in-situ ChlA data. The AUV is instructed to explore and characterize places with high ChlA concentration. This is achieved by optimizing the expected improvement (EI) of the ChlA variable among possible direction and depth trajectories. Known to balance exploration and exploitation, our simulation studies show that EI is more effective for hotspot sampling than using other data acquisition criteria. Importantly, EI has a closed-form expression for GRFs, facilitating real-time operations. While spatio-temporal processes models are sometimes hampered by large grids and the required matrix computations, our implementation is scalable because it relies on keeping only data locations and values. The needed small-size grid directions are only formed when local sampling designs are evaluated. The main contributions of the paper are;

- real-time algorithm using EI for hotspot sampling of ChlA in 4D (north, east, depth and time),
- efficient Gaussian computations and model updating enabled by a grid-free approach and thinning of far-away data.
- new framework for design evaluations along spider-web trajectories in lateral coordinates and depth,
- simulation study comparing the outlined approach with others,
- field test in Norwegian coastal waters.

In Section 2, we provide a background for our case with plankton sampling in coastal areas of Norway. In Section 3, we outline the statistical model used for ChlA variability in space and time. In Section 4, we develop EI algorithms for adaptive AUV sampling of hotspots in ChlA. In Section 5, we study properties of the suggested algorithm for hotspot sampling through a simulation study. In Section 6, we show results of the AUV field trials for ChlA sampling in a coastal region in Norway. Here, we also compare and discuss the data and sampling results with other data sources. In Section 7, we conclude and point to future work.

2 Background

At the base of the marine food web, phytoplankton produce more than half of the atmospheric oxygen (Boyce et al. 2010). Yet, there is still much to learn about the mechanisms that drive the spatial-temporal distribution of phytoplankton. Some of the key factors influencing the distribution are sunlight availability, nutrients, temperature, ocean currents, and zooplankton grazing, among other organisms. The sunlight changes throughout the day and year. Nutrients are transported with complex ocean currents, and they can be depleted by the phytoplankton and other organisms. Eddies, current characteristics and internal waves are being regarded as very important for the vertical dynamics, but these uncertain drivers have complicated interactions. As a result of the complex factor relations, the spatio-temporal distribution of phytoplankton is heterogeneous and patchy (McGillicuddy and Franks 2019).

The patchiness is noticed in satellite data of ocean color or ChlA (Schalles 2006). It is also seen in numerical ocean models that mimic the spatial-temporal dynamics in the upper water column. Fig. 1(a) shows a satellite image¹. Even though this is at coarse resolution and with much uncertainty, one notices regions that contain more ChlA than others. Fig. 1(b) shows numerical ocean model results for *Calanus finmarchicus*, an abundant and important species of zooplankton. This is a result of using both physical and biological oceanography models to study the spatio-temporal variability in plankton which again form patchy uncertain structures driven by eddies and currents.

Images like in Fig. 1 provide some indication of the lateral distribution and movement of phyto- and zooplankton, but they rarely show much of the vertical distribution. Depth variability is largely driven by nutrient availability, sunlight, temperature, and water density (Wirtz and Smith 2020), which vary with location and season (Cullen and MacIntyre 1998). Zooplankton exhibit movement in the lateral plane driven by some of the same influences as phytoplankton. However, their vertical movement often differs. While their main diet is phytoplankton (Baars and Fransz 1984), they are also an important food source of several fish types, and they can migrate to large depths during the day to avoid visual predation (Lampert 1989). This daily migration, often referred to as the diel vertical migration (DVM), is influenced by several factors and can be stronger or weaker during different seasons.

We conduct AUV experiments to gain more insight into the spatio-temporal variation of ChlA and then put this in the context of phytoplankton distributions and its relations to zooplankton density. AUVs use fluorometers to obtain measurements of ChlA. Similar to what has been seen in satellite images and ocean models, patchiness of ChlA in AUV samples has been discussed in e.g. Fossum et al. (2019) sampling off the coast of Norway and in Tippenhauer et al. (2021) sampling in the front of Atlantic and Arctic waters.

Our AUV experiments were conducted in the Norwegian Sea, specifically in Frohavet near Mausund (Fig. 1) in late Spring and early Summer. In this region there are often several phytoplankton peaks during this season, going to various depths.

¹Global Ocean Color Plankton and Reflectances MY L3 daily observations. E.U. Copernicus Marine Service Information (CMEMS). Marine Data Store (MDS). DOI: <https://doi.org/10.48670/moi-00282>

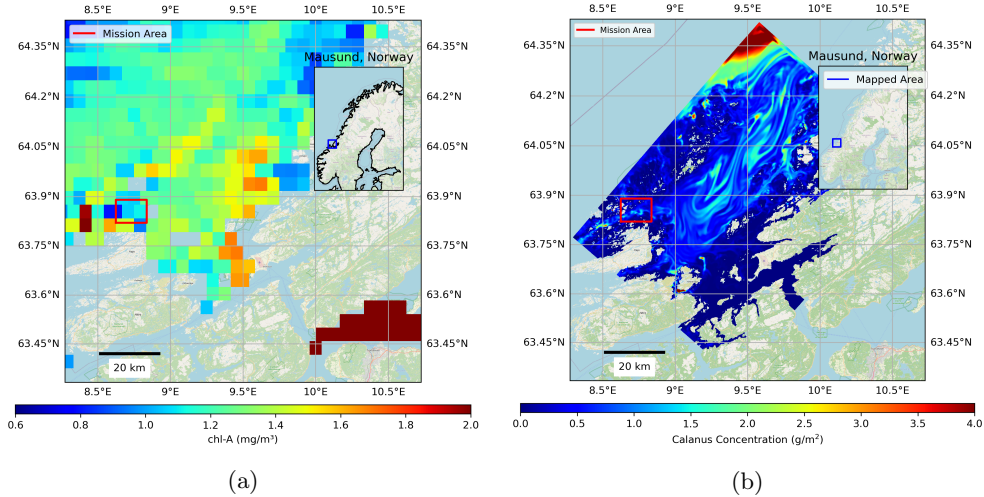


Fig. 1: (a) Estimated mass concentration of ChlA off the coast of Norway using satellite data on 2024.06.06 (Courtesy of E.U. Copernicus Marine Services). The estimate has a 4 km resolution. (b) Concentration of the zooplankton *Calanus finmarchicus* in the same area using a physical-biological oceanography numerical simulation. The simulation has a 160 m resolution. The field experiments were conducted within the red box depicted in both displays.

One reason why the phytoplankton subsides can be grazing by the species *Calanus finmarchicus* (Fragoso et al. 2024).

Building on what is seen in ocean models, satellite data and knowledge from DVM models, we mimic the spatio-temporal ChlA distribution using random field models. With explicit correlation structure, such models capture the patchiness of plankton. Limiting the scope to GRFs, we gain computational convenience because one can update the model and evaluate sampling designs with the computing units onboard the AUV in real-time.

In summary, the background modeling and sensor opportunities motivate the following steps for efficient plankton characterization;

- train a random field model to characterize the natural spatial distribution variability and patchiness of ChlA, as well as the vertical domain trends over time,
- update the random field description with in situ AUV measurements of ChlA,
- use the updated random field model to evaluate various AUV design strategies that guide the sampling effort to hotspots in ChlA,
- compare data results with that of other sensors, e.g. SilCam for zooplankton observations.

3 Model formulation

Ocean model data, satellite images and AUV measurements have shown that the distribution of ChlA in the ocean can be patchy, with strong spatio-temporal correlations. Therefore, we model the logarithmic ChlA as a GRF. The quantity is measured by the AUV at its sampling locations, and based on these measurements, a conditional GRF model is regularly updated during the mission. We here describe the mathematical notation, models and formulae for updating.

3.1 Gaussian random field model for ChlA

At location $\mathbf{s} = (s_x, s_y, s_z, t) \in \mathcal{D} \subset \mathcal{R}^4$, where (s_x, s_y) is lateral coordinate, s_z depth and t time, we denote log-ChlA by $x(\mathbf{s}) \in \mathcal{R}$. The logarithmic transform ensures a sample space on the real line for variables assumed to be Gaussian. Such an assumption has been used previously by e.g. Low (2009); Mo-Bjørkelund et al. (2024); Liu et al. (2026). The log-ChlA is modeled by a GRF, see e.g. Cressie and Wikle (2015). In its prior specification, before any AUV sampling efforts, the GRF model requires the specification of a mean $E(x(\mathbf{s})) = \mu(\mathbf{s})$ and variance $\text{Var}(x(\mathbf{s})) = \sigma^2(\mathbf{s})$ for $\mathbf{s} \in \mathcal{D}$, and a correlation function $\text{Corr}(x(\mathbf{s}), x(\mathbf{s}')) = \rho(|\mathbf{s} - \mathbf{s}'|)$ for $\mathbf{s}, \mathbf{s}' \in \mathcal{D}$. The mean is set from depth-dependent variation in ChlA (Wirtz and Smith 2020).

We use valid positive definite correlation functions, and can then formulate a Gaussian model for the log-ChlA $\mathbf{x}_Q = (x(\mathbf{s}_1), \dots, x(\mathbf{s}_n))$ at any size- n subset of locations $Q = (\mathbf{s}_1, \dots, \mathbf{s}_n)$:

$$\mathbf{x}_Q \sim \mathcal{N}(\boldsymbol{\mu}_Q, \boldsymbol{\Sigma}_Q). \quad (1)$$

Here, entries of the length- n mean vector are $\mu(\mathbf{s}_i)$ and matrix entry i, j of the $n \times n$ variance-covariance matrix is $\sigma(\mathbf{s}_i)\sigma(\mathbf{s}_j)\rho(|\mathbf{s}_i - \mathbf{s}_j|)$.

In our case, we assume stationary a priori variance σ^2 at all locations, and the correlation between two variables at locations \mathbf{s}_i and \mathbf{s}_j is defined to be separable in the lateral xy -plane, depth direction z , and time t . This means that

$$\Sigma_Q(i, j) = \sigma^2 \rho_{xy}(|\mathbf{s}_{i,xy} - \mathbf{s}_{j,xy}|) \cdot \rho_z(|\mathbf{s}_{i,z} - \mathbf{s}_{j,z}|) \cdot \rho_t(|\mathbf{s}_{i,t} - \mathbf{s}_{j,t}|), \quad (2)$$

where the distances between locations are computed in lateral, depth and time in the different correlation functions. Here, the spatial correlation functions (ρ_z and isotropic ρ_{xy}) are of a Matérn covariance function type with smoothness parameter $3/2$, similar to what was done in Fossum et al. (2019). The correlations at distance h are then $\rho_{xy}(h) = (1 + \phi_{xy}h) \exp(-\phi_{xy}h)$ for lateral domain and $\rho_z(h) = (1 + \phi_z h) \exp(-\phi_z h)$ for depth. The decay parameters ϕ_{xy} and ϕ_z are specified from ocean model data. Assuming the correlation in time is fluctuating according to a less smooth process, we use an exponential correlation function; $\rho_t(h) = \exp(-\phi_t h)$, with decay parameter ϕ_t again set from the data of the ocean models.

3.2 Conditioning on measurements of ChlA

Assume that we have measured ChlA data at m locations $\mathcal{S} = (\mathbf{s}_{\mathcal{S},1}, \dots, \mathbf{s}_{\mathcal{S},m})$. Denote the associated log-ChlA data by length- m vector $\mathbf{y}_{\mathcal{S}} = (y(\mathbf{s}_{\mathcal{S},1}), \dots, y(\mathbf{s}_{\mathcal{S},m}))$. These

data are modeled as conditionally independent, given the true ChlA values. We then have

$$y(\mathbf{s}_{S,k}) = x(\mathbf{s}_{S,k}) + \epsilon(\mathbf{s}_{S,k}), \quad k = 1, \dots, m, \quad (3)$$

where the zero-mean independent Gaussian noise terms $\epsilon(\mathbf{s}_{S,1}), \dots, \epsilon(\mathbf{s}_{S,m})$ have variance $\text{Var}(\epsilon(\mathbf{s}_{S,k})) = \sigma_\epsilon^2$.

Recall that the Gaussian multivariate distribution in Eq.(1) holds for any size- n and set \mathcal{Q} of locations. It could be identical to the measurement set \mathcal{S} or it could be different locations. Next, we define a prediction or design set \mathcal{P} with locations $\mathbf{s}_{\mathcal{P}} = (\mathbf{s}_{\mathcal{P},1}, \dots, \mathbf{s}_{\mathcal{P},n})$. In the adaptive sampling routine, we will be comparing several such sets of potential designs to optimize the AUV data acquisition towards hotspots. We are interested in predicting the ChlA variables at the design locations, i.e. $\mathbf{x}_{\mathcal{P}} = (x(\mathbf{s}_{\mathcal{P},1}), \dots, x(\mathbf{s}_{\mathcal{P},n}))$, given the available data $\mathbf{y}_{\mathcal{S}}$. Using standard formulae of the multivariate Gaussian distribution, the conditional distribution of $\mathbf{x}_{\mathcal{P}}$ given data $\mathbf{y}_{\mathcal{S}}$ is

$$\begin{aligned} \mathbf{x}_{\mathcal{P}} | \mathbf{y}_{\mathcal{S}} &\sim \mathcal{N}(\mathbf{m}_{\mathcal{P}|\mathcal{S}}, \mathbf{V}_{\mathcal{P}|\mathcal{S}}), \\ \mathbf{m}_{\mathcal{P}|\mathcal{S}} &= \boldsymbol{\mu}_{\mathcal{P}} + \boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}} (\boldsymbol{\Sigma}_{\mathcal{S}} + \mathbf{R})^{-1} (\mathbf{y}_{\mathcal{S}} - \boldsymbol{\mu}_{\mathcal{S}}), \\ \mathbf{V}_{\mathcal{P}|\mathcal{S}} &= \boldsymbol{\Sigma}_{\mathcal{P}} - \boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}} (\boldsymbol{\Sigma}_{\mathcal{S}} + \mathbf{R})^{-1} \boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}^t, \end{aligned} \quad (4)$$

where $\mathbf{R} = \sigma_\epsilon^2 \mathbf{I}_m$ is the covariance matrix of the measurement errors of log-ChlA. Moreover, $\boldsymbol{\Sigma}_{\mathcal{P},\mathcal{S}}$ is the size $n \times m$ cross-covariance matrix between log-ChlA variables at locations sets $\mathbf{s}_{\mathcal{P}}$ and $\mathbf{s}_{\mathcal{S}}$.

We note that the computational cost of these model updates depends on the number m of conditioning points and the number n of prediction points. In particular, matrix factorization of order $O(m^3)$ is needed. To ensure that computation cost remains moderate, one should keep the points in memory (\mathcal{S}) small and only predict the points (\mathcal{P}) that are relevant for the decision. Our framework keeps solely the observed points in memory and only the local design paths under consideration are predicted. This means that we do not use a regular grid covering the domain, but only points from the observed AUV trajectory and potential continuations of this path. This reduces the computation done at each step, but still keeps a high resolution where it matters. Figure 2 a) illustrates the framework; the blue points are in memory at step k , and the possible paths predicted are shown by the yellow points.

Even with this grid-free approach, computations can become infeasible to run on the AUV computer. Especially so when the set of locations we condition on grows very large. To keep the size of the conditioning set \mathcal{S} sufficiently small, we reduce its size during operation. This entails that conditioning points in space and time are thinned when the computation time exceeds a threshold. Thinning occurs in two ways: First, the sampling is typically denser than the spatio-temporal variation. Hence, every other point can be removed. Second, over operation time, we remove far-away points from the conditioning set. Due to the screening effect (Stein 2002), such thinning has very little effect on the accuracy of the Kriging predictions.

4 Sampling for hotspots

The goal of the AUV sampling is to find and explore hotspots of ChlA. Here, we use expected improvement (EI) as a criterion for choosing the sampling design. We first describe the EI and then put this into an adaptive sampling system.

4.1 Expected improvement

We define an improvement function that compares the predictive log-ChlA distribution at any sampling location to the largest log-ChlA measurement gathered so far. At the current stage of sampling, let x_{\max} denote the largest measurement so far in the sampling strategy. Building on the Gaussian predictive distribution for $x(\mathbf{s}_{p,i})$ at the current stage, see Expression (4), with mean m_i and variance v_i^2 , the probability of improvement is

$$P(x(\mathbf{s}_i) > x_{\max}) = \Phi\left(\frac{m_i - x_{\max}}{v_i}\right). \quad (5)$$

Here, $\Phi(z)$ is the standard normal cumulative distribution function. While the probability of improvement can guide sampling for hotspots, it is typically too focused on exploitation rather than exploration. Running this strategy one tends to prioritize sampling at locations in the vicinity of the current optimum, having probability near 0.5 of exceeding x_{\max} but very small variance. On these aspects associated with exploitation and exploration, there are benefits of instead running strategies based on expected values of improvements.

Letting the improvement at a location \mathbf{s}_i be $I[x(\mathbf{s}_i)] = \max[x(\mathbf{s}_i) - x_{\max}, 0]$, the EI is

$$E\{I[x(\mathbf{s}_i)]\} = (m_i - x_{\max})\Phi\left(\frac{m_i - x_{\max}}{v_i}\right) + v_i\phi\left(\frac{x_{\max} - m_i}{v_i}\right), \quad (6)$$

where $\phi(z)$ is the standard normal probability density function. The closed form solution of the expectation in (6) relies on a transformation of variables $z = (x(\mathbf{s}_i) - m_i)/v_i$, using that $\int z \exp(-z^2/2) dz = \exp(-z^2/2) + \text{const}$, and $\Phi(z) = 1 - \Phi(-z)$.

EI is commonly used in the optimization of complex functions (Jones 2001), where Bayesian optimization is used (Garnett 2023) over iterative stages of sampling. The EI has a tendency of declining over sampling stages. In situations where one can freely sample any location, an EI-based algorithm partly focuses sampling in high-concentration areas while also spanning the domain for potential hotspots. In our setting, the next-stage sampling design is constrained by the current AUV location and its possible trajectories. Depending on the spatio-temporal variation, the EI criterion strategy then sometimes guides the AUV to choose trajectories providing more local information near established hotspot areas, but at other times it sends the AUV to explore other parts of the domain. In doing so, the EI criterion provides sampling design trajectories that naturally balance exploration and exploitation.

4.2 Sequential sampling algorithm

The sampling is done sequentially. At each stage possible sampling designs are suggested with corresponding ChlA predictions conditional on previously observed

locations. This sequential design is shown in Fig. 2, where the blue points illustrate previously observed locations, the yellow points are possible design trajectories and the red points indicate the chosen next sampling design. These paths are evaluated

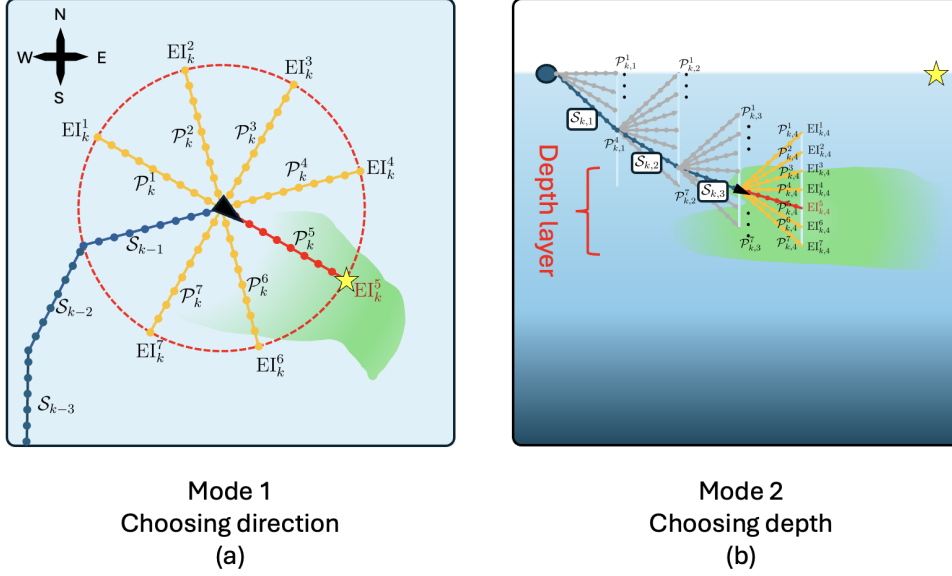


Fig. 2: Illustration of the adaptive path planning. (a) Mode 1, which is activated when the AUV (Black arrow) is at the surface and allows it to choose one of the seven possible lateral directions in which to move. (b) Mode 2, where the AUV moves in the selected direction from Mode 1 and now chooses the optimal depth along the way. The green area indicates locations with a high concentration of ChlA. The yellow star is the same point in both illustrations. It represents a surface location at the end of the transect.

based on the EI, and the AUV trajectory is hence adapted to the previous sampling locations and observations via the predicted mean and variance.

The sequential sampling procedure is split into two modes; Mode 1 (Fig. 2 (a)) chooses the lateral direction and Mode 2 (Fig. 2 (b)) chooses the sampling depth. In Mode 1 (Fig. 2 (a)), the agent is at the surface and makes a choice about which lateral direction to sample for the next stage. Here, the seven possibilities in a spider leg design are evaluated, and for the illustration in Fig. 2 (a) transect number 5 is selected (indicated in red). Mode 1 effectively chooses the direction and an endpoint for the transect. In Mode 2 (Fig. 2 (b)), the AUV adapts to sample at interesting depths while moving along the direction selected from Mode 1. In the second phase, again based on the EI, the best sampling depths are chosen (red line in the illustrative Fig. 2 (b)). The sampling depths are restricted by the diving angle of the AUV which again depends on its velocity and trim. For our field experiment the diving angle

range is 10 degrees giving 17 m vertical per 100 meters lateral. Based on this, we ended up evaluating nine possible depths. Depth flexibility is also restricted by the fact that the AUV returns to the surface at the end of the transect to get a GPS fix. For the practical implementation of the experiment, an important consideration is determining how long the AUV stays underwater. Equivalently, the distance of a transect (to yellow star in Fig. 2). As inaccuracies in the internal position estimate accumulate over time, especially with strong currents at the sampling location, too long a time under water can be risky. In the field experiment we used 800 m transects. Note also that the strategy with Mode 1 and Mode 2 means that the AUV avoids lateral turns when underwater, this can reduce the positioning error.

For a complete description of the suggested adaptive sampling algorithm, we need to introduce some more notation. Assuming we are at stage k of the sequential strategy, we denote the prediction points along a design path by $\mathcal{P}_k^j = (\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_m)$ for paths $j \in \{1, 2, \dots, J\}$. Further, all the previous sampling data are $\mathbf{y}(\mathcal{S}_{1:k-1})$ made at locations set $\mathcal{S}_{1:k-1}$ at stages $1, \dots, k-1$. Then we can use the onboard GRF model, see expression (4), to predict $\mathbf{x}(\mathbf{p}_i) | \mathbf{y}(\mathcal{S}_{1:k-1}) \sim \mathcal{N}(m_i, v_i^2)$ for $\mathbf{p}_i \in \mathcal{P}_k^j$. EI is used to select the best design path:

$$\text{Best path}_{EI,k} = \arg \max_{j \in \{1, 2, \dots, J\}} \max_{\mathbf{p}_i \in \mathcal{P}_k^j} \text{E} \{I[x(\mathbf{p}_i)]\} = \arg \max_{j \in \{1, 2, \dots, J\}} \text{EI}_k^j. \quad (7)$$

When the AUV has moved along the selected transect, the data dimension is increased and the onboard GRF has more data to assimilate for stage $k+1$, and so on.

Algorithm 1 summarizes the steps required for our sequential sampling and data assimilation. The algorithm also indicates that data are thinned over operation time, to avoid too large computation costs. At the core of the algorithm is the design evaluation and the selection of a best path. This has a function call with a separate last input parameter for Mode 1 and 2. Algorithm 2 summarizes the steps used to get the best path. Here, the EI evaluations along each potential transect is the critical acquisition function, providing the basis for optimal design choice and efficient adaptive sampling. When we compare the EI results with that of using other acquisition functions, this line is changed in the code.

5 Simulation Study

5.1 Simulation Setup

In the simulation study, we assume that an AUV moves over time in a 3D box measuring $4 \text{ km} \times 4 \text{ km}$ with a maximum depth of 75 m. In this box, we generate 100 different time-space varying random fields using the GRF described in Section 3 with $\phi_z = 1/8$, $\phi_{xy} = 1/400$ and $\phi_t = 1/14400$. The prior mean varies with depth and time of day, mimicking a DVM model. It has no trend in lateral space. The AUV has different fixed start points in each of the 100 fields. The AUV has 4 hours to sample inside the domain, and it moves with a speed of 1 m/s. To mimic real conditions, the AUV must return to the surface periodically, in this case, every 1000 m operation distance. While sampling a transect, the AUV can dive with a maximum angle of 13° .

Algorithm 1 Sampling ChlA hotspots in stages

Require: Current position \mathbf{s}_{now} , mean field $\mu(\mathbf{s})$, standard deviation σ , correlation function $\rho(h)$, and operational domain \mathcal{D} , transect length L_{trans} , depth seeking step length L_{depth} , maximum update time t_{update} .

- 1: Initialize $\mathcal{S}_0 = \emptyset$, $\mathbf{y}_0 = \emptyset$, $x_{\text{max},0} = 0$.
- 2: **for** each time step $k = 0, 1, \dots$ **do** ▷ Loop over stages
- 3: **Mode 1:** ▷ Fig 2 (a)
- 4: $\mathcal{P}_k = \text{GetBestPath}(\mathbf{s}_{\text{now}}, L_{\text{trans}}, \mathcal{S}_{1:k-1,0}, \mathbf{y}_{1:k-1,0}, \mathbf{m}_{1:k-1,0}, x_{\text{max}}, 1)$ ▷ Eq (7)
- 5: Set WP_{goal} at the surface at the end of \mathcal{P}_k
- 6: **for** each depth step $l = 0, 1, \dots$ until \mathbf{s}_{now} is close to WP_{goal} **do**
- 7: **Mode 2:** ▷ Fig 2 (b)
- 8: $\mathcal{P}_{k,l} = \text{GetBestPath}(\mathbf{s}_{\text{now}}, L_{\text{depth}}, \mathcal{S}_{1:k-1,l-1}, \mathbf{y}_{1:k-1,l-1}, \mathbf{m}_{1:k-1,l-1}, x_{\text{max}}, 2)$
▷ Eq (7)
- 9: Move agent along path $\mathcal{P}_{k,l}$ collecting measurements $\mathbf{y}_{k,l}$ on locations $\mathcal{S}_{k,l}$
- 10: Update $\mathcal{S}_{1:k,l} = (\mathcal{S}_{1:k,l-1}, \mathcal{S}_{k,l})$ and $\mathbf{y}_{1:k,l} = (\mathbf{y}_{1:k,l-1}, \mathbf{y}_{k,l})$
- 11: Compute $\mathbf{m}_{\mathcal{S}_{1:k,l}}$ and $\mathbf{V}_{\mathcal{S}_{1:k,l}}$ conditional on data $\mathbf{y}_{1:k,l}$ at $\mathcal{S}_{1:k,l}$ ▷ Eq. (4)
- 12: Update $x_{\text{max}} = \max[\mathbf{m}_{\mathcal{S}_{1:k,l}}]$
- 13: **if** Time to update model is greater than t_{update} **then**
- 14: The points in memory are thinned
- 15: **end if**
- 16: **end for**
- 17: **end for**

The depth-seeking step length is set to $L_{\text{depth}} = 100$ m, meaning the AUV chooses new depths 10 times during each transect.

For each of the replicate fields, adaptive AUV sampling is conducted with the suggested EI approach (Algorithm 1 - 2). We also compare the EI hotspot sampling criterion with other methods relying on the same type of sequential updating sketched in Fig. 2 and described in Algorithm 1. For this purpose, we simply switch EI with other acquisition functions, enabling a seamless comparison. We refer to this as using different agents. We here study

- maximum variance; picking the direction with the largest variances
- maximum expected intensity; picking the direction with largest expectation values
- lawnmower; systematically moving through the domain, spanning the area
- probability of improvement (see Equation 5),
- EI (see Equation 6)

All strategies are adaptive, except the pre-scripted lawnmower strategy.

Example paths of four agents are shown in Fig. 3. The displays show results for one of the replicate runs. The start, end and sample paths are illustrated in a map view, while the depths of the agents are indicated by the color-codes. We observe characteristic behavior for the different agents. Fig. 3 (a) illustrates how the lawnmower follows a yo-yo path; down to the maximum depth and up again, etc., along the path. In Fig. 3 (b), the maximum variance agent spans the spatial domain, moving to the maximum depth, and then back up, while attempting to sample locations far away

Algorithm 2 GetBestPath

Require: The current position \mathbf{s}_{now} , the distance of the path L , previous data $\mathbf{y}_{\mathcal{S}}$ at points \mathcal{S} with estimated intensity $\mathbf{m}_{\mathcal{S}}$, maximum estimated intensity x_{max} and the *mode* for making the paths.

- 1: Define the paths $\{\mathcal{P}^j\}_{j=1}^J$ from \mathbf{s}_{now} with a length L according to the *mode* ▷ Modes illustrated in Fig 2
- 2: Set $\text{EI}_{\text{max}} = 0$
- 3: Set $\mathcal{P}^* = \emptyset$
- 4: **for** each depth path $\mathcal{P}^j, j = 1, \dots, J$ **do** ▷ Design evaluation
- 5: Predict $\mathbf{m}_{\mathcal{P}^j}$ and $\mathbf{V}_{\mathcal{P}^j}$ conditional on data $\mathbf{y}_{\mathcal{S}}$ at \mathcal{S} . ▷ Eq. (4)
- 6: Compute $\text{EI}^j = \max_{\mathbf{p}_i \in \mathcal{P}^j} E\{\max[x(\mathbf{p}_i) - x_{\text{max}}, 0]\}$ ▷ Eq. (6)
- 7: **if** $\text{EI}^j > \text{EI}_{\text{max}}$ **then**
- 8: $\text{EI}_{\text{max}} = \text{EI}^j$
- 9: Set $\mathcal{P}^* = \mathcal{P}^j$
- 10: **end if**
- 11: **end for**
- 12: **Return** \mathcal{P}^* ▷ Returning the best path

from those previously sampled. Fig. 3 (c) shows the EI agent, which balances exploration and exploitation. It spans a rather large part of the domain in this case. The survey depth varies along the path, occasionally doing more exploration to full depth. The maximum expected intensity agent is shown in Fig. 3 (d). This agent tries to sample the region with the highest expected values in ChlA, and if it finds a region with high values, it can get stuck there. In the plot, we can see an extreme case of this greedy behavior where the agent is very focused in the south-eastern parts of the domain. Notice that it also remains at nearly the same depth, except when it needs to rise to the surface. The agent shows a high degree of exploitation, without exploring new locations.

We compute several metrics to study the properties of the various agents. When evaluating performance, we are mainly interested in assessing two attributes: i) the agents' ability to find and measure hotspots, and ii) how well these regions are explored. For each replicate run and for each agent, we evaluate metrics related to i) and ii). The comparison of different agents is then based on attribute statistics over the 100 replicate runs.

For the ability to find and measure hotspots, we register the proportion of time that the AUV samples at a location that can be considered a hotspot in ChlA. Here, a hotspot is defined as a location with a log-ChlA value above the 90th percentile of the distribution for the whole field and the entire run length. The computed attribute is then

$$\text{Prop}(x_{0.90}) = \frac{1}{n_{\text{samples}}} \sum_{i=1}^{n_{\text{samples}}} \mathbf{1}_{x(\mathbf{s}_i) > x_{0.90}}, \quad (8)$$

where the indicator function $\mathbf{1}_A$ equals 1 if event A occurs and is 0 otherwise.

We also look at the conditional variance for regions that are above the 90th percentile threshold, as provided by the ground truth of each replicate. At regular time

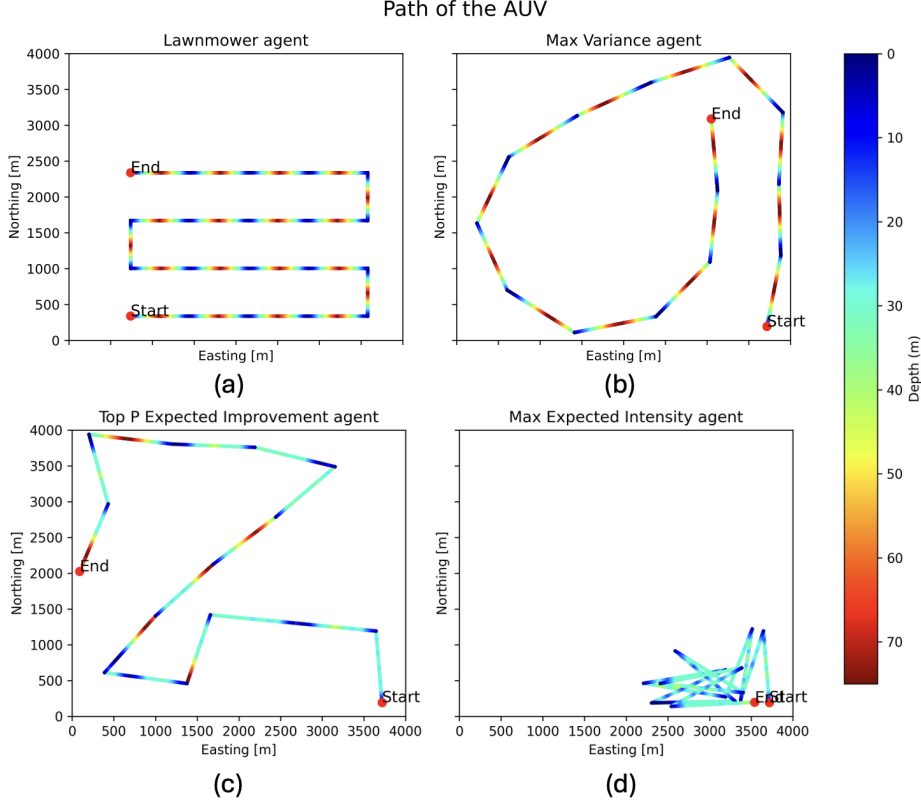


Fig. 3: Example paths of four different agents on the same replicate ChlA field. The start and end points are marked, and the depth is indicated by the color-code along the path.

intervals, we compute the conditional variance for these locations using equations (4). The mean variance for the points above the threshold is recorded.

Related to the second performance attribute the ability to explore distinct high ChlA value regions is explored. The domain is split into N smaller non-overlapping regions $\mathbf{R} = \{R_i \subset \mathcal{D} | i = 1, \dots, N\}$. The dimensions of each of these regions R_i are $h_x \times h_y \times h_z \times h_t$, with lengths defined using the correlation functions $h_t = \rho_t^{-1}(0.5)$, $h_x = h_y = \rho_{xy}^{-1}(0.5)$ and $h_z = \rho_z^{-1}(0.5)$ from equation (2). This choice means that regions are split up such that measuring one region only gives partial information of a neighboring region. For each of the regions, we define $\bar{\mathbf{x}}_{R_i}$ as the average value. Then the hotspot regions are defined as $\mathbf{H} = \{R_i \subset \mathbf{R} | \bar{\mathbf{x}}_{R_i} > \bar{\mathbf{x}}_{0.9}\}$, e.g. the 10 % regions with the highest average values. Thus, we count how many hotspot regions an agent has visited along a path \mathcal{S} using

$$\mathbf{HV}_{\mathcal{S}} = \{R_i \subset \mathbf{H} | \exists \mathbf{s} \in \mathcal{S} \text{ such that } \mathbf{s} \subset R_i\}. \quad (9)$$

An illustration of the hotspot regions \mathbf{H} , hotspot regions visited \mathbf{HV}_S , and the AUV path for a 2D illustration can be found in Figure 4. The number of hotspots is here $n(\mathbf{H}) = 10$ and the number visited is $n(\mathbf{HV}_S) = 6$. Note that this display is for illustrative purposes only. In the simulation study these boxed also have depth and time dimensions. From this illustration there are 2 distinct hotspots clusters. These clusters can be defined as a set of hotspot regions neighboring at least one other in the same cluster, therefore each hotspot region belongs to one and only one cluster $\mathcal{C}_k = (H_1^{(k)}, H_2^{(k)}, \dots, H_j^{(k)})$. We keep track of how many clusters $n(\mathbf{CV}_S)$ that are visited by the AUV, and build on this to study how well each cluster is explored. The proportion of regions visited in a cluster is then $n(\mathbf{HV}_S \in \mathcal{C}_K)/n(\mathbf{H} \in \mathcal{C}_k)$. Because the various regions have very different sizes, we look at the weighted exploration of each cluster;

$$WCE = \frac{1}{\sum_{\mathcal{C}_k \in \mathcal{C}} \log(n(\mathcal{C}_k))} \sum_{\mathcal{C}_k \in \mathcal{C}} \frac{n(\mathbf{HV}_S \in \mathcal{C}_k)}{n(\mathcal{C}_k)} \cdot \log(n(\mathcal{C}_k)). \quad (10)$$

Here, $n(\cdot)$ simply counts the elements in the set. In summary, we look at how many clusters are explored and how well the clusters are explored.

5.2 Simulation Results

The main results of the simulations are shown in Fig. 5 and 6. Fig. 5 (a) shows the proportion of sampling time above the threshold as defined in Equation (8). Fig. 5 (b) shows the mean conditional variance for regions in the 90th percentile. In display a) we notice that the plot increases for all acquisition functions except lawnmower and the maximum variance strategies, the latter settles at spending about 10 % of the time in the 90th percentile regions. Maximum EI increases the fastest and stabilizes at the highest level, surpassing the probability of improvement and the maximum expected intensity. One might expect that maximum expected intensity would perform a bit better here, but it can often get stuck in local maximum that are below the threshold value. Note that the shape of EI and other curves fluctuates at the rate of surfacing. When doing so, the agent goes away from the interesting areas. The results look a bit different when looking at Fig. 5 (b). For all agents there is a steady drop in variance at true hotspot locations. Not surprisingly maximum expected intensity drops the slowest because it does not explore well. While the lawnmower and maximum variance spend little time in the 90th percentile of the field, the conditional uncertainty of these regions are comparable with agents that spent much more time in these regions, like probability of improvement and EI. On average EI still performs slightly better. One can also look at the conditional variance of the whole field and not just the top 90th percentile, then maximum variance and lawnmower perform better than EI.

Fig. 6 shows the mean proportion of visited clusters (a) and the mean weighted cluster exploration (b) over the operation time. The maximum variance ends up visiting the most clusters, but it does not explore them very well. EI visits the second most clusters and explores them the best. Interestingly, the strategy aiming for the maximum expected intensity of ChlA is doing quite well in Fig. 5, but in Fig. 6 the line is almost flat, meaning that once it finds a cluster it fails to explore more of the field.

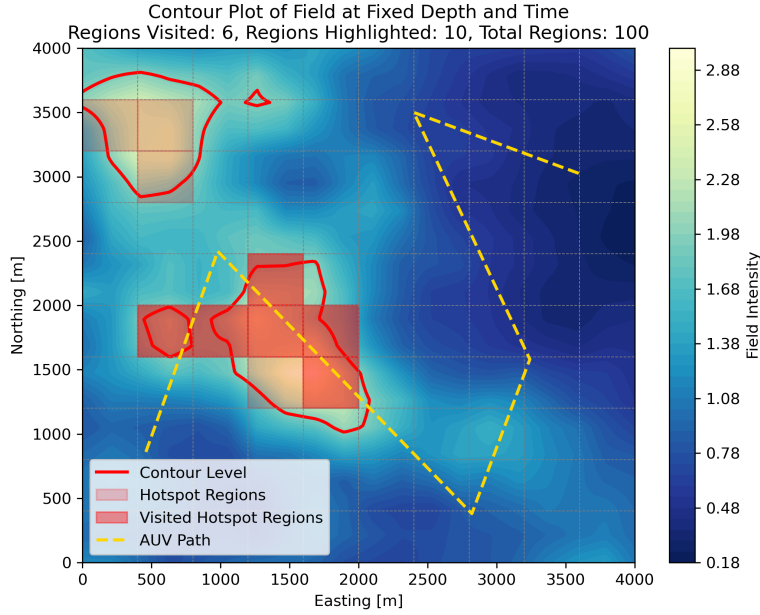


Fig. 4: Map view of a simulated field. The yellow line shows the path of the AUV, the red boxes show the hotspot regions \mathcal{R} , the highlighted red boxes show the hotspot regions visited by the AUV, the red line is a contour line at the 90 percentile. This also shows two separate hotspot clusters. Note that this is just a 2D illustration, in the simulation study depth and time is also taken into account.

As mentioned, one can measure other ocean variables with sensors on the AUV. It is then worth asking if additional data would be informative of ChlA. To study this closer, we analyzed field data of ChlA, temperature and salinity. There appears to be a significant correlation between ChlA and the vertical temperature gradient. That is, ChlA tends to be large then the temperature change with depth is not very negative, while ChlA gets smaller when the temperature declines with depth. There is also some correlation with salinity, but not so clear. Based on this, we construct a bi-variate GRF model with temperature change with depth and log-ChlA. This bivariate model is build on similar correlation functions for each variable and a Kronecker product assumption so that the variances remain stationary and the cross-correlation between temperature change with depth and log-ChlA remain constant at all locations. Measurements of both attributes are used to update the bi-variate fields, assuming the yo-yo movement of the AUV would gauge the temperature gradient accurately. EI in log-ChlA was used to guide the sampling. Results indicate that EI strategies with the bi-variate fields are reducing the variance for hotspot regions if there is large variance in log ChlA measurements and high cross-correlation between log ChlA and temperature gradient. However, for realistic levels of ChlA measurement error, the cross-correlation values




Fig. 5: a) Mean proportion of the agents' sampling above the threshold for the duration of the simulation. The shaded area is a 95 % confidence interval for the mean computed over the 100 replicates. Some of the lines are a bit jagged due to the agents periodically going to the surface, rather than where the intensity is the highest. b) Conditional mean variance above the 90th percentile threshold




Fig. 6: a) Mean proportion of clusters visited and b) Weighted exploration of the clusters.

must typically exceed 0.9 for this bi-variate modeling to have a visible effect. In our field data, there was not more than 0.5 cross-correlation.

Going beyond multi-variate modeling, it is here interesting to conduct multi-objective optimization (Gunantara 2018), where one is interested in hotspots of multiple variables. On this topic, the expected hypervolume improvement is a popular acquisition function, see e.g. Yang et al. (2019), but it can be challenging to compute this quickly in spatio-temporal settings that we have considered here.

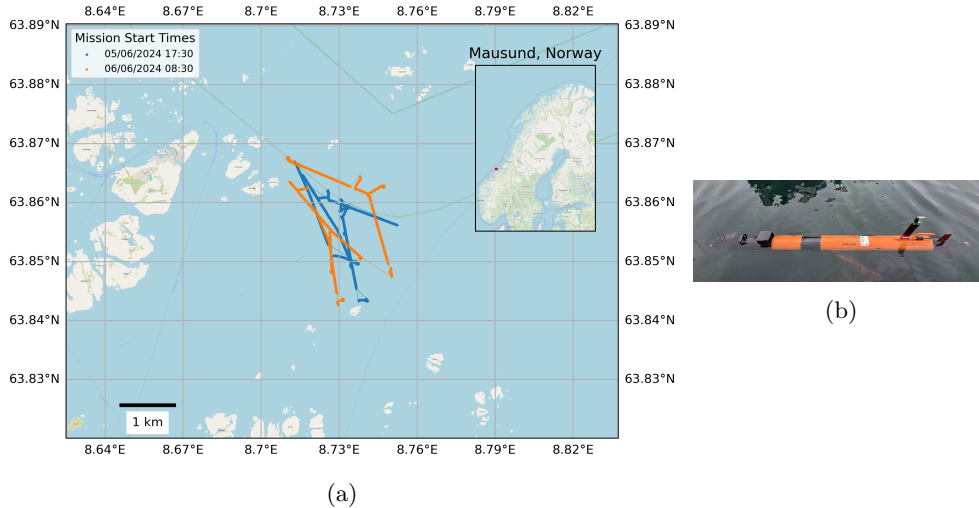


Fig. 7: (a) The path of the AUV in latitude and longitude for the two deployments. The deployments happened on June 6 (blue line) and 7 (orange line). (b) The AUV used for the deployments.

6 Results of mission in Norwegian coastal waters

We now describe results from the AUV field deployment in Norway.

6.1 Description

The experiments were carried out in early Summer of 2024. The survey area is off the west coast of mainland Norway (Fig. 7 (a)), about 100 km from Trondheim. We conducted surveying on two different days.

In both deployments, the adaptive EI algorithm (see Algorithm 1-2) was applied to select the path directions and survey depths, see Fig. 2. These experiments were part of a larger field deployment that involved other vehicles. One of the overall goals was to run several robots with only a few engineers in the field while most of the vehicle operators and the scientific staff were working remotely from a control center in Trondheim (Barstein 2024). Such operations could make future missions more cost-effective and minimize risk.

6.2 AUV setup and software

The AUV used in the two deployments is shown in Fig. 7 (b). It is a Light Autonomous Underwater Vehicle (LAUV). 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 System (ROS) (Quigley 2009) and DUNE (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. Pre-launch protocol tests consisted of standard remote control verification.

During the missions, the AUV was carrying several sensors. For detecting ChlA concentration, a RBR Tuner Cyclops7 Fluorometer sensor was used. This the sensor used for adaptive sampling. Additionally, used for off-line comparison, an AML Smart X CTD was used to measure salinity and temperature, and a SilCam was used for detecting zooplankton copepods. This SilCam was developed by SINTEF and mounted on the AUV. It works by taking pictures of a small water volume that is illuminated from below, similar to bright-field microscopy, but with lower magnification.

6.3 Adaptive sampling mission

For the field experiments the maximum depth was set at 70 m, $L_{\text{trans}} = 800$ m, maximum submerge time 15 minutes, dive angle to 10° and AUV speed 1.1 m/s. The maximum depth was set based on the topography of the seabed and L_{trans} was set so that the AUV could reach the maximum depth of 70 m with a 10° and return to the surface. The maximum submerge time of 15 minutes is more of a safety mechanism if the AUV fails to make fast enough progress towards the end of the transect. The number of directions for Mode 1 was set to 7 and the number of depths for Mode 2 was set to 9.

Fig. 8 shows the diving depths for the AUV for the two experiments on two different days. As described in Section 4.2, the AUV started out with a yo-yo transect to a maximum depth of 70 m. After this initial dive, the AUV was free to dive at whatever depth was best according to the adaptive sampling approach described in Algorithm 1. After the initial dive, it never returned as deep as 70 m but rather stayed at depths around 10-30 m, which is where the ChlA measurements were highest (color-coded in the displays). Note that the AUV was programmed to return to the surface every 15 minutes or 800 m, whichever came first, to correct the position using GPS. In one of the missions, there was also a longer pause where the AUV stayed on the surface. The implementation assumed deterministic values for estimated pose, but this could likely be extended by applying a kernel weight accounting for the pose error estimate from e.g. the extended Kalman filter or unscented Kalman filter (Allotta et al. 2016).

In Fig. 9, scatter plots of ChlA measurements and depth are visualized for both survey days. The displays indicate the depths with the highest ChlA concentration. The AUV sampled quite clear layers with high ChlA concentrations in both missions. On the first day, Fig 9 (a), the highest concentration is found at 0-25 m depth, while it is somewhat deeper at around 10-30 m for the second survey day shown in Fig 9 (b). It also appears that the layer is somewhat narrower on the second day.

6.4 Other sensors

Off-line, we looked at the other data attributes gathered during the missions; temperature, salinity and zooplankton copepod counts captured by the SilCam. We now discuss their connections to ChlA.

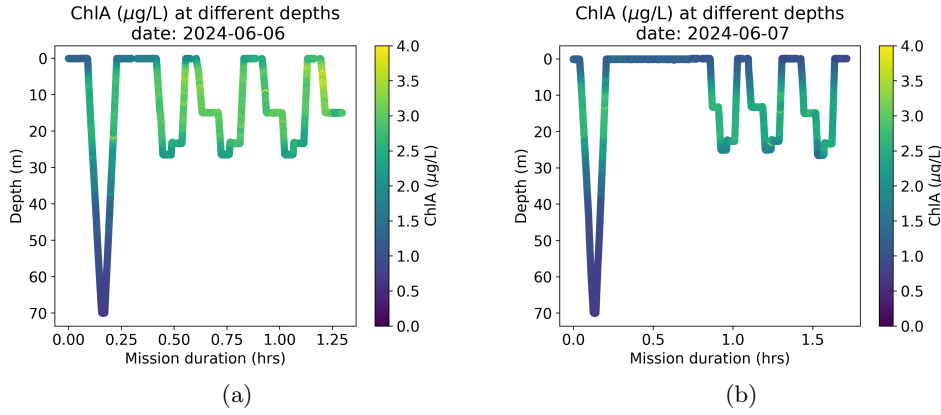


Fig. 8: Depth of the AUV during the two missions, (a) 2024.06.06 and (b) 2024.06.07. For the first dive, the AUV was instructed to go down to 70 m. Subsequently, the depth was chosen adaptively. In this case, the AUV focused on depths ranging from 10 to 30 m on both survey days. The colors indicate the measured ChlA, indicating a layer of ChlA.




Fig. 9: Scatter plot showing the depths with the highest ChlA measurements for both missions. There is a clear layer where the concentration was higher. On 2024.06.06 (a), it was at a depth of 0-25 m. On 2024.06.07 (b) it was around 10-30 m.

We first investigate how salinity and temperature vary with depth during the two missions. In Fig. 10, we notice a transition in temperature and salinity (first axes in the displays) at around 40 m. Especially for temperature, this is rather sharp and it could indicate a layer shift from almost constant temperatures to much colder regimes. This seasonal thermocline is typical during Spring in the North Atlantic (Brody and Lozier 2015), creating a separation between the warmer fresh water at the top and the colder salt water in the deep. The thermocline remains relatively stable during the




Fig. 10: Salinity and temperature measurements against the depth for the two experiments, (a) 2024.06.06 and (b) 2024.06.07, the color indicates the ChlA measurements.

two experiments (Fig. 10a-b). The ChlA concentration (color-coded in the displays) is rapidly decreasing when moving deeper than the well-mixed zone in the upper layer of the two water masses. This correlation in ChlA and vertical temperature gradient inspired our bivariate studies in the simulation study, and we think more nuanced bivariate modeling and multi-objective optimization methods would be useful to better understand this phenomenon.

Next, we study the distribution of copepods captured by the SilCam. Specifically, we look at *Calanus Finmarchicus* which is an abundant zooplankton species of *Calanus* in the survey area. *Calanus* may graze on the available phytoplankton. The frequency of the SilCam was set at 1 Hz during the experiments. The raw SilCam images were processed off-line in two steps: the first is to segment the images into relevant particles, and then use a convolutional neural network (CNN) to classify the particles.

Fig. 11 shows some of the pictures taken from the AUV during the two missions. We compare these with images taken in a laboratory setting. The laboratory images are much clearer than those taken in the field deployments, but we still recognize the general shape and size of the zooplankton. Note that the pictures from the field in Fig. 11 (b) are in various orientations, and for some of them, the thinner parts of the *Calanus* are partially or fully cropped by the processing software.

After examining the particles captured by the AUV SilCam, we filter out those that are likely *Calanus*. These images are then time-stamped, allowing us to correlate them with the data from other sensors. By merging data from the two deployments, we study relations between *Calanus* count and depth (Fig. 12 a) as well as *Calanus* count and ChlA (Fig. 12 b). For the comparison between depth and *Calanus* count, the leftmost display shows the mean number of *Calanus* per image for various depths (second axis). The mean (red) is estimated by fitting a Poisson generalized additive model (GAM) (Rigby and Stasinopoulos 2005). A 90 % bootstrap confidence interval is shown in shaded blue. The rightmost display in a) shows a scatter plot of the number

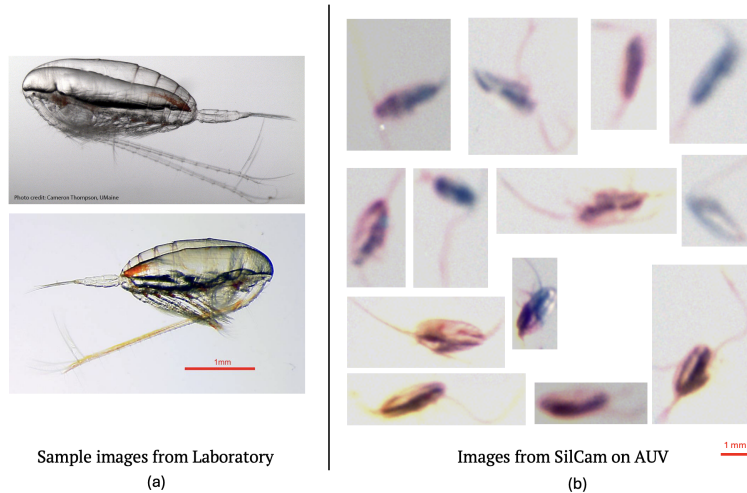


Fig. 11: (a) Images of *Calanus Finmarchicus* from a lab setting (Photo: top *Cameron Thompson/UM-Maine*, bottom: *Dag Altin, BioTrix*). (b) Images that could be *Calanus Finmarchicus* taken from the AUV using the onboard SilCam. The images in (b) are all on the same scale.

of *Calanus* that can be identified per image and the depth at which the images are taken. Similar to the ChlA, the *Calanus* intensity seems to form a layer at shallow depths, around 10-25 m. Fig. 12 b) shows a GAM fit of the mean number of *Calanus* per image as a function ChlA (second axis) and a scatter plot of *Calanus* count and ChlA measurements at the location where the image was captured. Note that the bootstrap interval is cropped for large values of ChlA because the uncertainty is very high. The plots strengthen our belief that there were more *Calanus* at places with high ChlA values. This suggests that there was a layer of *Calanus* grazing on the phytoplankton near the surface during the experiments.

There are some drawbacks to using this SilCam on the AUV that are worth mentioning. First, the AUV must maintain a speed of at least 1 m/s for maneuverability. However, this speed is still quite high for capturing sharp SilCam images. Hence, it turns into a balance of exposure time, motion blur, and image noise. This makes accurate species identification difficult. Second, there is the possibility that the *Calanus* has an avoidance reaction to the hydrodynamic disturbances caused by the AUV (*Buskey et al. 2002*), resulting in a rapid movement away from the AUV and the imaging chamber. Despite these challenges, we foresee that real-time operations assimilating SilCam data as well as ChlA data (and temperature and salinity) would be a very interesting avenue for gaining further insight in the multivariate dependence and for doing multi-objective adaptive AUV sampling. With the mentioned issues, it is however challenging to transition from the relative intensity of *Calanus* to an estimate of biomass or abundance. Only via net samples in select AUV sampling locations could one attempt to calibrate the AUV measurements.




Fig. 12: Calanus counts from SilCam images for various depths (a) and for ChlA values (b). The images with the most Calanus can be found at depths of 10-25 m and at ChlA levels of around 2-4. The red line indicates the mean Calanus per image, the shaded area is a 90 % bootstrap confidence interval.

7 Closing remarks

We have presented an adaptive AUV sampling algorithm designed to efficiently map hotspots of ChlA in both space and time. The algorithm relies on a spatio-temporal GRF model running onboard the AUV. This model is updated with new ChlA measurements and fast computations of the EI acquisition function are used when the AUV chooses its next design path segment.

A key merit of our approach is its scalability - it relies on data values alone and no big grids. It forms the needed small-size grid directions on the fly, for the required sampling evaluations. For long-time operations, we further thin the data samples and remove far-away data.

A simulation study demonstrates that using EI yields more efficient sampling than other criteria. In the field test off the coast of Norway, adaptive sampling guides the AUV to depths around 5-30 m, where there appears to be the most ChlA. Off-line, we relate the ChlA samples to zooplankton images made in the field. Despite the difficulties of labeling image data, we notice clear correlations between the intensity of high-quality zooplankton images made by a SilCam and the high ChlA sample values.

In the future, we aim to connect more of the multivariate ocean data for longer-term sampling. This entails more nuanced prior modeling of both the ChlA variations and Calanus density, and fitting realistic multivariate space-time variation in the GRF models. We also aim to improve image classification algorithms for the intensities of various plankton types, and embed these in real-time operations decisions for effective sampling. For future tests, several high-quality net samples should be done to correlate with the different sensors. Only then can one aim to grasp more of the mechanisms underlying the plankton migrations. Note that there can also be a feedback loop here,

where the AUV samples and associated predictions inform us about where and when to gather water samples.


References

- Allotta B, Caiti A, Costanzi R, et al (2016) A new auv navigation system exploiting unscented kalman filter. *Ocean Engineering* 113:121–132
- Baars M, Fransz H (1984) Grazing pressure of copepods on the phytoplankton stock of the central north sea. *Netherlands Journal of Sea Research* 18(1):120–142
- Barstein KH (2024) Robots collaborate on remote sampling of zooplankton. URL <https://blog.sintef.com/ocean/robots-collaborate-on-remote-sampling-of-zooplankton/#comments>
- Berget GE, Eidsvik J, Alver MO, et al (2023) Dynamic stochastic modeling for adaptive sampling of environmental variables using an auv. *Autonomous Robots* 47(4):483–502
- Boyce DG, Lewis MR, Worm B (2010) Global phytoplankton decline over the past century. *Nature* 466(7306):591–596
- Brewin RJ, Ciavatta S, Sathyendranath S, et al (2017) Uncertainty in ocean-color estimates of chlorophyll for phytoplankton groups. *Frontiers in Marine Science* 4:104
- Brody SR, Lozier MS (2015) Characterizing upper-ocean mixing and its effect on the spring phytoplankton bloom with in situ data. *ICES Journal of Marine Science* 72(6):1961–1970
- Buskey E, Lenz P, Hartline D (2002) Escape behavior of planktonic copepods in response to hydrodynamic disturbances: high speed video analysis. *Marine Ecology Progress Series* 235:135–146
- Chen W, Khardon R, Liu L (2024) Adaptive robotic information gathering via non-stationary gaussian processes. *The International Journal of Robotics Research* 43(4):405–436
- Coon T, Lopez M, Richerson P, et al (1987) Summer dynamics of the deep chlorophyll maximum in lake tahoe. *Journal of Plankton Research* 9:327–344
- Cornec M, Claustre H, Mignot A, et al (2021) Deep chlorophyll maxima in the global ocean: Occurrences, drivers and characteristics. *Global Biogeochemical Cycles* 35(4):e2020GB006759
- Cressie N, Wikle CK (2015) *Statistics for Spatio-Temporal Data*. John Wiley & Sons

- Cullen JJ (1982) The deep chlorophyll maximum: Comparing vertical profiles of chlorophyll a. *Canadian Journal of Fisheries and Aquatic Sciences* 39(5):791–803
- Cullen JJ, MacIntyre JG (1998) Behavior, physiology and the niche of depth-regulating phytoplankton. *Nato Asi Series G Ecological Sciences* 41:559–580
- Das J, Py F, Harvey JB, et al (2015) Data-driven robotic sampling for marine ecosystem monitoring. *The International Journal of Robotics Research* 34(12):1435–1452
- Field CB, Behrenfeld MJ, Randerson JT, et al (1998) Primary production of the biosphere: Integrating terrestrial and oceanic components. *Science* 281(5374):237–240
- Flaspohler G, Preston V, Michel AP, et al (2019) Information-guided robotic maximum seek-and-sample in partially observable continuous environments. *IEEE Robotics and Automation Letters* 4(4):3782–3789
- Fonseca J, Bhat S, Lock M, et al (2023) Adaptive sampling of algal blooms using autonomous underwater vehicle and satellite imagery: Experimental validation in the baltic sea. *arXiv preprint arXiv:230500774*
- Fossum TO, Fragoso GM, Davies EJ, et al (2019) Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics* 4(27):eaav3041
- Fragoso GM, Dallolio A, Grant S, et al (2024) The role of rapid changes in weather on phytoplankton spring bloom dynamics from mid-norway using multiple observational platforms. *Journal of Geophysical Research: Oceans* 129(3)
- Garnett R (2023) *Bayesian Optimization*. Cambridge University Press
- Ge Y, Eidsvik J, Mo-Bjørkelund T (2023) 3d adaptive auv sampling for classification of water masses. *IEEE Journal of Ocean Engineering* 48:626–639
- Gunantara N (2018) A review of multi-objective optimization: Methods and its applications. *Cogent Engineering* 5(1):1502242
- Jones DR (2001) A taxonomy of global optimization methods based on response surfaces. *Journal of Global Optimization* 21:345–383
- Karakose P, Bal C (2024) Efficient chlorophyll prediction and sampling in the sea: A real-time approach with ucb-based path planning. *IEEE Access* 13:8127–8139
- Lampert W (1989) The adaptive significance of diel vertical migration of zooplankton. *Functional Ecology* 3(1):21–27
- Li Y, Zhou Q, Zhang Y, et al (2021) Research trends in the remote sensing of phytoplankton blooms: Results from bibliometrics. *Remote Sensing* 13(21)

- Liu J, Ding K, Yue L, et al (2026) Simultaneous retrieval of chlorophyll-a, total suspended solids, and colored dissolved organic matter in inland waters based on multi-task learning constrained by probability density distributions. *IEEE Transactions on Geoscience and Remote Sensing* 64:1–18
- Low KH (2009) Multi-robot adaptive exploration and mapping for environmental sensing applications. PhD thesis, Carnegie Mellon University, USA, aAI3374752
- LSTS (2022) Inter module communication protocol. URL <https://lsts.pt/docs/imc/master>
- Manjanna S, Li AQ, Smith RN, et al (2018) Heterogeneous multi-robot system for exploration and strategic water sampling. In: 2018 IEEE international conference on robotics and automation (ICRA), IEEE, pp 4873–4880
- McGillicuddy D, Franks PJ (2019) Models of plankton patchiness. *Encyclopedia of Ocean Sciences* 5:536–546
- Mo-Bjørkelund T, Fossum TO, Norgren P, et al (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, <https://doi.org/10.1109/IEEECONF38699.2020.9389324>
- Mo-Bjørkelund T, Majaneva S, Fragoso G, et al (2024) Multi-vehicle adaptive 3d mapping for targeted ocean sampling. *PLOS ONE* 19:e0302514
- Pinto J, Dias PS, Martins R, et al (2013) The lsts toolchain for networked vehicle systems. In: 2013 MTS/IEEE OCEANS - Bergen, pp 1–9, <https://doi.org/10.1109/OCEANS-Bergen.2013.6608148>
- Quigley M (2009) Ros: an open-source robot operating system. In: *ICRA 2009*
- Rigby RA, Stasinopoulos DM (2005) Generalized additive models for location, scale and shape. *Journal of the Royal Statistical Society Series C: Applied Statistics* 54(3):507–554
- Schalles JF (2006) Optical remote sensing techniques to estimate phytoplankton chlorophyll a concentrations in coastal waters with varying suspended matter and cdom concentrations. *Remote Sensing of Aquatic Coastal Ecosystem Processes: Science and Management Applications* pp 27–79
- Sekerci Y, Petrovskii S (2015) Mathematical modelling of plankton–oxygen dynamics under the climate change. *Bulletin of Mathematical Biology* 77:2325–2353
- Stankiewicz P, Tan YT, Kobilarov M (2021) Adaptive sampling with an autonomous underwater vehicle in static marine environments. *Journal of Field Robotics* 38(4):572–597

- Stein ML (2002) The screening effect in kriging. *The Annals of Statistics* 30(1):298–323
- Tippenhauer S, Janout M, Chouksey M, et al (2021) Substantial sub-surface chlorophyll patch sustained by vertical nutrient fluxes in Fram Strait observed with an autonomous underwater vehicle. *Frontiers in Marine Science* 8:605225
- Wirtz K, Smith SL (2020) Vertical migration by bulk phytoplankton sustains biodiversity and nutrient input to the surface ocean. *Scientific Reports* 10(1):1142
- Yang K, Emmerich M, Deutz A, et al (2019) Multi-objective bayesian global optimization using expected hypervolume improvement gradient. *Swarm and evolutionary computation* 44:945–956
- Zhang Y, Kieft B, Hobson BW, et al (2019) 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
- Zhang Y, Ryan JP, Hobson BW, et al (2021) A system of coordinated autonomous robots for lagrangian studies of microbes in the oceanic deep chlorophyll maximum. *Science Robotics* 6(50):eabb9138



RRT*-enhanced long-horizon path planning for AUV adaptive sampling using a cost valley

Yaolin Ge¹*, Jo Eidsvik¹, André Julius Hovd Olaisen

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

ARTICLE INFO

Dataset link: <https://github.com/MASCOT-NTNU/RTCV>

Keywords:

Autonomous underwater vehicles
Adaptive sampling
Oceanographic sampling
Long-horizon path planning

ABSTRACT

Despite advances in adaptive sampling, existing methods predominantly rely on myopic (greedy) strategies and single-objective criteria, which inadequately balance long-term exploration and exploitation. Moreover, obtaining real-time computations with complex, time-varying models remains challenging. With the goal of effective sampling of oceanographic variables by autonomous underwater vehicles, we propose a long-horizon adaptive sampling system that integrates a flexible cost valley concept with a non-myopic path planner. Our method addresses autonomous navigation within a fixed time frame while adaptively sampling ocean variables and avoiding obstacles, aiming to reduce the expected variability or classification error at river plume fronts. The novelty of our approach lies in combining variance and classification metrics as sampling objectives into a weighted cost surface that guides the vehicle along its minimal-cost path. We implement this concept using a rapidly exploring random trees (RRT*) strategy for non-myopic path planning. Simulation results based on 100 replicates demonstrate differences in traffic flow, root mean square error, variance reduction (VR), and integrated Bernoulli variance (IBV) under various cost weightings for RRT* versus a myopic approach. The equal weight cost valley appears robust, yielding prediction metrics closer to those in extreme IBV or VR-dominant cases. Statistical results further show that RRT*-based planning achieves only slightly better numerical scores than the myopic method—for example, an IBV of 75.76 (SD 7.26) compared to 75.93 (SD 6.4). A 2.5-hour field trial in a Norwegian fjord confirms that the AUV successfully runs the long-horizon adaptive sampling algorithm in real time on its onboard computing units.

1. Introduction

With the advent of edge computing and the Internet of Things, there has been a drive towards embedded sensing and computing systems, facilitating active learning for exploration and monitoring, using for instance autonomous underwater vehicles (AUVs) to investigate various oceanographic phenomena. Recent overviews of opportunities for AUV sampling are provided by [1,2]. The 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, adaptive sampling forms AI-based decision systems onboard the AUV for choosing optimal sampling locations. This choice of where to sample make deployments more effective in the under-sampled ocean. Various adaptive sampling methods have gained interest in oceanographic applications such as quantification of plankton [3], chlorophyll [4,5], benthic habitat mapping [6,7] and frontal zones characterization [8–10]. However, these currently used algorithms are greedy (myopic) and/or optimized towards a single objective. Building on earlier contributions, our goal is to develop new long-horizon algorithms for AUV sampling using multiple objectives.

1.1. Problem statement

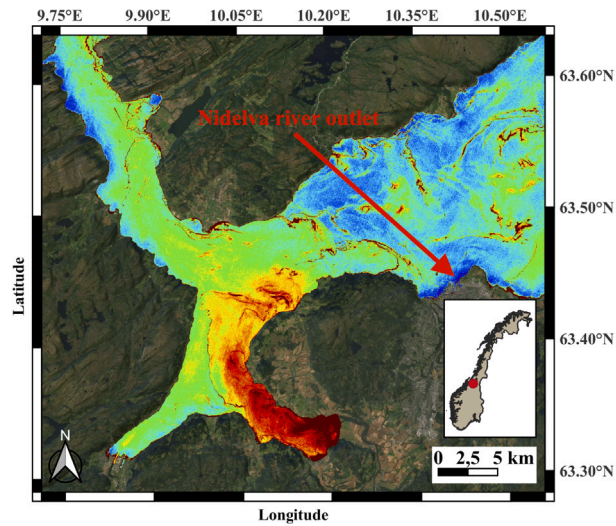
Frontal regions have been recognized as important for physical and biological oceanography [11]. We are particularly interested in river plumes in this paper. There has recently been much research on the transformation of river freshwater discharge and its dissolution into the ambient ocean, see e.g. [12] for a review. This is important for instance in the context of agricultural run-off and understanding how riverine nutrients reach the open ocean [13].

Fig. 1(a) shows a satellite image illustrating typical water mass mixing during Spring in the Trondheim fjord, Norway. Although it is not obvious how to interpret the different colors in an image like this, one recognizes river water masses with different sediment types. 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 (Fig. 1(a)) in this study, with the goal of characterizing the river plume region using ocean models and an AUV. Fig. 1(b) shows the salinity field from numerical ocean

* Corresponding author.


E-mail addresses: geyaolin@gmail.com (Y. Ge), jo.eidsvik@ntnu.no (J. Eidsvik), andre.j.h.olaisen@ntnu.no (A.J.H. Olaisen).



(a)



(b)



(c)

Fig. 1. (a) Sentinel-2 image of ocean water mixing on June-4th, 2021 in the Trondheim fjord. The arrow points out the Nidelva river outlet which we are interested in. (Courtesy of the Copernicus, ESA.) (b) Numerical ocean model data for surface salinity on May 11th, 2022. (c) In the selected operational area, the outer polygon draws the border for the AUV operation whereas the inner polygon shows an unsafe region. The small red dot shows the start location and the big blue dot shows where to end the deployment.

model SINMOD [14] 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 Fig. 1(b) are constructed by averaging numerical ocean model data over time. It shows multiple river outlets causing regional river plumes.

For the AUV field deployment, the operational area is chosen based on the interesting parts of the outputs from the numerical ocean model (Fig. 1(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 in Fig. 1(c)) and an end location (big blue dot in Fig. 1(c)). 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 while sampling the salinity field informatively, and ensure operational constraints such as avoiding the shallow region near the island in the center and that of time or traveling distance budgets.

The start location and end destination are chosen to enable reasonable mapping of the entire river plume frontal region, well within the




Fig. 2. (Left) The adaptive agent follows the sense-plan-act framework. (Right) An illustration of myopic (near-sighted) and non-myopic (long-horizon) path planning strategies, shows the challenges posed by local optima and the pursuit of the global optimum.

specified distance budget. A critical challenge lies in developing robust path-planning strategies that balance exploration and exploitation objectives to achieve effective AUV deployment for salinity sampling.

1.2. Literature review

In using only one agent (AUV in our case) to conduct adaptive sampling, dominating methods can be grouped into either myopic (greedy) or non-myopic approaches. Myopic strategies guide the agent towards the most informative location selected from a subset of candidate locations within the myopic (near-sighted) neighborhood radius or short horizon. Spatial sampling must balance exploration and exploitation [15], and greedy algorithms can 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 [16]. We will focus on non-myopic approaches in this presentation. [17] demonstrated 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. [18] presented 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. There are marine field applications using RRT*. [19] demonstrated autonomous navigation for surface vessels and [20] used RRT* onboard an AUV but only for standard occupancy maps. It has shown difficulty 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 exploration 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 agent to recognize high or low costs and rewards. 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 [21]. But in many cases, the reward is more ambiguous, say via occupancy grids for presence or absence of a feature [22], reduced variance or entropy of a field variable [23–26] or by identifying hot spots in the field or the different water masses [9,10,27]. 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 incorporate both information rewards and operational costs. Others have looked at multi-objective cost functions for path planning, albeit not for AUV sampling. For instance, [28] applied a

cost map approach to optimize path planning in different case studies within a predefined configuration space. Similarly, [29] showed the feasibility of utilizing a blended cost map based on the weighted sum of sub-components for rough terrain path planning using RRT. For an AUV setting, a simple multi-objective is a balance between hotspots (large prediction) and large variance. [30] used this model and a branch-and-bound approach for non-myopic sampling of dissolved oxygen. Maybe most similar to our suggested approach, [31] showed the benefits of multiple objectives in the setting of evolving mission policies where an AUV goes back and forth between two locations while avoiding obstacles and staying in a thermal layer. Recent advances have further emphasized the need for multi-objective planning. For example, [32] introduced a multi-strategy adaptable ant colony optimization algorithm that dynamically adjusts to multiple planning criteria, and [33] proposed enhancements in heuristic mechanisms for handling dynamic environments. These studies underscore the growing trend toward integrating adaptive weighting and long-horizon considerations, motivating our introduction of a flexible cost valley that seamlessly combines operational safety with informative sampling objectives. We go beyond this in the way we update a statistical model onboard and use RRT* for planning.

1.3. Novelty and outline

Fig. 2 (left) illustrates the common sense-plan-act framework where an agent over stages senses the environment, plans where to go next and then commits to that plan (acts) for the current stage. Using non-myopic (long-horizon) rather than myopic (greedy) planning (right display of Fig. 2), the agent is less prone to search for local optima but rather aims for the long-term goal of sampling. In this study, we design a non-myopic adaptive sampling system for agents mapping boundaries or fronts. Here, one aims to characterize the front accurately while also exploring the domain. In this sense, we define our objective function as a balance between reducing (i) the classification error or frontal set uncertainty and (ii) the integrated variance. Building on these two along with collision avoidance and time constraints, we form the cost valley. This is used in the onboard path planner.

The main contributions of this paper are:

- A versatile multi-criterion cost surface where its cost valley defines promising agent sampling paths.
- A RRT* algorithm for exploring the cost valley in a long-horizon planner.
- An AUV field deployment sampling a river frontal zone based on RRT* and the cost valley concept.

In Section 2, we outline our statistical model running onboard the agent. In Section 3, we present the cost valley concept. In Section 4, we study the properties of the suggested approach using simulations. In Section 5, we show the results of the AUV field deployment. In Section 6, we conclude and point to future research directions.

2. 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.

2.1. 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 [34].

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 0.5 m. Times are indicated by $t = 0, 1, \dots$, discretized in a regular sampling time interval. We denote the spatio-temporal salinity variable by $\xi_t = (\xi_{t,u_1}, \dots, \xi_{t,u_n})^T$.

The initial state is Gaussian distributed $\xi_0 \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, 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; $\Sigma(i, i') = \sigma^2(1 + \phi h(i, i')) \exp(-\phi h(i, i'))$, 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 [34] of SINMOD data.

The temporal variation is defined by an autoregressive process:

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

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 $\xi_t \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ 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.

2.2. Updating

The prior model described in Eq. (1) is updated by in-situ AUV measurements, where the AUV is assumed to cruise at a constant depth of 0.5 m. We model the AUV measurement y_t at stage or time $t = 1, \dots$ by

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

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

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

updated mean $\mathbf{m}_{t|t}$ and covariance $\mathbf{S}_{t|t}$ are available by the recursive Kalman filter formulae defined by

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

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

If the AUV is pausing, the last three steps in Eq. (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 Eq. (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.

3. Long-horizon path planning using a 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 as in Fig. 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 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.

3.1. 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. Fig. 3 illustrates the obstacle cost (left) and the time budget cost (right).

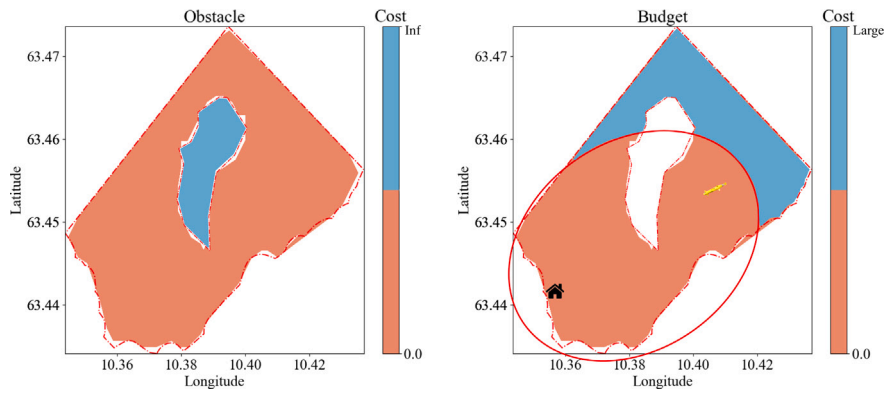


Fig. 3. Operational cost fields. The left display shows an obstacle cost field that assigns infinite cost to the island in the middle of the operation domain and zero cost to other areas. The right display shows a budget cost field that assigns high costs to areas outside the budget ellipse and zero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

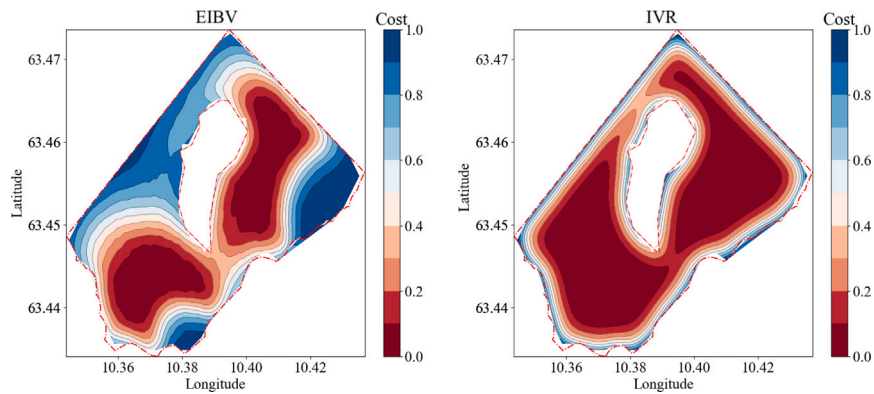


Fig. 4. Informative cost fields at the initial stage. The left display shows the EIBV cost field, which assigns low cost to areas where the river plume front might exist at this initial time. The right display shows the IVR cost field, which assigns low cost almost everywhere as it has not yet sampled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

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.

3.2. Informative cost fields

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

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 Eq. (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 $p(\mathbf{u}_i) = P(\xi_{t,\mathbf{u}_i} < \zeta) = \Phi_1\left(\frac{\zeta - m_{t|t}(\mathbf{u}_i)}{\sqrt{S_{t|t}(\mathbf{u}_i, \mathbf{u}_i)}}\right)$, 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 Φ_2 [10]. 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 to the threshold have a large Bernoulli variance, and therefore carry much information, as shown in Fig. 4.

Algorithm 1 shows the informative cost fields computation steps. 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 sampling individual locations. Spatial site \mathbf{u}_i has an associated cost value for each criterion. The closed form solutions for both IVR and EIBV provided in Algorithm 1 ensure that the computations are relatively fast to do on the onboard computing units.

3.3. Cost valley construction

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

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

The operational cost fields, which include the obstacle cost field $\mathbf{C}_{\text{obstacle}}$ and the budget cost field $\mathbf{C}_{\text{budget}}$, are important for 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, the agent must return to the base when the budget cost field

Algorithm 1 Calculate informative cost fields for EIBV and IVR

Require: $m_{t|t-1}, \mathbf{S}_{t|t-1}$
EIBV = $\mathbf{0}^{n \times 1}$ \triangleright Initialize the expected integrated Bernoulli variance
IVR = $\mathbf{0}^{n \times 1}$ \triangleright Initialize the integrated variance reduction

for $i \in 1 \dots n$ **do**
 $\mathbf{f} = \mathbf{0}^{n \times 1}, \quad \mathbf{f}(i) = 1$ \triangleright Design point
 $\mathbf{R} = \mathbf{S}_{t|t-1} \mathbf{f} (\mathbf{f}^T \mathbf{S}_{t|t-1} \mathbf{f} + r^2)^{-1} \mathbf{f}^T \mathbf{S}_{t|t-1}$ \triangleright Covariance reduction matrix
IVR(i) = $\sum_{i'=1}^n \text{diag}(\mathbf{R})$ \triangleright Integrated variance reduction
 where, $\mathbf{a} = \begin{bmatrix} \zeta \\ -\zeta \end{bmatrix}$ \triangleright Threshold ζ evaluation vector
 $\mathbf{b}(i') = \begin{bmatrix} m_{t|t-1}(i') \\ -m_{t|t-1}(i') \end{bmatrix}$ \triangleright Threshold calculation mean
 $\mathbf{W}(i', i') = \begin{bmatrix} S_{t|t-1}(i', i') & -R(i', i') \\ -R(i', i') & S_{t|t-1}(i', i') \end{bmatrix}$ \triangleright Threshold calculation covariance
EIBV(i) = $\sum_{i'=1}^n \Phi_2(\mathbf{a}, \mathbf{b}(i'), \mathbf{W}(i', i'))$, \triangleright Expected integrated Bernoulli variance
end for
 $\mathbf{C}_{\text{EIBV}} = (\mathbf{EIBV} - \min(\mathbf{EIBV})) / (\max(\mathbf{EIBV}) - \min(\mathbf{EIBV}))$ \triangleright Normalized EIBV cost
 $\mathbf{C}_{\text{IVR}} = 1 - (\mathbf{IVR} - \min(\mathbf{IVR})) / (\max(\mathbf{IVR}) - \min(\mathbf{IVR}))$ \triangleright Normalized IVR cost
return $\mathbf{C}_{\text{EIBV}}, \mathbf{C}_{\text{IVR}}$

is nearly exhausted. Therefore, locations located outside the allocated budget ellipse are heavily penalized, regardless of their perceived worth in the informative cost fields.

Fig. 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 Eq. (4). The weighting aids the AUV sampling in balancing exploration of uncertain locations and detecting the salinity boundaries, while maintaining the operational constraints for the vehicle.

3.4. 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 e.g. [35]. This planner computes the least-cost trajectory 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 Eq. (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 a RRT* path planner to determine the next waypoint, see [36].

In Fig. 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 facilitates long-horizon planning, and the planner is able to navigate away from the high-cost area. However, due to insufficient convergence in the RRT* path planner, the path taken is sub-optimal, as evidenced by the detour in the trajectory. Increasing the number of iterations in the tree expansion can improve the optimality of the path. However, one must consider the trade-off between the computational cost and path optimality.

Algorithm 2 Informative long-horizon path planning algorithm

```

Set  $m_{1|0} = \mu$  and  $S_{1|0} = \Sigma$ .           ▷ Initialize mean and covariance
Set start waypoint  $D_0$                        ▷ Initialize design set
 $\mathcal{Y}_0 = \emptyset$                                ▷ Initialize data vector
Budget=MaxDistanceRange
 $t = 1$ .
while Budget >= allowance do
  Plan :
    CV = updateCostValley( $m_{t|t-1}, S_{t|t-1}, \text{Budget}, D_{t-1}$ ) ▷ Cost valley
    calculation (Algorithm 1)
     $u_t = \text{argmin}_{u \in \mathcal{M}}(\text{CV})$ 
     $T_t = \text{RRT}^*(\text{CV}, D_{t-1}, u_t)$  ▷ Run rapidly exploring random trees
     $D_t = T_t\{1\}$ 
    Budget = Budget -  $\|D_t - D_{t-1}\|_2$ 
  Act :
    Go to waypoint  $D_t$ .                       ▷ Commit to chosen design
  Sense :
     $\mathcal{Y}_t = (\mathcal{Y}_{t-1}, y_t)$ .                     ▷ Gather data  $y_t$  and augment data set
     $G_t = S_{t|t-1} f_t (f_t^T S_{t|t-1} f_t + r^2)^{-1}$ .
     $m_{t|t} = m_{t|t-1} + G_t (y_t - f_t^T m_{t|t-1})$ .   ▷ Updated mean vector
    (Eq. (3))
     $S_{t|t} = S_{t|t-1} - G_t f_t^T S_{t|t-1}$ . ▷ Updated covariance matrix (Eq. (3))
     $m_{t+1|t} = \mu + \rho(m_{t|t} - \mu)$ .           ▷ Predictive mean vector (Eq. (3))
     $S_{t+1|t} = \rho^2 S_{t|t} + (1 - \rho^2)\Sigma$ .   ▷ Predictive covariance matrix
    (Eq. (3))
     $t = t + 1$ .
end while

```

3.5. System architecture

Fig. 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 an updated state model including an uncertainty map, which is essential for the cost valley calculation. After this model update, an optimal trajectory is determined from the current location to the minimum cost location, producing the next waypoint. The ROS-IMC bridge is then used to send instructions to the AUV, which is discussed in more detail in the field deployment description. DUNE executes the control command, allowing the AUV to sample autonomously.

4. 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.

4.1. Simulation setup

Fig. 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 Fig. 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 h from SINMOD as the mean value field. For the variation in the replicate runs, the Matérn covariance matrix coefficients are specified as $\sigma = 1.0$ and $\phi = 0.0064$. The average salinity over the 100 replicates is shown in Fig. 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. 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 240 m and are allowed to run for 8 h, 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 $r = 0.5$.

4.2. 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.

4.2.1. Myopic

Fig. 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. Fig. 8(a) 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.

Figs. 8(b) and 8(c) 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 (Fig. 8(d)), the eastern parts remain 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.

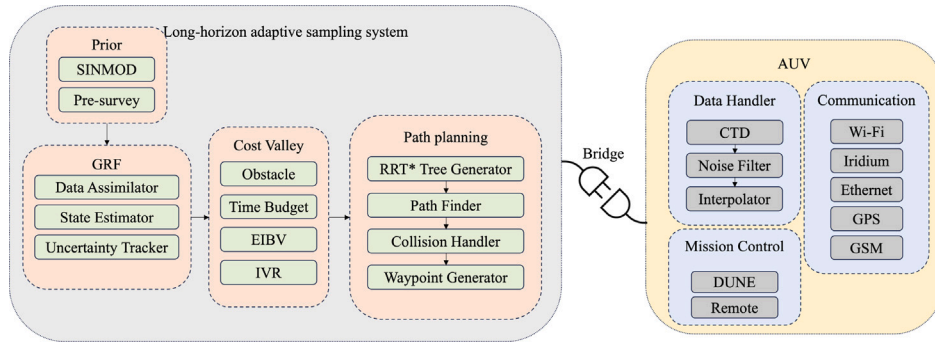


Fig. 6. System architecture of our proposed adaptive sampling system using RRT* for path planning and a cost valley for guidance.

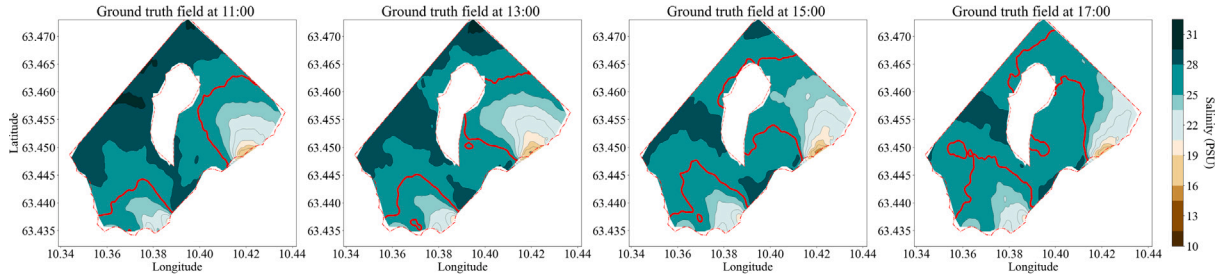


Fig. 7. The surface salinity ground truth field at different time steps. The red contour shows the boundary of different water masses classified by the threshold.

4.2.2. Long-horizon RRT*

Fig. 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.

Fig. 9(a) 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 (Figs. 9(b) and 9(c)), the AUV now moves between the west and east parts of the plume front. Hence, in contrast to the myopic approach in Fig. 8, the adaptive agent utilizing the long-horizon RRT* strategy exhibits a sampling behavior that covers much more of the spatial domain. Running RRT* onboard, the agent's ability to escape local attractions is 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 Fig. 9(d), 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.

4.3. 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.

The comparison of traffic-flow density maps in Figs. 10 and 11 reveal distinct behaviors of the myopic and long-horizon RRT* planners across different cost valleys. Under the EIBV dominant cost valley, the myopic planner exhibits concentrated sampling in southern regions, whereas the RRT* planner achieves a more balanced distribution across both southern and northern areas. This suggests that the RRT* planner

optimizes sampling efficiency by covering a broader area, minimizing the risk of over-exploitation.

For the IVR dominant cost valley, the myopic planner shows localized behavior, focusing on high-density regions, while the RRT* planner more effectively spreads its efforts throughout the field. The latter's balanced coverage reduces uncertainty across the domain, supporting the RRT* planner's superiority in handling large, dynamic fields. In the equal weight cost valley, the myopic planner's sampling remains localized around the initial position, with limited reach into northern areas. In contrast, the RRT* planner demonstrates a more adaptive approach, adjusting its coverage to include both southern and northern regions as environmental conditions, such as tides, shift. This adaptability highlights the RRT* planner's capability for balancing exploration and exploitation, ensuring comprehensive field coverage and enhanced sampling efficiency.

4.4. Metric analysis

In the replicate simulations, we monitor key metrics: integrated Bernoulli variance (IBV), variance reduction (VR), and root mean squared error (RMSE). For each metric, we take the average over all spatial grid cells, at every time point. Table 1 shows the last step, i.e., after 120 steps. Overall, the long-horizon RRT* planner does not consistently outperform the myopic planner across all cost valleys. For the EIBV dominant case, RRT* shows a marginal improvement in IBV, suggesting a slight advantage in exploring high-value areas (i.e. the boundary region). However, in both the IVR dominant and equal-weighted cases, both planners achieve comparable results in all metrics, reflecting a similar capacity to reduce variance across the field. These results indicate that while the long-horizon approach offers a more balanced exploration strategy, the performance gains may be minor, especially in highly structured fields like the IVR dominant case.

5. Experiment in the Trondheim fjord

We present the case of river plume exploration in the Trondheim fjord, Norway (Fig. 1). The suggested algorithm using RRT* with a cost valley runs onboard an AUV in this field experiment.

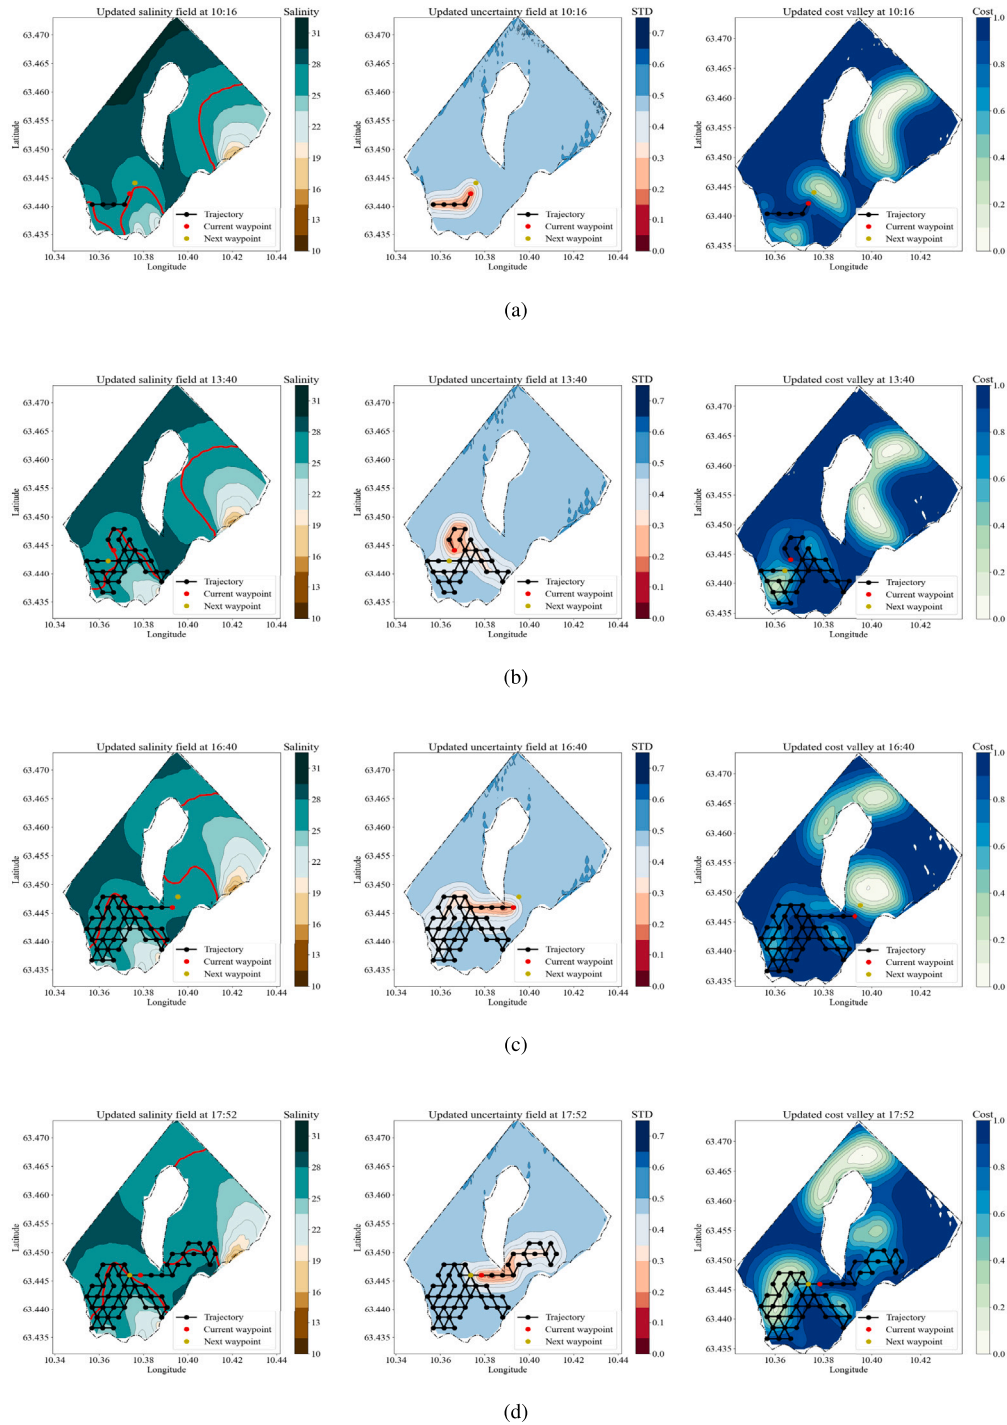


Fig. 8. The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using the myopic strategy over a series of time iterations (a–d). In this particular example, most of the effort was focused on the southwest side before the agent moved eastward.

Table 1
Comparison of Myopic and RRT* performance under different cost valleys. The metrics include mean and standard deviation for IBV, IVR, and RMSE.

Cost valley	IBV		VR		RMSE	
	Myopic mean std	RRT* mean std	Myopic mean std	RRT* mean std	Myopic mean std	RRT* mean std
EIBV dominant	70.82 5.85	71.76 5.69	241.12 3.43	241.93 7.66	1.04 0.11	1.05 0.11
IVR dominant	75.93 6.40	75.76 7.26	229.83 0.00	232.93 3.84	0.96 0.09	0.98 0.10
Equal weight	71.54 5.70	72.45 6.88	234.27 3.10	236.62 7.06	1.01 0.11	1.02 0.11

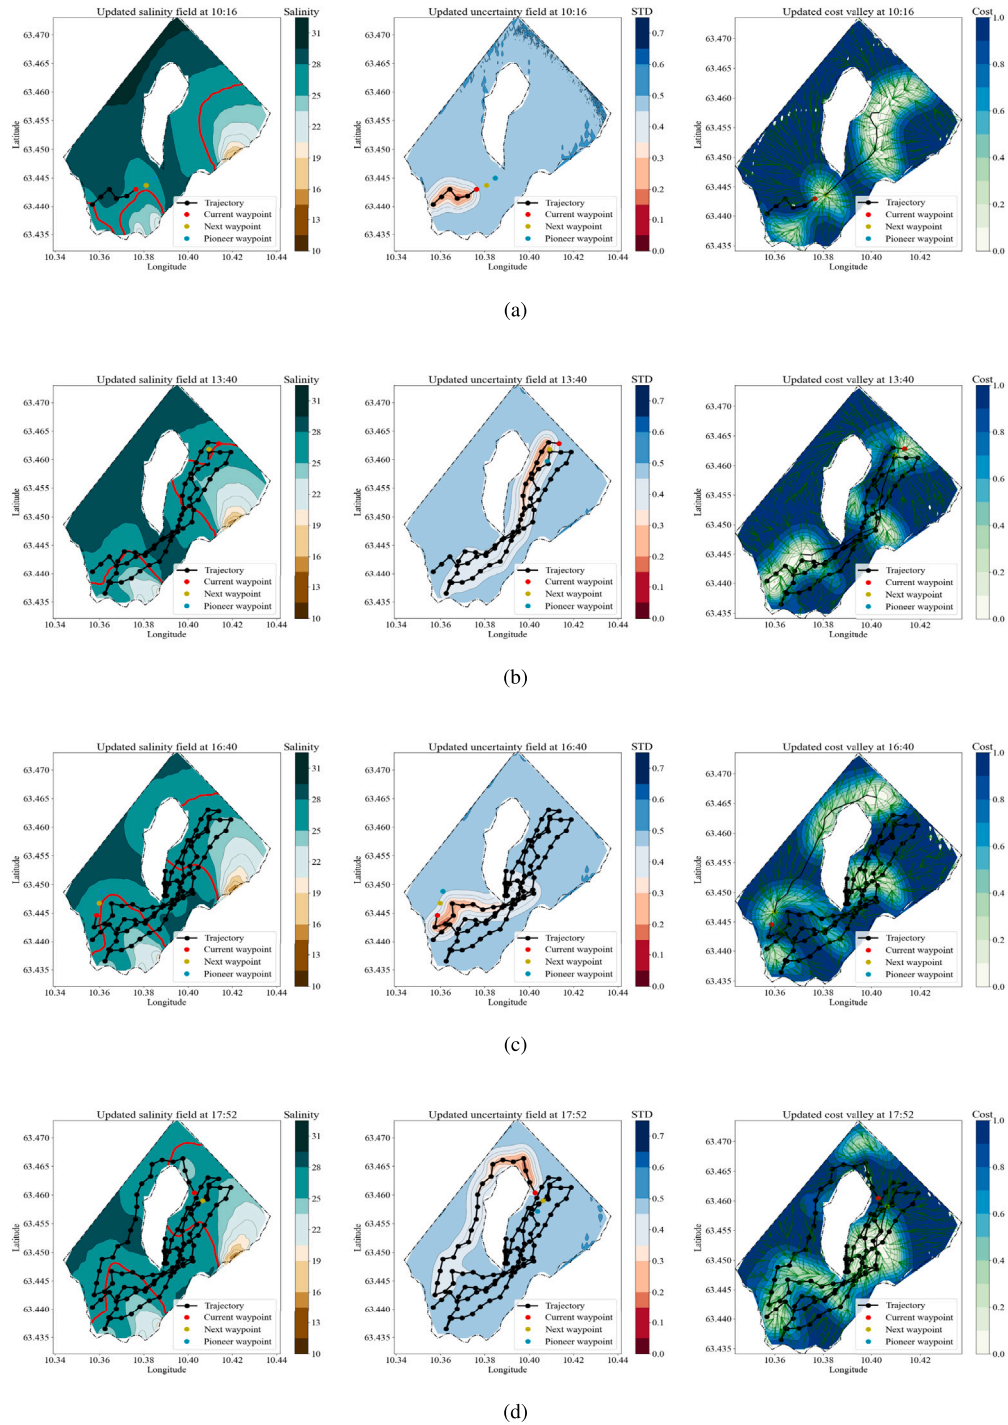


Fig. 9. The updated mean salinity field (left), uncertainty field (middle) and cost valley (right) after the adaptive sampling using our proposed long-horizon RRT* strategy over a series of time iterations (a–d). The RRT* trees are displayed in the cost valley column as well to better illustrate the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

5.1. Experimental setup

The map view in Fig. 1(b) 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.

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 Fig. 12(a).

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 [41], which manages the messaging between ROS and DUNE. Our algorithm interfaces directly with Robot Operating Systems (ROS) [40], and its messages are then relayed to the ROS-IMC bridge in the vehicle, which incorporates DUNE (DUNE: Unified Navigation Environment [37]). The components within the LAUV communicate using the Inter Module Communication (IMC) message protocol [38]. Fig. 12(b) summarizes the connections between software elements. We implement the same integration scheme as in [10], 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

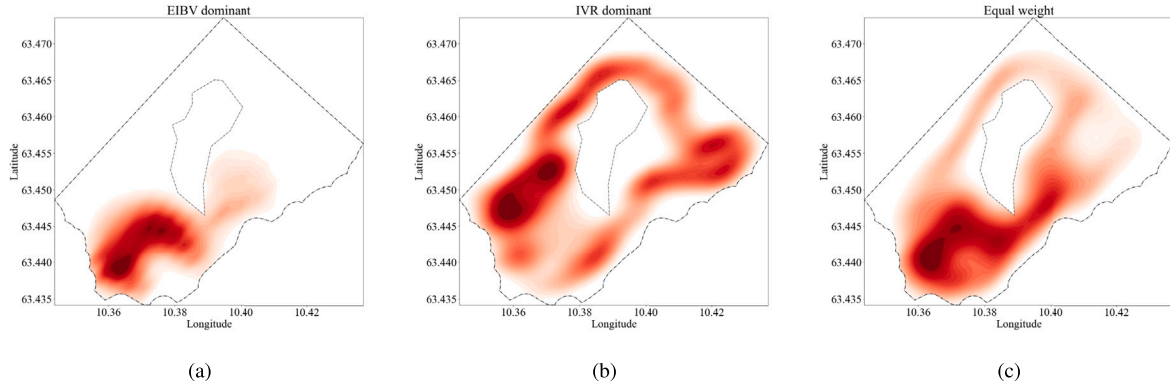


Fig. 10. Traffic flow density map for the myopic agent under (a) EIBV dominant, (b) IVR dominant, (c) equal weight cost valleys.

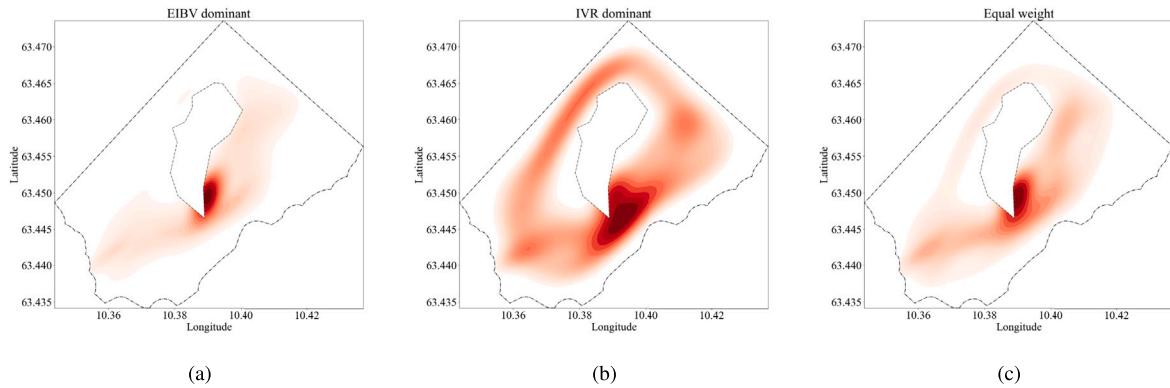




Fig. 11. Traffic flow density map for the long-horizon RRT* agent under (a) EIBV dominant, (b) IVR dominant, (c) equal weight cost valleys.



(a)



(b)

Fig. 12. (a) The LAUV named Thor is under remote operation check before its expedition. (b) Main software components in the communication between the adaptive code and the vehicle. DUNE [37] is running on the main CPU of the AUV while the IMC [38] messages are transmitted via TCP [39] to an auxiliary CPU, where ROS [40] and the adaptive code is run.

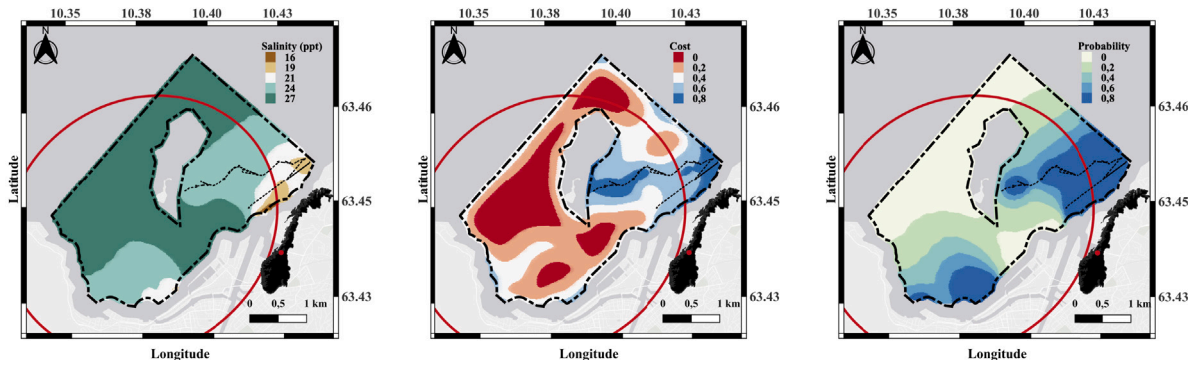


Fig. 13. The updated mean field for the salinity (left), the cost valley (middle), and the excursion probability (right) after sampling the region from 12:30 to 13:41. The remaining budget starts to take an effect, shown as the red ellipse.




Fig. 14. 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).

resurface every 10 min to correct its navigational errors. The operation took place on May 11th, 2022, and it lasted for more than 2.5 h. 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.

5.2. Results

Fig. 13 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 Fig. 13 (middle-right). The path taken by the LAUV (black line) indicates that it has used the first part of its adaptive mission to explore the area close to the river mouth. As the AUV travels through the region, it adheres to the guidance provided by the cost valley and endeavors to minimize expenses. The budget ellipse (red color in Fig. 13) diminishes as time goes by.

At step 44 (15:10), Fig. 14 shows that the LAUV has avoided the obstacle and discovered surprisingly low salinity values in the western parts. The excursion probability field (right display) 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.

Fig. 15(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. Fig. 15(b) presents the cross plot of these two data sources. The kernel

density estimate shows that SINMOD data has two modes near salinity 23 and 27, while the AUV data is distributed around salinity 25. The majority of the SINMOD data is higher than the actual measurements, indicating the need for practical adjustments of SINMOD data to avoid bias stemming from the numerical solver.

6. 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 the start location to the home destination while adaptively sampling, in real-time, 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 objectives, the path planning achieves a balance between exploitation and exploration while the hard constraints on safety and punctuality are all considered, as demonstrated in the results of the field deployment.

In terms of the algorithm properties, 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.

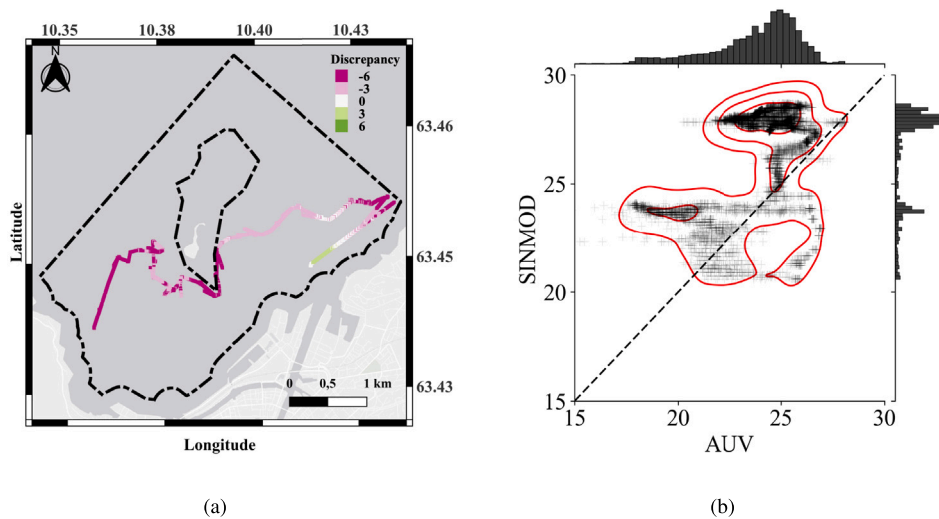


Fig. 15. (a) The difference between the AUV measurements and the SINMOD prediction, referred to as the residual. The in-situ data reveal a considerable river plume in the western part, characterized by low salinity values, which contradicts the SINMOD prediction. (b) The cross plot between the SINMOD 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.

We believe that non-myopic sampling for multiple objectives will see more use in the future, for ocean exploration and other applications. 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 grows, there is an enhanced capability to address multiple areas of interest, contingent upon each component receiving ample allocation for exploration. A potential improvement over what we suggested in this paper is the implementation of dynamic weighing mechanisms among the cost valley components, say based on time of day.

A plethora of long-horizon path planning algorithms exist; other variants of RRT*, A*, the probabilistic road map and ant-colony optimization [32,33] 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 contexts with computational or hardware constraints. We managed to run real-time adaptive sampling with a long-horizon planner on a two-dimensional time-varying model. Adding depth variation would make it more computationally demanding, and so would a more realistic temporal dynamical model (going beyond the auto-regressive model), and it is yet unclear which algorithm would enable efficient and timely sampling in such situations.

We only considered static obstacles, such as the island in the field deployment. Future work will focus on incorporating dynamic obstacle avoidance through real-time AIS data and implementing adaptive weighting schemes for further optimization of the exploration-exploitation trade-off. Additionally, integrating predictive machine learning techniques and modern reinforcement learning-based methods will be explored to enhance scalability and operational efficiency.

CRedit authorship contribution statement

Yaolin Ge: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jo Eidsvik:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Funding acquisition, Conceptualization. **André Julius Hovd Olaisen:** Writing – original draft, Methodology, Investigation, Conceptualization.

Declaration

During the preparation of this work, the authors used ChatGPT to improve the English language. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We acknowledge support from the The Research Council of Norway (RCN) through the MASCOT project 305445. The authors thank AU-RLab 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. Data and code is available in the repository: <https://github.com/MASCOT-NTNU/RRTCVC>

Data availability

Data and code is available in the repository: <https://github.com/MASCOT-NTNU/RRTCVC>.

References

- [1] J. Hwang, N. Bose, S. Fan, AUV adaptive sampling methods: A review, *Appl. Sci.* 9 (2019) 3145.
- [2] N. Basil, M.E. Alqaysi, M. Deveci, A.S. Albahri, O.S. Albahri, A.H. Alamooodi, Evaluation of autonomous underwater vehicle motion trajectory optimization algorithms, *Knowl.-Based Syst.* 276 (2023) 110722.
- [3] J. Das, F. Py, J.B.J. Harvey, J.P. Ryan, A. Gellene, R. Graham, D.A. Caron, K. Rajan, G.S. Sukhatme, Data-driven robotic sampling for marine ecosystem monitoring, *Int. J. Robot. Res.* 34 (2015) 1435–1452.
- [4] T.O. Fossum, G.M. Fragoso, E.J. Davies, J.E. Ullgren, R. Mendes, G. Johnsen, I. Ellingsen, J. Eidsvik, M. Ludvigsen, K. Rajan, Toward adaptive robotic sampling of phytoplankton in the coastal ocean, *Sci. Robot.* 4 (2019) eaav3041, <http://dx.doi.org/10.1126/scirobotics.aav3041>, arXiv:<https://www.science.org/doi/pdf/10.1126/scirobotics.aav3041>.

- [5] Y. Zhang, 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, C.A. Scholin, Autonomous tracking and sampling of the deep chlorophyll maximum layer in an open-ocean eddy by a long-range autonomous underwater vehicle, *IEEE J. Ocean. Eng.* 45 (2020) 1308–1321.
- [6] P. Rigby, O. Pizarro, S.B. Williams, Toward adaptive benthic habitat mapping using gaussian process classification, *J. Field Robot.* 27 (2010) 741–758.
- [7] S. Anyosa, J. Eidsvik, O. Pizarro, Adaptive spatial designs minimizing the integrated Bernoulli variance in spatial logistic regression models-with an application to benthic habitat mapping, *Comput. Statist. Data Anal.* 179 (2023) 107643.
- [8] Y. Zhang, J.G. Bellingham, J.P. Ryan, B. Kieft, M.J. Stanway, Autonomous four-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle, *J. Field Robot.* 33 (2016) 67–81.
- [9] T.O. Fossum, C. Travelletti, J. Eidsvik, D. Ginsbourger, K. Rajan, Learning excursion sets of vector-valued Gaussian random fields for autonomous ocean sampling, *Ann. Appl. Stat.* 15 (2021) 597–618.
- [10] Y. Ge, J. Eidsvik, T. Mo-Bjørkelund, 3D adaptive AUV sampling for classification of water masses, *IEEE J. Ocean. Eng.* 48 (2023) 626–639.
- [11] I.M. Belkin, P.C. Cornillon, K. Sherman, Fronts in large marine ecosystems, *Prog. Oceanogr.* 81 (2009) 223–236.
- [12] A.R. Horner-Devine, R.D. Hetland, D.G. MacDonald, Mixing and transport in coastal river plumes, *Annu. Rev. Fluid Mech.* 47 (2015) 569–594.
- [13] J. Sharples, J.J. Middelburg, K. Fennel, T.D. Jickells, What proportion of riverine nutrients reaches the open ocean? *Glob. Biogeochem. Cycles* 31 (2017) 39–58.
- [14] D. Slagstad, T.A. McClimans, Modeling the ecosystem dynamics of the barents sea including the marginal ice zone: I. Physical and chemical oceanography, *J. Mar. Syst.* 58 (2005) 1–18.
- [15] H. Yang, Y. Yu, J. Cheng, Z. Lei, Z. Cai, Z. Zhang, S. Gao, An intelligent metaphor-free spatial information sampling algorithm for balancing exploitation and exploration, *Knowl.-Based Syst.* 250 (2022) 109081.
- [16] S. Bai, T. Shan, F. Chen, L. Liu, B. Englot, Information-driven path planning, *Curr. Robot. Rep.* 2 (2021) 177–188.
- [17] C. Xiao, J. Wachs, Nonmyopic informative path planning based on global Kriging variance minimization, *IEEE Robot. Autom. Lett.* 7 (2022) 1768–1775.
- [18] J. Suh, J. Gong, S. Oh, Fast sampling-based cost-aware path planning with nonmyopic extensions using cross entropy, *IEEE Trans. Robot.* 33 (2017) 1313–1326.
- [19] T.T. Enevoldsen, R. Galeazzi, Grounding-aware RRT* for path planning and safe navigation of marine crafts in confined waters, *IFAC-PapersOnLine* 54 (2021) 195–201.
- [20] L. Zucchini, M. Franchi, A. Ridolfi, Sensor-driven autonomous underwater inspections: A receding-horizon RRT-based view planning solution for AUVs, *J. Field Robot.* 39 (2022) 499–527.
- [21] J. Eidsvik, T. Mukerji, D. Bhattacharjya, Value of Information in the Earth Sciences: Integrating Spatial Modeling and Decision Analysis, Cambridge University Press, 2015.
- [22] M. Jakuba, D.R. Yoerger, Autonomous search for hydrothermal vent fields with occupancy grid maps, in: *Proc. of ACRA*, vol. 8, Citeseer, 2008, p. 2008.
- [23] J. Binney, A. Krause, G.S. Sukhatme, Optimizing waypoints for monitoring spatiotemporal phenomena, *Int. J. Robot. Res.* 32 (2013) 873–888.
- [24] G.E. Berget, T.O. Fossum, T.A. Johansen, J. Eidsvik, K. Rajan, Adaptive sampling of ocean processes using an AUV with a Gaussian proxy model, *IFAC-PapersOnLine* 51 (2018) 238–243, <http://dx.doi.org/10.1016/j.ifacol.2018.09.509>, 11th IFAC Conference on Control Applications in Marine Systems, Robotics, and Vehicles CAMS 2018.
- [25] K.-C. Ma, L. Liu, H.K. Heidarsson, G.S. Sukhatme, Data-driven learning and planning for environmental sampling, *J. Field Robot.* 35 (2018) 643–661.
- [26] T.O. Fossum, J. Eidsvik, I. Ellingsen, M.O. Alver, G.M. Fragoso, G. Johnsen, R. Mendes, M. Ludvigsen, K. Rajan, Information-driven robotic sampling in the coastal ocean, *J. Field Robot.* 35 (2018) 1101–1121, <http://dx.doi.org/10.1002/rob.21805>, [arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21805](https://onlinelibrary.wiley.com/doi/pdf/10.1002/rob.21805).
- [27] J. Das, K. Rajany, S. Frolovy, F. Pyy, J. Ryany, D.A. Caronz, G.S. Sukhatme, Towards marine bloom trajectory prediction for AUV mission planning, in: 2010 IEEE International Conference on Robotics and Automation, IEEE, 2010, pp. 4784–4790.
- [28] L. Jaillet, J. Cortés, T. Siméon, Sampling-based path planning on configuration-space costmaps, *IEEE Trans. Robot.* 26 (2010) 635–646.
- [29] A. Ettlin, H. Bleuler, Rough-terrain robot motion planning based on obstacleness, in: 2006 9th International Conference on Control, Automation, Robotics and Vision, 2006, pp. 1–6, <http://dx.doi.org/10.1109/ICARCV.2006.345116>.
- [30] P. Stankiewicz, Y.T. Tan, M. Kobilarov, Adaptive sampling with an autonomous underwater vehicle in static marine environments, *J. Field Robot.* 38 (2021) 572–597.
- [31] G. Okopal, Multi-objective autonomy for auv adaptive sampling missions, in: *Oceans 2019 MTS/IEEE Seattle*, IEEE, 2019, pp. 1–6.
- [32] J. Cui, L. Wu, X. Huang, D. Xu, C. Liu, W. Xiao, Multi-strategy adaptable ant colony optimization algorithm and its application in robot path planning, *Knowl.-Based Syst.* 288 (2024) 111459.
- [33] C. Liu, L. Wu, W. Xiao, G. Li, D. Xu, J. Guo, W. Li, An improved heuristic mechanism ant colony optimization algorithm for solving path planning, *Knowl.-Based Syst.* 271 (2023) 110540.
- [34] N. Cressie, C.K. Wikle, *Statistics for Spatio-Temporal Data*, John Wiley & Sons, 2015.
- [35] G.A. Hollinger, G.S. Sukhatme, Sampling-based robotic information gathering algorithms, *Int. J. Robot. Res.* 33 (2014) 1271–1287.
- [36] Y. Ge, A.J.H. Olaisen, J. Eidsvik, R.P. Jain, T.A. Johansen, Long-horizon informative path planning with obstacles and time constraints, *IFAC-PapersOnLine* 55 (2022) 124–129.
- [37] J. Pinto, P.S. Dias, R. Martins, J. Fortuna, E. Marques, J. Sousa, The LSTS toolchain for networked vehicle systems, in: 2013 MTS/IEEE Oceans - Bergen, 2013, pp. 1–9, <http://dx.doi.org/10.1109/OCEANS-Bergen.2013.6608148>.
- [38] LSTS, Inter module communication protocol, 2022, URL: <https://lsts.pt/docs/imc/master>.
- [39] V. Cerf, R. Kahn, A protocol for packet network intercommunication, *IEEE Trans. Commun.* 22 (1974) 637–648.
- [40] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A.Y. Ng, ROS: An open-source robot operating system, in: *ICRA Workshop on Open Source Software*, vol. 3, 2009, p. 5.
- [41] T. Mo-Bjørkelund, T.O. Fossum, P. Norgren, M. Ludvigsen, Hexagonal grid graph as a basis for adaptive sampling of ocean gradients using AUVs, in: *Global Oceans 2020: Singapore – U.S. Gulf Coast*, 2020, pp. 1–5, <http://dx.doi.org/10.1109/IEEECONF38699.2020.9389324>.

Spatio-temporal Calanus Mapping Using Different Vehicles and Sensors in a Campaign off the Norwegian Coast

André Julius Hovd Olaisen¹, Jo Eidsvik¹, Ahmed Abdelgayed²,
Erik Liu², Karoline Barstein², and Martin Ludvigsen²

¹Norwegian University of Science and Technology, Department of
Mathematical Sciences

²Norwegian University of Science and Technology, Department of
Marine Technology

Abstract

This paper presents data and analysis results from a field campaign conducted in Spring off the coast of Norway. A goal of this campaign was to use multiple ocean vehicles to study the phytoplankton and zooplankton densities in space and time. Platforms included two AUVs and two ASVs, in addition to sampling done from the operational boat. The AUVs use Chl *a* as a proxy for phytoplankton and a particle camera for zooplankton. The ASVs estimate vertical zooplankton distribution using echosounders. Lastly, there are optical sensors winched from the operational boat as well as nets for physical samples.

We build spatio-temporal statistical models for the upper and lower boundaries of the plankton layers, as well as for the integrated biomass over depth. Separate models are fitted for each ocean variable and for every vehicle. Results are then compared across platforms and variables. Although the different platforms are never in exactly the same location at the same time, or for the AUVs at the same depth, we can still compare data from periods when they are in close proximity. We learn that there are evident correlations when observing phytoplankton or zooplankton between vehicles, in particular of the same type. Especially notable is the correlation between the optical zooplankton sensors on the AUVs and the acoustic sensors on the ASVs. We notice a negative correlation, if any, between the phytoplankton and zooplankton results across all platforms and sensor types.

1 Introduction

In spring, increasing light and seasonal changes in the water column promote a phytoplankton bloom in the Coastal Norwegian Sea (Rey, 2004). This is again followed by a bloom of grazing zooplankton that has been hibernating deep in the sea (Broms and Melle, 2007). The dominant species of zooplankton in the Coastal Norwegian Sea is *Calanus finmarchicus* (Wiborg, 1955; Melle et al., 2004), which grazes on phytoplankton and smaller microzooplankton (Yeh et al., 2020). It is an important part of the food chain of planktivorous fish (Kaartvedt, 2000; Utne et al., 2012).

While zooplankton and phytoplankton form the basis of the marine ecosystem, they can, for many reasons, be challenging to observe. First, they have a heterogeneous spatial and temporal distribution, and this patchiness means that the population has to be sampled at several locations over time to get a good understanding of the distribution and its dynamics (Fossum et al., 2019; Zhang et al., 2021). Secondly, the noise levels in different sensor types make it difficult to calibrate the quantities. The most robust way to observe the abundance of phytoplankton is to measure the primary pigment Chlorophyll *a* (Chl *a*) (Huot et al., 2007). Phytoplankton and primary production is today assessed from Chl *a* measurements at sea surface using remote sensing by satellites, planes, or drones (Wu et al., 2022). This data has been invaluable for understanding the temporal and spatial dynamics, but it requires good weather and cloud conditions to obtain high-quality data, and it is restricted to near surface abundance. To get observations deeper in the water column, an in-situ fluorometer can be used to measure Chl *a*. Net samples or water samples are the most accurate ways of measuring plankton intensity, but they are very sparse in space and time. Autonomous underwater vehicles (AUVs) provide more flexibility in obtaining Chl *a* measurements in the ocean column. Note that in-situ measurements and models for the vertical distribution can also be used to calibrate satellite data of Chl *a*, enabling prediction of phytoplankton distribution where satellite data would normally not be usable (Ardyna et al., 2013). Still, there is the issue of low-resolution sampling from satellites.

Observing the zooplankton, specifically *Calanus*, presents new challenges. The first method relies on collecting physical samples in the field; this can be done with water samples or nets. A net can be trawled behind a vessel, and changing the size of the net will change the size of the collected zooplankton (Wiebe et al., 2009). The physical samples then need to be analyzed back in the lab (Goswami, 2004). This accurate method is rather costly. Similarly to the green pigment *chl a* in phytoplankton, it has been suggested that the red pigment can be detected using satellite images (Basedow et al., 2019; Shunmugapandi et al., 2025). This mostly works for surface swarms. However, it is common for *Calanus* to move to the surface during spring. Another method for observing zooplankton is using acoustic back-scattering techniques to observe the abundance and size of marine organisms (Lavery et al., 2010), this method works well for the whole water column. With platforms that can move in the water column, optical sensors such as a SilCam can be used to count zooplankton,

record shape, and possibly classify species (Kydd et al., 2018). These optical sensors can also be used to speed up lab analysis.

In the end, no sensor will provide a perfect solution to the problem of observing plankton. Therefore, a goal should be to combine different sensors and compare them. This can be challenging because sensors do not necessarily provide the same type of data. In addition to the sensors that target the plankton, several supporting sensors can be valuable. Salinity and temperature are important for ocean currents, the mixing of water masses, and other factors important for ocean life (Lalli and Parsons, 1997).

Another important factor is what ocean platform the sensors are placed on. The main focus will be on autonomous surface vehicles (ASVs) and AUVs. While ASVs only move on the surface, they can still gather important information from the water column through acoustics. AUVs have the ability to dive down and collect data that is difficult or impossible to get from the surface. They also have different constraints on how they operate, utilizing the different platforms effectively is therefore important to get the best results.

Dynamic models for the plankton can help borrow information across space and time samples. This can be achieved by coupling numerical ocean models with bio-oceanography understanding. For plankton, there are relatively stable vertical time cycles during the day and night, as discussed in Wirtz and Smith (2020), and tuned to the North Sea conditions and variations (Baars and Fransz, 1984).

In this paper, we report findings from using multiple platforms in experiments conducted in Mausund, Norway. The platforms include AUVs, ASVs and net and rig samples. Our target variable is the biomass integrated over depth and the upper and lower layer boundaries. Both Chl *a* and Calanus measurements are analyzed. We build separate statistical models for each of the variables and sensors and compare the results of different platform data when they are in close proximity to each other in space and time.

In Section 2, we present the operational context, the biomass, zooplankton intensity, and lower and upper depth layer variables, as well as the various data sources. In Section 3, we outline the statistical models used for the variables of interest and the data. In Section 4, we show results and discuss different data attributes. In Section 5, we provide conclusions and point to future work.

2 Operational setting and data sources

2.1 Mausund field campaign and objective

During Spring 2025, we acquired various ocean data in the Mausund archipelago west of Trondheim, Norway (Figure 1). This area is where the fjord waters from the Trondheim fjord meet the Norwegian coastal current waters from further south and the vast Atlantic Ocean. Fishermen have recognized this area as rich in fish resources, possibly due to elevated water mixing providing much nutrient transport and good conditions for primary production (Sætre, 1999).

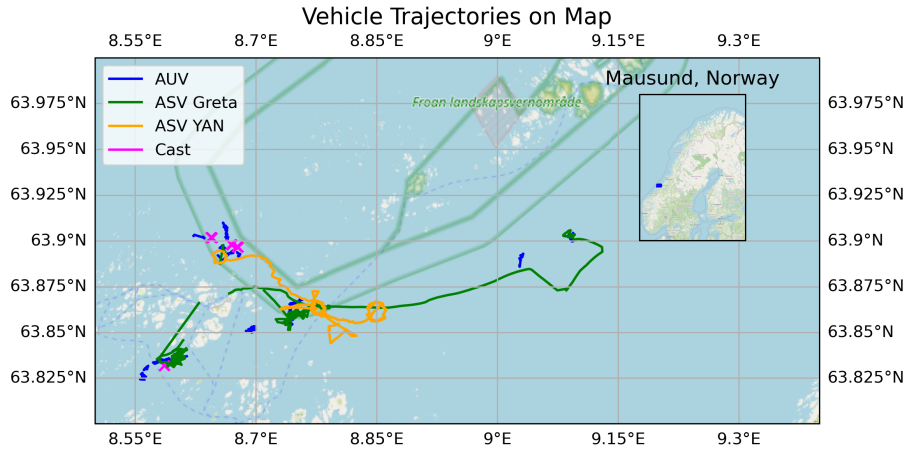


Figure 1: Map of the experimental area off the Western coast of Norway. This shows the path of the vehicles and where the casts are taken

To gain further understanding of the basic driving mechanisms of this complex cold water marine ecosystem, several tests have already been conducted in this region (Fossum et al., 2018; Fragoso et al., 2019, 2024).

The mission in 2025 focused on multi-platform sampling of *Calanus*, which has been observed to be abundant during Spring. From the variety of sensors used in the field sampling of *Calanus*, an objective was to infer the *Calanus* biomass from the different sensors used and to assess the depth interval of *Calanus* zooplankton. This aids our understanding of how the zooplankton intensity changes with space and time, and sparks opportunities for guided sampling efforts to gain further insight.


2.2 Ocean observation platforms

Multiple platforms were used to collect data during the mission.





For this mission, all vehicles were controlled from a control room located approximately 100 km away in Trondheim (Barstein, 2024). With the ability for both data transfer and communication, cooperation between limited personnel in the field and the remaining staff in the operation room was possible. This meant that the boats in the field were not too crowded, and the personnel in the control room could communicate without the added stress of being in the field while maintaining a view of all vehicles and platforms in operation.

Operation boats

The whole operation was managed by two smaller boats, usually with 2 to 4 people on board. Their task was to operate the boats, manage equipment, communicate with vehicles and the control room, and launch platforms. The



Sensors

-  SilCam
-  Echosounder
-  CTD + Chl *a*
-  Physical Samples

Vehicles





-  Boat with operators
-  ASV Grethe
-  ASV AutoNaut
-  LAUVs

Figure 2: The green layer indicates the significant Calanus activity, limited by its upper and lower bound. Five main vehicles are used in the field, with their corresponding sensors.

boats were fitted with a satellite internet connection and served as a connector between the control room and the vehicles in the field. Data were also collected directly from the boats. Physical samples were collected using vertical and horizontal nets. A winch could move a sensor rig up and down in the water column.

ASV Grethe

ASV Grethe is an electric, propeller driven vehicle. The vehicle has a satellite based internet connection, meaning it can send data and be operated from the control room. The ASV has two modes of operation; remote control or autonomous. In remote control mode, live pictures from the forward facing camera and GPS location are used to steer the vehicle; in autonomous mode, the vehicle chooses where to go based on the programmed algorithms or instructions.

ASV AutoNaut

The ASV AutoNaut is also an electric ASV; however, it is driven by wave energy and not by a propeller. With the addition of solar panels, the vehicle can operate continuously, but it does have reduced maneuverability and speed compared with the ASV Grethe.

AUV

Two AUVs were deployed in the field. These vehicles navigate submerged and autonomously collect data as deep as 100m. They are therefore important for collecting bio-chemical data in the water column. A somewhat limiting factor is that the AUVs were launched by hand. Mostly, this is good because it enables fast deployments and retrievals from a small boat, but waves can make retrieval difficult and even dangerous. Thus, there is a limit to the wave height that is safe to operate in.

2.3 Vertical Biomass estimation

To accurately describe what is meant by biomass mapping from sensor data, we introduce the required notation. We let (s_E, s_N) denote east and north coordinates, s_z is depth and t is time. The underlying true but unknown variables of interest at any lateral and time coordinates (s_E, s_N, t) are

- $x_{\text{upper}}(s_E, s_N, t)$: The upper bound of the Calanus layer.
- $x_{\text{lower}}(s_E, s_N, t)$: The lower bound of the Calanus layer.
- $x_{\text{bio}}(s_E, s_N, s_t) = \int_{x_u}^{x_l} x_{\text{bio}}(s_E, s_N, s_z, s_t) ds_z$: The latent log intensity field for the Calanus biomass per square meter over the entire depth column of interest from $x_u = x_{\text{upper}}(s_E, s_N, s_t)$ to $x_l = x_{\text{lower}}(s_E, s_N, s_t)$.

When there is a clear layer of Calanus, the true lower bound $x_{\text{lower}}(s_E, s_N, t)$ is much deeper than the upper bound $x_{\text{upper}}(s_E, s_N, t)$. In practice, we did not need to impose any strict constraints for the lower and upper bounds in our modeling. Figure 2 illustrates the situation with a vertical cross-section having lateral coordinate on the first axis and depth on the second axis. Here, the upper and lower layers vary in depth and the biomass as well.

2.4 Data sources and observation models

Table 1 summarizes the various data sources used in the campaign. SINMOD

Type	Sensors	Measured variables
SINMOD	Numerical bio-ocean model	Calanus
2 AUVs	Turner Cyclops7 SilCam AML (CTD)	Chl <i>a</i> Zooplankton Salinity, Temperature, Depth
2 ASVs	EK80 (Echosounder)	Zooplankton
Net- and rig water samples	SilCam	Zooplankton

Table 1: Overview of the platforms and sensors used to gather data.

is a numerical physical oceanography model for the fluid dynamics in the ocean domain. It is connected to a biological oceanographic model that provides the

Calanus intensity. Figure 3 shows a numerical ocean model prediction from the area. While such numerical models provide much insight into the phenomenon,

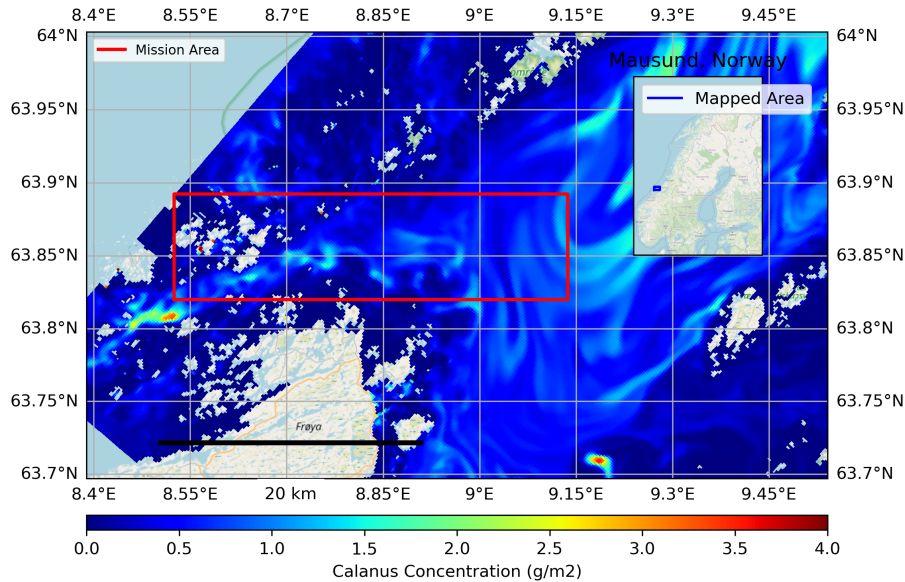


Figure 3: Numerical ocean model prediction of Calanus at a selected time.

they can be quite biased in space and time. In-situ sensor data are the opposite, offering no skill and often sparse data samples, but they have accurate snapshots of the sampled variables at the observation locations. In situ observations are provided here by AUVs and ASVs in addition to net samples at select locations.

Note that the sensors measure at different scales: While AUVs move in the lateral domain and depth, providing rich data at point observations in north, east, and depth coordinates over time, the echosounders installed on the ASVs provide integrated biomass observations and the depth to the upper and lower bounds of this Calanus biomass layer.

In Figure 2, we sketch AUV and ASV data acquisition settings. The ASVs carry a 200 kHz transducer from Kongsberg Discovery AS, Horten, Norway. These sensors are used for detecting the sound scattering layer (SSL) of the zooplankton. The strength of the SSL can be used to estimate the characteristics of the zooplankton distribution in the water column (Liao et al., 1999; Iida et al., 1996) The AUVs carry several sensors of value; the SilCam, Chl *a* fluorometer, salinity, and temperature sensors. The SilCam is used for taking photographs of the Calanus, and when linked with depth and position, this enables us to create depth profiles for the Calanus with relative biomass. To obtain estimates of the water column distribution of the biomass from the AUV measurements, one needs to fit the data to a model for the true variables and then predict for the water column. From the operation boats, a rig fitted with a SilCam

was moved down and up in the water column, collecting particle images. In addition, horizontal nets were trawled to collect physical samples of zooplankton, primarily *C. finmarchicus*.

The various in-situ data sources are discussed next in more detail using mathematical notation. The upper and lower bound data are modeled by

$$\begin{aligned} y_{\nu, \text{lower}}(s_N, s_E, t) &= x_{\text{lower}}(s_N, s_E, t) + \epsilon_{\nu, \text{lower}}(s_N, s_E, t), \\ y_{\nu, \text{upper}}(s_N, s_E, t) &= x_{\text{upper}}(s_N, s_E, t) + \epsilon_{\nu, \text{upper}}(s_N, s_E, t), \end{aligned} \quad (1)$$

where ν refers to the sensor used. The ASVs echosounders provide lower and upper bound depths from onboard processing. For the AUVs, the measurements are computed from a threshold in the Calanus counts, using a window of depth measurements and lateral extent.

For the ASV biomass measurements, we assume the following observation model for the log-biomass:

$$\log y_{\nu, \text{bio}}(s_N, s_E, t) = \alpha_{\nu} x_{\text{bio}}(s_N, s_E, t) + \beta_{\nu} + \epsilon_{\nu, \text{bio}}(s_N, s_E, t). \quad (2)$$

The correction parameters α_{ν} and β_{ν} are needed for the differing types of sensors that are mounted onboard vehicles.

Note that the observation models in equation (1)-(2) are of a conditional type. That is, for the biomass in (2) for instance, we can equivalently write this as

$$\begin{aligned} \log y_{\nu, \text{bio}}(s_N, s_E, t) | x_{\text{bio}}(s_N, s_E, t) &\sim N(\mu_{\nu, \text{bio}}(s_N, s_E, t), r_{\nu, \text{bio}}^2), \\ \mu_{\nu, \text{bio}}(s_N, s_E, t) &= \alpha_{\nu} x_{\text{bio}}(s_N, s_E, t) + \beta_{\nu}, \end{aligned} \quad (3)$$

where r_{ν}^2 terms denote the variance of the observation process for different sensors. With multiple measurements we then get a variance-covariance matrix for the observations, denoted by $\mathbf{R}_{\eta, \nu}$, where η refers to the field type. In our setting with independent error terms in space and time, the matrix becomes diagonal. The models in (1)-(2) are Gaussian and linear, which is the most common assumption in such settings.

SilCam observations of biomass introduce a complication because of Calanus count numbers rather than having a continuous variable like depth or Chl *a* measurements. For SilCam data, one moves out of the Gaussian model assumptions. For the counts data we use a Poisson model with the intensity as a latent variable. The observation model is then defined by

$$y_{\nu, \text{bio}}(s_N, s_E, s_Z, t) | x_{\text{bio}}(s_N, s_E, s_Z, t) \sim \text{Poisson}(\exp(x_{\text{bio}}(s_N, s_E, s_Z, t))). \quad (4)$$

Here, we assume that the samples are over the same water volumes and that the quality of the SilCam data is the same for both AUV vehicles. Only the spatio-temporally varying intensity $x_{\text{bio}}(s_N, s_E, s_Z, t)$ matters in the observation model.

3 Methodology

We now describe our approach for probabilistic assessments of the upper / lower layers and the biomass distribution. Our approach relies on Bayesian methods, using a priori assumptions about the layer depths and the biomass distribution along with the observation models described in Section 2.4.

Prior models are chosen to be Gaussian random fields. Coupled with linear Gaussian observation models, this gives a conjugate Gaussian posterior distribution, conditional on the in-situ observations. The AUV data with SilCam counts are assumed to be Poisson distributed, giving a more complex approximate conjugate model.

The platforms work differently, and sensors also need different processing techniques. We have here chosen to assess the data separately, and then only compare results in the end. Model parameters are then also estimated separately for each platform and each field. For this specification, we use cross-validation.

3.1 Prior models

We build a model for the variables of interest to enable spatio-temporal prediction. This enable estimation and uncertainty assessments about the copepod biomass per square meter, and at what depths (upper and lower) the layer of copepods are located. The model that is used here consists of three independent Gaussian random fields (GRF). Focusing on the ASV platform, the locations $\mathbf{s} \in \mathcal{D} \subset \mathbb{R}^2 \times \mathbb{R}^+$, where \mathcal{D} is the operational space-time domain. The variables are then $x_{\text{bio}}(\mathbf{s})$, $x_{\text{upper}}(\mathbf{s})$ and $x_{\text{lower}}(\mathbf{s})$, at space-time points $\mathbf{s} \in \mathcal{D}$. The GRF models allow for trends in space and time coordinates, variability in realized field variables and spatio-temporal correlations in the fields (Banerjee et al., 2015).

Sets of domain points are denoted by $\mathcal{S} = (\mathbf{s}_{\mathcal{S},1}, \mathbf{s}_{\mathcal{S},2}, \dots, \mathbf{s}_{\mathcal{S},n})$ or $\mathcal{P} = (\mathbf{s}_{\mathcal{P},1}, \mathbf{s}_{\mathcal{P},2}, \dots, \mathbf{s}_{\mathcal{P},m})$. This will be useful for us when we gather data at a collection of sampling point locations at times, set \mathcal{S} , and when the goal is to predict at the same or other locations and times, set \mathcal{P} .

Mathematically, a GRF model means that the distribution of chosen field variables at set \mathcal{P} can expressed using a multivariate Gaussian model:

$$\mathbf{x}_{\eta,\mathcal{P}} \sim N(\boldsymbol{\mu}_{\eta,\mathcal{P}}, \boldsymbol{\Sigma}_{\eta,\mathcal{P}}), \quad (5)$$

where $\boldsymbol{\mu}_{\eta,\mathcal{P}}$ and $\boldsymbol{\Sigma}_{\eta,\mathcal{P}}$ are the a priori mean and covariance matrix, respectively, for any of the field η being log biomass, upper or lower depths. When the cardinality of the set \mathcal{P} is n , the mean is a length n vector while the covariance-matrix is size $n \times n$.

Often, the mean is modeled using a set of covariates or from data. In our case the mean values are typically depending on the depth coordinate. Of course there is also a time of day dependence, but the experiments were done during daytime. The entries of the covariance function are set from correlation functions $\rho(\cdot)$, such that between two points \mathbf{s}_i and \mathbf{s}_j we get covariance

$$\boldsymbol{\Sigma}_{\eta}(i, j) = \sigma_{\eta}^2 \cdot \rho_{\eta,\text{EN}}(|\mathbf{s}_{i,\text{E},\text{N}} - \mathbf{s}_{j,\text{E},\text{N}}|; \varphi_{\eta,\text{EN}}) \cdot \rho_t(|\mathbf{s}_{i,t} - \mathbf{s}_{j,t}|; \varphi_{\eta,t}), \quad (6)$$

where the field variance is σ_η^2 and we assume that the covariance is separable in space and time, and with the spatial correlation function $\rho_{\eta,\text{EN}}(\cdot)$ and in time $\rho_{\eta,t}(\cdot)$. These correlation function are in our case chosen from a Matern class with smoothness 3/2 (Banerjee et al., 2015), which guarantees positive definite covariance for any size n and any subset of points going into the set \mathcal{P} . The correlation parameters $\varphi_{\eta,t}$ in time and $\varphi_{\eta,\text{EN}}$ in the lateral domain must be estimated from data.

For the AUV point sensor, we add depth coordinates for the GRF representation. With the separable covariance structure defined in Equation 6 we then augment the model with a depth component $\rho_z(|s_{i,z} - s_{j,z}|; \varphi_{\eta,z})$.

3.2 Conditioning

Provided field- η observations $\mathbf{y}_{\eta,\nu}$ of data type ν at points \mathcal{S} , the next step is to condition the field mean and covariance. Assuming Gaussian linear models, we get conjugate models (Omre et al., 2024). The conditional Gaussian distribution of field variables $\mathbf{x}_\eta(\mathcal{P})$ at prediction locations \mathcal{P} , given the measurements $\mathbf{y}_{\eta,\nu}(\mathcal{S})$ is:

$$\begin{aligned} \mathbf{x}_\eta(\mathcal{P}) & \mid \mathbf{y}_{\eta,\nu}(\mathcal{S}) \sim \mathcal{N}(\mathbf{m}_{\eta,\mathcal{P}|\mathcal{S}}, \mathbf{V}_{\eta,\mathcal{P}|\mathcal{S}}), \\ \mathbf{m}_{\eta,\mathcal{P}|\mathcal{S}} & = \boldsymbol{\mu}_\eta(\mathcal{P}) + \boldsymbol{\Sigma}_\eta(\mathcal{P}, \mathcal{S}) [\boldsymbol{\Sigma}_\eta(\mathcal{S}) + \mathbf{R}_{\eta,\nu}]^{-1} (\mathbf{y}_{\eta,\nu}(\mathcal{S}) - \boldsymbol{\mu}_\eta(\mathcal{S})), \\ \mathbf{V}_{\eta,\mathcal{P}|\mathcal{S}} & = \boldsymbol{\Sigma}_\eta(\mathcal{P}) - \boldsymbol{\Sigma}_\eta(\mathcal{P}, \mathcal{S}) [\boldsymbol{\Sigma}_\eta(\mathcal{S}) + \mathbf{R}_{\eta,\nu}]^{-1} \boldsymbol{\Sigma}_\eta(\mathcal{P}, \mathcal{S})^\top. \end{aligned} \quad (7)$$

Here, $\mathbf{m}_{\eta,\mathcal{P}|\mathcal{S}}$ is the conditional mean and $\mathbf{V}_{\eta,\mathcal{P}|\mathcal{S}}$ is the conditional covariance matrix. In the updating formulas in Equation (8), $\boldsymbol{\Sigma}_\eta(\mathcal{P}, \mathcal{S})$ is the cross-covariance between the variables in prediction point set \mathcal{P} and observation point set \mathcal{S} .

With Poisson distributed counts (for the AUV sampling), there is not a similar closed form for the posterior. See Appendix A for how a posterior approximation is fit in this situation. For posterior prediction of biomass at unsampled locations, the nonlinear observaton model result in a non-Gaussian posterior. To estimate the biomass integral we use Monte-Carlo sampling. From the posterior approximation, we draw K samples from the water column

$$\mathbf{x}_{\eta, \mathcal{W}_c^{(i)} \approx \mathbf{y}_\eta(\mathcal{S})}^{(k)} \sim \mathcal{N}(\mathbf{m}_{\eta, \mathcal{W}_c^{(i)}|\mathcal{S}}, \mathbf{V}_{\eta, \mathcal{W}_c^{(i)}|\mathcal{S}}), \quad k = 1, \dots, K. \quad (8)$$

Then we compute the numerical integral over the water column

$$I^{(k)}(\mathbf{s}) = \sum_{j=1}^n \exp(x_\eta^{(k)}(z_j)) \cdot \Delta z, \quad (9)$$

for a selected spacing Δz . The samples $\{I^{(k)}(\mathbf{s})\}_{k=1}^K$ are Monte Carlo draws of the integrated biomass. Statistics such as the mean integrated biomass, variance and credible intervals can be computed from the samples.




Figure 4: Illustration of how observations along one path are split into spatio-temporal blocks and folds. Left: spatial partitioning at one time step. Right: one path split into blocks; each color corresponds to a cross-validation fold.

3.3 Parameter estimation and selection

In both prior models and observation models there are several parameters that must be specified. Starting with the observation models, there are the two biomass correction parameters; α_ν and β_ν , for the four different vehicles. Then there are variance parameters in the measurement noise terms $\epsilon_{\eta,\nu}$, for the four vehicles and for the biomass, upper and lower bound variables. For the prior models we must also specify lateral, depth and time correlation parameters and variance parameters for each field and for each vehicle because we are assessing the fields separately.

We employ k -fold cross-validation with spatio-temporal blocks to estimate model parameters. In contrast to standard k -fold cross-validation, where individual observations are randomly assigned to folds, our approach first partitions the data into spatio-temporal blocks. Due to the spatio-temporal dependence structure of the data, such partitioning has been reported to outperform the standard random-fold approach for spatio-temporal data (Valavi et al., 2019). The random fold tends to underestimate predictive uncertainty.

Figure 4 illustrates our method of dividing observations along a single path into spatio-temporal blocks. Once the blocks have been assigned to folds, the usual k -fold cross-validation is performed by iteratively holding out one fold for validation and using the remaining folds for training.

Model performance is evaluated using the predictive negative log-likelihood on the hold-out set. Doing so, the objective function is defined as the mean predictive negative log-likelihood across all folds, which is minimized with respect to the model parameters.

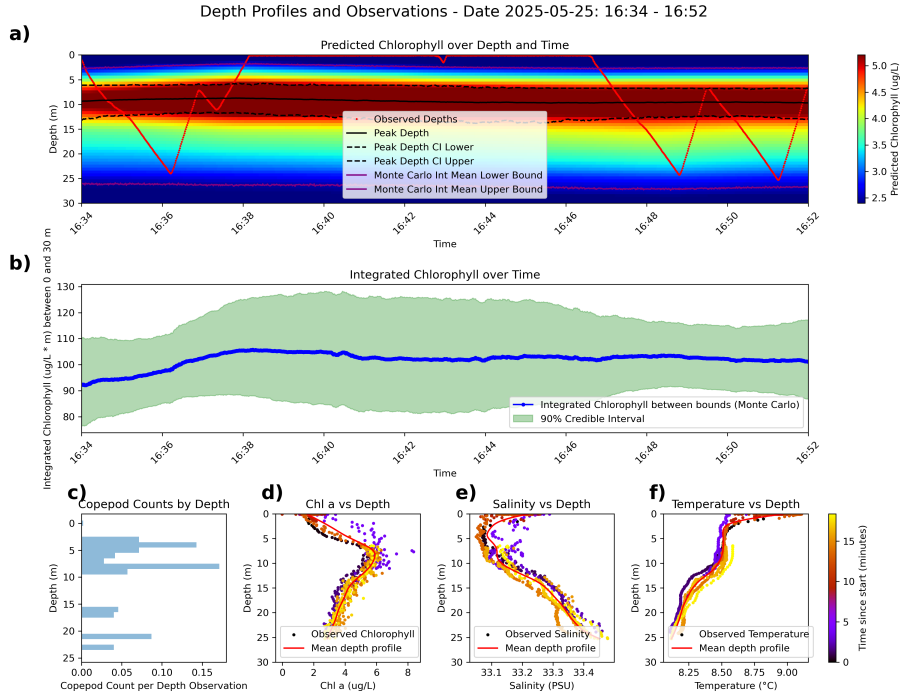


Figure 5: **a)** AUV’s movement in the water column and the predicted Chl *a* values from 0m to 30m. **b)** Integrated Chl *a* between the bounds and a 90 % credible interval. **c)** Number of copepods observed per depth observation. **d)**, **e)** and **f)** Sensor readings for the different depths throughout the transect, the color of the dots indicates how far into the mission the sample is taken

4 Results

The results here compare water column statistics for different vehicles and sensors. The data was collected between the 20th and 28th of May 2025.

4.1 Comparison AUVs

Before comparing observations across the different platforms, it is useful to first examine the data collected by each platform individually. The AUVs are equipped with four sensors that will be considered here: Chl *a*, salinity, temperature, and the SilCam particle camera. It should be noted that SilCam observations are not available for all transects, whereas the other sensors provide measurements for every transect.

Figure 5 presents an example transect performed by one of the AUVs and illustrates several key features of the collected data. In Figure 5(**a**), the vertical trajectory of the AUV through the water column is shown, with color indicating the predicted Chl *a* concentration. From this panel, it can be observed that Chl

a varies relatively little in the horizontal direction over the spatial scale of the transect, whereas substantial variation occurs with depth. The depth-integrated Chl a along the transect is shown in Figure 5(b). A slight increase is observed at the beginning of the transect, after which the values remain relatively stable. The credible interval widens when the AUV approaches the surface, reflecting increased uncertainty in the estimates in this region.

The lower four panels in Figure 5(c)–(f) present the measurements from the four sensors as functions of depth. These plots illustrate how the observed variables vary throughout the water column. In particular, the profiles suggest a relationship between variations in salinity, temperature, and Chl a . Furthermore, there appears to be a higher occurrence of copepods in regions where Chl a concentrations are elevated. It should be noted, however, that SilCam observations can be highly variable, as relatively few captured images contain identifiable copepods.

Because the SilCam was operating during this transect, it is also possible to visualize some of the observed organisms. A collage of selected SilCam images of *Calanus* recorded during the transect is shown in Figure 6. From the data



Figure 6: Collage of Copepods pictured by an AUV during one transect on 2025-05-25.

in Figure 5, it does seem that the zooplankton are roughly at the same depth as the phytoplankton, but this is not necessarily the case. Figure 7 shows the mean integrated Chl a against the mean integrated copepod intensity for different transects. This shows that there was an inverse relationship between the abundance of Chl a and copepods.

In this case, we have deployed two different AUVs, and while they are never




Figure 7: Mean integrated Chl *a* vs. mean integrated copepod intensity for different transects. The error bars indicate the variance for each transect.

at the same location at the same time, they are sometimes close enough to make a comparison. We conduct a pairwise comparison of data from the two AUVs at different locations within a range of distances in space and time. The results are shown in Figure 8, where the three plots show predictions for AUV 1 (first axis) and AUV 2 (second axis) for upper-bound (a), lower-bound (b) and integrated biomass of Chl *a* (c). Each marker in the plot represents the prediction from

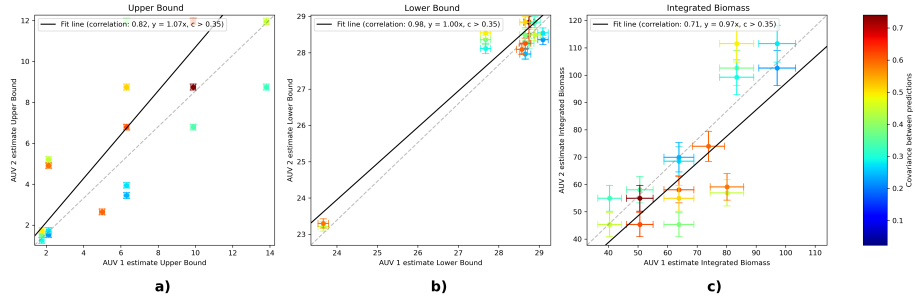


Figure 8: Comparison between the water column statistics for Chl *a* for the two AUVs. **a)** Crossplot and correlation of the estimated upper bound, **b)** Crossplot and correlation in the estimated lower bound and **c)** Crossplot and correlation in the integrated Chl *a* biomass. The color of the points indicates the spatio-temporal correlation as shown in the colorbar.

AUV 1 and AUV 2. The lengths of the vertical and horizontal bars at each mark show the prediction 90 % uncertainty derived from the standard deviation and the normal distribution percentile. The color of a marker is the correlation between the points as defined by the spatio-temporal statistical model; $C_\eta(i, j) = \Sigma_\eta(i, j) / \sigma_\eta^2$ where $\Sigma_\eta(i, j)$ is defined in Equation (6). The color

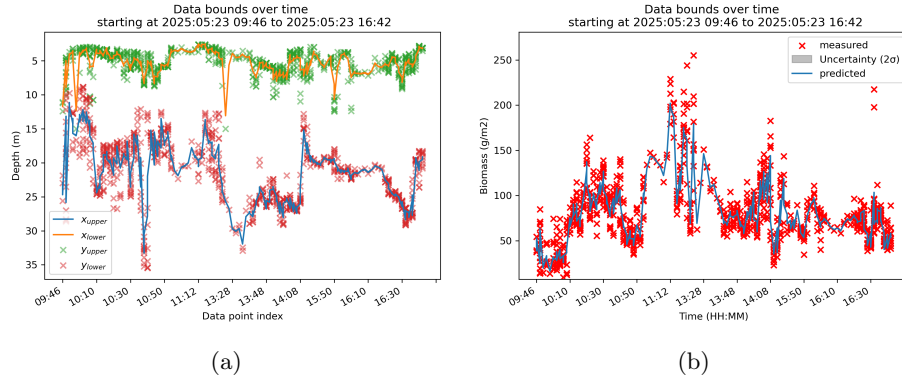


Figure 9: Estimated upper and lower limit (a) changes with time and also the biomass (b).

can then be interpreted as a closeness between the two AUV values, a red color indicates that the AUVs were close in space and time for that measurement. A threshold of 0.25 is chosen for points to be displayed.

The crossplots show that water column statistics from the two AUVs are correlated when they are relatively close. The solid fit line is not very far from the dashed straight line through the origin, and the correlation values (0.82, 0.98, and 0.71) are substantial. There is a tendency for scatter points closer to the diagonal to show more spatio-temporal correlation (red-orange-yellow colors). The correlation in the two AUV predictions is not very surprising, considering that the observations are made by similar vehicles with the same sensors. However, the transects performed by the vehicles are not the same, nor are they regular in the sampling design. This means that water column statistics extracted from an AUV can be compared with statistics from another vehicle. Despite the similarity, we notice that the uncertainty is somewhat underestimated, especially for the upper and lower bounds, as the width of the vertical and horizontal bars does not cover the diagonal.

4.2 Results of ASV sampling

The next major platform used is the ASV. For each location of the trajectory, we can obtain the upper and lower layers, as well as the biomass. This is shown for one vehicle in a select time and spatial domain in Figure 9.

Figure 9 (a) shows how the upper and lower bounds change. The crossed markers are sensor observations, while the solid lines are predictions made by the Gaussian random field model. Figure 9 (b) shows the biomass observations and predictions over the same time interval.

4.3 Comparing AUV and ASV

During the field campaign, the ASV and the AUVs were sometimes in close proximity. This means that we can compare statistics from different vehicle types and sensor types. Figure 10 shows the Chl *a* statistics for an AUV (first axis) and the echosounder biomass estimate from the ASV (second axis). The right plot shows the correlation between the abundance of Chl *a* in the water column measured by the AUV and the biomass estimate estimated from the echosounder on the ASV; if anything, it looks like there is a negative correlation here. In the locations with the highest amount of Chl *a* there was little back-scattering from the echosounder. This is similar to the results found for the single AUV results.

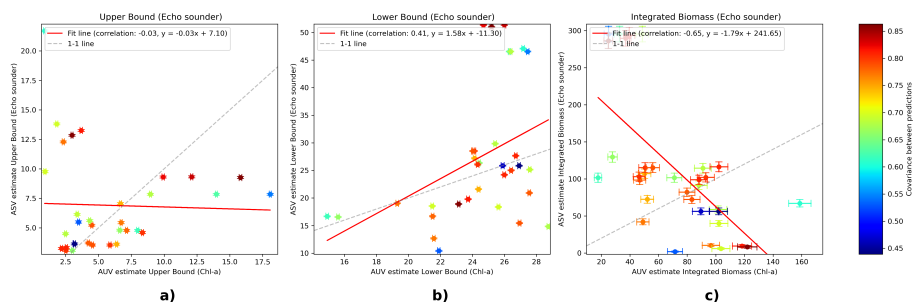


Figure 10: Comparison between the water column statistics for Chl *a* for the one AUV and one ASV. **a)** Crossplot and correlation of the estimated upper bound, **b)** Crossplot and correlation in the estimated lower bound and **c)** Crossplot and correlation in the integrated Chl *a* biomass. The color of the points indicates the spatio-temporal correlation as shown in the colorbar.

The AUVs are also equipped with the SilCam, which we can compare with the echosounder; this comparison is shown in Figure 11. For this figure, we can see that there is a correlation between the biomass estimate in the water column from the AUVs SilCam and the ASV echosounder. This can indicate that the echosounder and the SilCam can be used as complementary sensors to find copepods and estimate the biomass in the water column. It is important to note that a SilCam mounted on an AUV can be a noisy way of detecting copepods. Additionally, there was a smaller selection of transects where the SilCam was turned on and functioning properly.

4.4 SilCam Rig

For two days during the field campaign, the rig with the SilCam was used. This rig setup can be compared with the other vehicles. The issue is that only the AUVs were close enough for a comparison, but the SilCams were not running properly at the same time. Thus, it is mostly just the integrated Chl *a* and integrated Copepod abundance. The resulting correlation between the AUVs

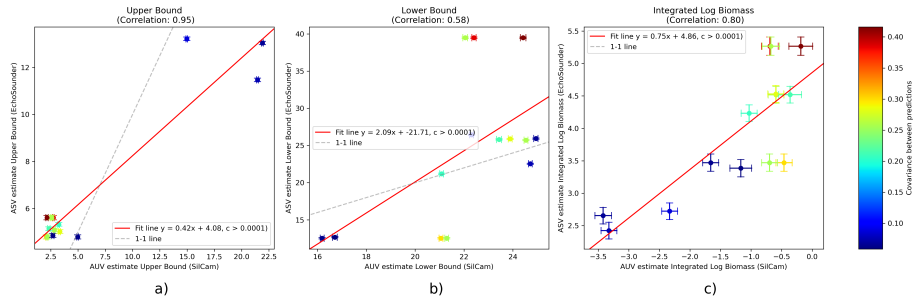


Figure 11: Comparison between the water column statistics for SilCam for the one AUV and one ASV. **a)** Crossplot and correlation of the estimated upper bound, **b)** Crossplot and correlation in the estimated lower bound and **c)** Crossplot and correlation in the integrated SilCam biomass. The color of the points indicates the spatio-temporal correlation as shown in the colorbar.

and the rig can be seen in Figure 12; similarly to what we have seen before, there was an inverse relationship between Chl *a* abundance and copepod abundance.




Figure 12: Comparison between integrated biomass for *a* from AUVs and integrated biomass for copepods from SilCam on rig.

5 Closing remarks

We have reported findings from a multi-vehicle, multi-attribute field deployment in the ocean off the coast of Norway around Mausund during the spring season. Focusing on phytoplankton and the zooplankton *C. finmarchicus*, we compare results obtained from AUVs, ASVs, net and rig samples, and numerical ocean models.

The comparison builds on the use of water column statistics, such as the up-

per and lower depths of the biomass layer and the vertically integrated biomass at different lateral locations. Statistical modeling based on Gaussian random fields and Bayesian statistics is used to infer property attributes based on the various sensors. The AUVs successfully mapped the vertical distribution of phytoplankton and zooplankton biomass in the area, which typically occurred between depths of 5 and 20 m. It should be noted that deeper layers than 30 m would be difficult to detect because of the vehicles' operational depths during the campaign. From SilCam images of Calanus, we noticed correlations in Chl *a*, Calanus counts, and variability in temperature and salinity with depth. Results from comparing AUVs measurements correlated reasonably well when they were close in space and time. Similarly, the ASV was able to map the vertical biomass and the layers of Calanus. Interestingly, the biomass estimate for Calanus aligned reasonably well with the estimates from the AUVs SilCam, but some calibration is necessary to make them truly comparable. These two sensors and vehicles work in different ways, making this comparison noteworthy. Across all platforms, an inverse relationship between Chl *a* and copepod abundance was found.

For future work, we are interested in building a unified model based on all data types and guiding this for hot-spot sampling, as well as for intelligent spatio-temporal exploration of the domain.

Acknowledgments

We acknowledge support from the Norwegian Research Council (RCN) through the MASCOT project 305445 and the SFI Harvest project 309661.

References

- Ardyna, M., Babin, M., Gosselin, M., Devred, E., Bélanger, S., Matsuoka, A., and Tremblay, J.-É. (2013). Parameterization of vertical chlorophyll *a* in the arctic ocean: impact of the subsurface chlorophyll maximum on regional, seasonal, and annual primary production estimates. *Biogeosciences*, 10(6):4383–4404.
- Baars, M. and Fransz, H. (1984). Grazing pressure of copepods on the phytoplankton stock of the central north sea. *Netherlands Journal of Sea Research*, 18(1):120–142.
- Banerjee, S., Carlin, B. P., and Gelfand, A. E. (2015). *Hierarchical modeling and analysis for spatial data*. Chapman and Hall/CRC.
- Barstein, K. H. (2024). Robots collaborate on remote sampling of zooplankton.
- Basedow, S. L., McKee, D., Lefering, I., Gislason, A., Daase, M., Trudnowska, E., Egeland, E. S., Choquet, M., and Falk-Petersen, S. (2019). Remote sensing of zooplankton swarms. *Scientific reports*, 9(1):686.

- Broms, C. and Melle, W. (2007). Seasonal development of calanus finmarchicus in relation to phytoplankton bloom dynamics in the norwegian sea. *Deep Sea Research Part II: Topical Studies in Oceanography*, 54(23):2760–2775. Effects of Climate Variability on Sub-Arctic Marine Ecosystems.
- Fossum, T. O., Eidsvik, J., Ellingsen, I., Alver, M. O., Fragoso, G. M., Johnsen, G., Mendes, R., Ludvigsen, M., and Rajan, K. (2018). Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics*, 35(7):1101–1121.
- Fossum, T. O., Fragoso, G. M., Davies, E. J., Ullgren, J. E., Mendes, R., Johnsen, G., Ellingsen, I., Eidsvik, J., Ludvigsen, M., and Rajan, K. (2019). Toward adaptive robotic sampling of phytoplankton in the coastal ocean. *Science Robotics*, 4(27):eaav3041.
- Fragoso, G., Dallolio, A., Grant, S., Garrett, J., Ellingsen, I., Johnsen, G., and Johansen, T. (2024). The role of rapid changes in weather on phytoplankton spring bloom dynamics from mid-norway using multiple observational platforms. *Journal of Geophysical Research: Oceans*, 129(3):e2023JC020415.
- Fragoso, G. M., Davies, E. J., Ellingsen, I., Chauton, M. S., Fossum, T., Ludvigsen, M., Steinhovden, K. B., Rajan, K., and Johnsen, G. (2019). Physical controls on phytoplankton size structure, photophysiology and suspended particles in a norwegian biological hotspot. *Progress in Oceanography*, 175:284–299.
- Goswami, S. (2004). *Zooplankton methodology, collection & identification-A field manual*. National Institute of Oceanography, Goa.
- Huot, Y., Babin, M., Bruyant, F., Grob, C., Twardowski, M., and Claustre, H. (2007). Does chlorophyll a provide the best index of phytoplankton biomass for primary productivity studies? *Biogeosciences discussions*, 4(2):707–745.
- Iida, K., Mukai, T., and Hwang, D. (1996). Relationship between acoustic backscattering strength and density of zooplankton in the sound-scattering layer. *ICES Journal of Marine Science*, 53(2):507–512.
- Kaartvedt, S. (2000). Life history of calanus finmarchicus in the norwegian sea in relation to planktivorous fish. *ICES Journal of Marine Science*, 57(6):1819–1824.
- Kydd, J., Rajakaruna, H., Briski, E., and Bailey, S. (2018). Examination of a high resolution laser optical plankton counter and flowcam for measuring plankton concentration and size. *Journal of Sea Research*, 133:2–10.
- Lalli, C. M. and Parsons, T. R. (1997). Chapter 3 - phytoplankton and primary production. In *Biological Oceanography: An Introduction (Second Edition)*, pages 39–73. Butterworth-Heinemann, Oxford, second edition edition.

- Lavery, A. C., Chu, D., and Moum, J. N. (2010). Measurements of acoustic scattering from zooplankton and oceanic microstructure using a broadband echosounder. *ICES journal of marine science*, 67(2):379–394.
- Liao, C.-H., Lee, K.-T., Lee, M.-A., and Lu, H.-J. (1999). Biomass distribution and zooplankton composition of the sound-scattering layer in the waters of southern east china sea. *ICES Journal of Marine Science*, 56(5):766–778.
- Melle, W., Ellertsen, B., Skjoldal, H. R., et al. (2004). Zooplankton: the link to higher trophic levels. *The Norwegian sea ecosystem*, 1:148–149.
- Omre, K. H., Fjeldstad, T. M., and Forberg, O. B. (2024). *Bayesian Spatial Modelling with Conjugate Prior Models*. Springer.
- Rey, F. (2004). Phytoplankton: the grass of the sea. *The Norwegian Sea Ecosystem*, 6(1):97–136.
- Sætre, R. (1999). Features of the central norwegian shelf circulation. *Continental Shelf Research*, 19(14):1809–1831.
- Shunmugapandi, R., McCarry, C. L., McKee, D., and Mitchell, C. (2025). Ocean color anomaly detection to estimate surface calanus finmarchicus concentration in the gulf of maine. *Frontiers in Marine Science*, 12:1507638.
- Utne, K. R., Hjøllø, S. S., Huse, G., and Skogen, M. (2012). Estimating the consumption of calanus finmarchicus by planktivorous fish in the norwegian sea using a fully coupled 3d model system. *Marine Biology Research*, 8(5-6):527–547.
- Valavi, R., Elith, J., Lahoz-Monfort, J. J., and Guillera-Arroita, G. (2019). blockcv: An r package for generating spatially or environmentally separated folds for k-fold cross-validation of species distribution models. *Methods in Ecology and Evolution*, 10(2):225–232.
- Wiborg, K. F. (1955). Zooplankton in relation to hydrography in the norwegian sea.
- Wiebe, P., Benfield, M., Steele, J., Thorpe, S., and Turekian, K. (2009). Zooplankton sampling with nets and trawls. *Measurement Techniques, Sensors and Platforms (Elsevier, 2009)*, pages 70–86.
- Wirtz, K. and Smith, S. L. (2020). Vertical migration by bulk phytoplankton sustains biodiversity and nutrient input to the surface ocean. *Scientific Reports*, 10(1):1142.
- 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.

Yeh, H. D., Questel, J. M., Maas, K. R., and Bucklin, A. (2020). Metabarcoding analysis of regional variation in gut contents of the copepod calanus finmarchicus in the north atlantic ocean. *Deep Sea Research Part II: Topical Studies in Oceanography*, 180:104738.

Zhang, Y., Ryan, J. P., Hobson, B. W., Kieft, B., Romano, A., Barone, B., Preston, C. M., Roman, B., Raanan, B.-Y., Pargett, D., et al. (2021). A system of coordinated autonomous robots for lagrangian studies of microbes in the oceanic deep chlorophyll maximum. *Science Robotics*, 6(50):eabb9138.

A Posterior approximation of Calanus intensity

Due to the count data of Calanus, there is no closed form posterior model conditional on the observed counts. Instead, we here approximate the posterior model with a Gaussian fit at the posterior mode.

The approximation is based on iteratively fitting a Gaussian model by linearizing the Poisson likelihood model for count data to a quadratic form. For simplicity, the prediction location set \mathcal{P} is here assumed to be the same as the observation location set \mathcal{S} . At each step in the iterative method, Equation 8 is then slightly modified as

$$\begin{aligned} \mathbf{x}_{\text{bio}}(\mathcal{S})|\mathbf{y}_{\text{bio}}(\mathcal{S}) &\approx N(\hat{\mathbf{m}}_{\mathbf{x}_A|\mathbf{y}_A}, \hat{\mathbf{V}}_{\mathbf{x}_A|\mathbf{y}_A}), \\ \hat{\mathbf{m}}_{\mathbf{x}_{\text{bio},\mathcal{S}}|\mathbf{y}_{\text{bio},\mathcal{S}}} &= \boldsymbol{\mu}_{\text{bio},\mathcal{S}} + \boldsymbol{\Sigma}_{\text{bio},\mathcal{S}}(\boldsymbol{\Sigma}_{\text{bio},\mathcal{S}} + \mathbf{P})^{-1}(\mathbf{z} - \boldsymbol{\mu}_{\text{bio},\mathcal{S}}), \\ \hat{\mathbf{V}}_{\mathbf{x}_{\text{bio},\mathcal{S}}|\mathbf{y}_{\text{bio},\mathcal{S}}} &= \boldsymbol{\Sigma}_{\text{bio},\mathcal{S}} - \boldsymbol{\Sigma}_{\text{bio},\mathcal{S}}(\boldsymbol{\Sigma}_{\text{bio},\mathcal{S}} + \mathbf{P})^{-1}\boldsymbol{\Sigma}_{\text{bio},\mathcal{S}}. \end{aligned} \tag{10}$$

Now, the values \mathbf{z} and \mathbf{P} are recalculated at every iteration in a Newton-Raphsons algorithm. Algorithm 1 indicates the steps of the procedure and for outputting the approximate posterior mode and covariance matrix in Equation (10).

With a different prediction set, one can introduce a linear sampling operator that picks the observation sites but leaves the prediction sites unobserved. This would entail building a sparse matrix which consists of mostly 0 entries and with some 1 entries. The 1 indicates a sampling site in a grid representation of the domain.

Algorithm 1 Newton-Raphson Algorithm for posterior approximation in the situation with a Poisson observation model.

Require:

Observations \mathbf{y}_{bio} from a path \mathcal{S} , prior-mean $\boldsymbol{\mu}_{\text{bio}}(\mathcal{S})$, and covariance matrix $\boldsymbol{\Sigma}_{\text{bio}}(\mathcal{S}, \mathcal{S})$ and a tolerance δ

$x^{(0)} = \log(\mathbf{y}_{\text{bio}} + \mathbf{1}\epsilon)$ as the initial guess (minuscule $\epsilon > 0$)

repeat For each time $k = 0, 1, \dots$: ▷ Loop until convergence

Get \mathbf{z} with entries $z_i = \left(y_i - \exp(x_i^{(k)}) + x_i^{(k)} \cdot \exp -x_i^{(k)} \right) \cdot \exp(x_i^{(0)})$

Get \mathbf{P} with entries $P_{i,i} = \exp(x_i^{(k)})$ and $P_{i,j} = 0$ when $i \neq j$

Get $\hat{\mathbf{m}}_{\mathbf{x}_{\mathbf{A}}|\mathbf{y}_{\mathbf{A}}}^{(k)}$ and $\hat{\mathbf{V}}_{\mathbf{x}_{\mathbf{A}}|\mathbf{y}_{\mathbf{A}}}^{(k)}$ from Equation (10)

$\mathbf{x}^{(k)} = \hat{\mathbf{m}}_{\mathbf{x}_{\mathbf{A}}|\mathbf{y}_{\mathbf{A}}}^{(k)}$

if $\|\mathbf{x}^{(k)} - \mathbf{x}^{(k-1)}\| < \delta$ **then**

Return $\hat{\mathbf{m}}_{\mathbf{x}_{\mathbf{A}}|\mathbf{y}_{\mathbf{A}}}^{(k)}$, $\hat{\mathbf{V}}_{\mathbf{x}_{\mathbf{A}}|\mathbf{y}_{\mathbf{A}}}^{(k)}$, \mathbf{z} and \mathbf{P}

end if

until
