



Dynamic stochastic modeling for adaptive sampling of environmental variables using an AUV

Gunhild Elisabeth Berget^{1,2} · Jo Eidsvik³ · Morten Omholt Alver¹ · Tor Arne Johansen¹

Received: 28 June 2021 / Accepted: 10 February 2023
© The Author(s) 2023

Abstract

Discharge of mine tailings significantly impacts the ecological status of the sea. Methods to efficiently monitor the extent of dispersion is essential to protect sensitive areas. By combining underwater robotic sampling with ocean models, we can choose informative sampling sites and adaptively change the robot's path based on in situ measurements to optimally map the tailings distribution near a seafill. This paper creates a stochastic spatio-temporal proxy model of dispersal dynamics using training data from complex numerical models. The proxy model consists of a spatio-temporal Gaussian process model based on an advection–diffusion stochastic partial differential equation. Informative sampling sites are chosen based on predictions from the proxy model using an objective function favoring areas with high uncertainty and high expected tailings concentrations. A simulation study and data from real-life experiments are presented.

Keywords Adaptive sampling · Gaussian process · Stochastic modeling · AUV · Mine tailings

We acknowledge support from the Norwegian Research Council (RCN) through the Centers of Excellence funding scheme, Project Number 223254 - Centre for Autonomous Marine Operations and Systems (NTNU-AMOS), and the INDORSE Project 267793. The authors would like to thank Tore Mo-Bjørkelund from NTNU and Frédéric Py from SINTEF Digital for help related to the field experiments, and Tor Nordam and Finn Are Michelsen from SINTEF Ocean for running the numerical models SINMOD and DREAM and supplying the data sets used in this paper.

✉ Gunhild Elisabeth Berget
gunhild.berget@ntnu.no

Jo Eidsvik
jo.eidsvik@ntnu.no

Morten Omholt Alver
morten.alver@ntnu.no

Tor Arne Johansen
tor.arne.johansen@ntnu.no

¹ Department of Engineering Cybernetics and Centre for Autonomous Marine Operations and Systems, NTNU, Trondheim, Norway

² Sintef Community, Trondheim, Norway

³ Department of Mathematical Sciences, NTNU, Trondheim, Norway

1 Introduction

The world's oceans are heavily impacted by human activities, such as pollution and overfishing (Halpern et al., 2008). One example is the harmful effects of mine tailings, drill cuttings, and mud discharge. Along the Norwegian continental shelf, there are sensitive species like cold-water coral reefs and sponges which are negatively impacted by such deposition (Tranum et al., 2010), but still, there exist several active mining seafills and new seafills are in the planning phase. Improved methods for monitoring the extent of the contamination from seafills are important to ensure that sensitive areas are protected. In this paper, we use a seafill located in Frønfjorden (see Fig. 1) as a case study where we validate our proposed method for monitoring the tailings distribution close to the outlet.

When mapping the tailings distribution, models of ocean dynamics can assist us with helpful information. However, building such models is challenging because of the oceans' large-scale nonlinear processes and high spatio-temporal variability. Existing models use numerical methods to improve their accuracy (Griffies et al., 2000) continuously, but they still have errors due to simplifications, limitations in numerical implementation, uncertain parameters, and initial and boundary conditions. Thus, data assimilation using in situ measurements can help improve the model accuracy.




Fig. 1 Map of the Frænfjorden area. The pink contours in the figure show the location of the factory Omya Hustadmarmor, and the black lines indicate the pipes of the seafill. The numbers (2.1–2.4) and (1.1–1.2) mark possible outlets for the tailings. At the time of the field experiment, outlet 2.3 was the operational discharge. The blue rectangle represents the operational area for the simulation study and field experiments (Color figure online)

Measurements are commonly obtained using remote sensing, ships, or buoys, but this data is expensive to collect, making it hard to observe the environment in detail. Recently, robotic vehicles like autonomous underwater vehicles (AUVs) have become more robust and less expensive, and by equipping them with sensors, one can collect measurements (sample) more efficiently. Still, the possible observations are sparsely distributed, leading to significant gaps in our understanding of the ocean (Stewart, 2008). This measurement limitation prompts the need for intelligent and targeted observation designs, and an important question is when and where to sample.

This paper addresses this question and suggests an adaptive sampling method utilizing prior data from ocean models and real-time in situ measurements from an AUV. Because the already existing ocean dynamics models have a high computational load, making them unfit for running on embedded robotic systems, an onboard low-complexity proxy model is created so that real-time adaption of the mission can be obtained. The proxy model represents the state of the ocean at the time and can be updated when new information is available. We use a spatio-temporal Gaussian process (GP) with a kernel describing spatial dependencies and a temporal model built using an advection–diffusion stochastic partial differential equation (SPDE). The linear SPDE model and the Gaussian measurement model are written as a state-space model, and the Kalman filter is used to update the state recursively over time. Two existing numerical models, SINMOD (ocean model) and DREAM (particle dynamics model) are used to train the compact GP model creating a proxy model of the tailing distribution. Sensor readings on the AUV can then be used to update the proxy model onboard in real-time. An objective function for waypoint selection is presented to maximize the value of information from the measurements.

This function ensures that the area is explored by choosing locations assumed to be information-rich and also considers the travel length limitations of the AUV, leading to an adaptive sampling strategy.

In our case study, the method is tested in simulation using the Frænfjorden area (Norway), a fjord containing a seafill for mine tailings. The fjord can be seen in Fig. 1, where the location is marked with a red circle. We compare the proposed method to four simpler strategies in the simulation study. As a proof of concept, real-life experiments are carried out in the same area.

The paper is organized as follows: Sect. 2 presents background information on the Frænfjorden area, the ocean models, and adaptive sampling. In Sect. 3, the development of the proxy model is presented. Section 4 presents a combined state space model explaining how the data assimilation is done before the sampling strategies are discussed in Sect. 5. Section 6 contains simulation results comparing our proposed method to four simpler strategies, and results from experimental field trials are presented in Sect. 7. Section 8 discusses the results before Sect. 9 gives a conclusion and future work.

1.1 Related work and contribution

With the improved technology and robustness of AUVs, adaptive sampling using these robotic platforms has had increased focus lately, and a review of the existing adaptive sampling methods for AUV missions can be found in Hwang et al. (2019). Most adaptive sampling methods utilize GPs (Cressie & Wikle, 2011), which are widely used to create non-parametric and computationally efficient models. They are popular not only for adaptive sampling but used in various applications like neuroimaging (Hyun et al., 2016), machine learning (Rasmussen & Williams, 2005) and sensor placement (Krause et al., 2008). Combining GPs with robotic vehicles is amongst others explored in Zhang et al. (2012) where an AUV is used to track an upwelling front, and in Luo and Sycara (2018), which explore a method of multiple vehicles using a mixture of GP models. Foss et al. (2022) use a spatio-temporal model for the same case we are studying here, but with another objective function relying on specifying a threshold for the particle concentration.

This work presents an implementation of an adaptive sampling method for environmental sensing of mine tailings close to a seafill. We show how to utilize prior information from ocean models to create GP models. Similar approaches, also utilizing prior knowledge to learn GP models, can be found in Das et al. (2015) who use an AUV to collect samples for ex-situ analysis, selecting the sampling locations based on previous mission data and maximizing a utility function, and in Ma et al. (2018), who create a model of a spatio-temporal environment by learning and refining the hyperparameters of

a GP kernel using prior ocean model data. In our approach, the ocean model data are used to develop a non-stationary correlation kernel as in Berget et al. (2018) and Fossum et al. (2018) and set the parameters in a temporal model based on the stochastic advection–diffusion partial differential equation (Sigrist et al., 2015). This allows the model to capture critical dynamic phenomena such as tides. Estimated solutions of the SPDE can be written in a linear form so that the Kalman filter can be used to update the model state. An objective function based on variance reduction, hotspot detection, and AUV operation time constraints is used; see also Berget et al. (2018). The approach sequentially selects waypoints based on the current model state, and the method is tested in both simulations and field experiments near a mining seafill in a Norwegian fjord.

The main contributions of this paper are:

- Propose an adaptive sampling method for environmental sensing of mine tailings.
- Use ocean model data to set the parameters in a linear temporal model based on a stochastic advection–diffusion partial differential equation.
- Present results from real-life experiments using the proposed adaptive sampling method.

2 Background

This section presents definitions and background information on ocean models, the spatio-temporal model, and data assimilation.

2.1 The Frønfjorden area

In Norway, several active and planned seafills have implications for the marine environment (Ramirez-Llodra et al., 2015; Morello et al., 2016). The tailings typically consist of particulates and residual chemicals from processing, and this waste will smother the sea bed close to the sea fill. In addition, some of the waste may be transported by ocean currents away from the outlet, thus affecting a larger area of the sea bed and the water column.

As the case study for this paper, we focus on a seafill located in Frønfjorden, Norway (see Fig. 1). A calcium carbonate processing plant in the factory has been operating for several decades. This processing of ore from mining causes fine-ground waste (tailings) (Kvassnes & Iversen, 2013), disposed of on the ocean floor, in a seafill close to the factory.

The outlets in Frønfjorden are located at approximately 20 m depth, and tailings are continuously discharged with a time-varying discharge rate. A typical discharge rate from the outlets is around 400 m³/h. The discharge alternates between two pipes, illustrated by black lines in Fig. 1. The operational

discharge for the data in the simulation study and the field experiments is on the western pipe. This outlet is marked as 2.3 in the figure.

2.2 Ocean models

Numerical oceanographic models are used in this paper to build a prior belief of the state of the ocean and to predict the movement of particles with the currents. Ocean models provide information on salinity, currents, temperature, density, and pressure and describe the state of the ocean at a given time. The models use a set of hydrodynamic and thermodynamic equations called the primitive equations, and these are commonly solved numerically using finite difference methods. This makes the models computer intensive; hence, high-resolution predictions can only be computed for limited areas, depending on the available computer resources. Input to ocean models typically includes tides, sea-level pressure, wind, heat exchange, bathymetry, and freshwater runoff. Ocean model outputs are uncertain due to limits in model resolution, limits to our knowledge of the processes resolved by the model, errors in initial values, boundary conditions and parameter values, and inaccuracies in numerical implementations (Lermusiaux et al., 2006). It is, therefore, necessary to validate the output against observations, and data assimilation can be used to correct the model state estimates. For monitoring of releases of mine tailings or other particles or dissolved material, there is a need to develop enabling technology for efficient and targeted ocean sampling. Observations from different platforms such as buoys, ship-based sampling, or robotic sampling from, for example, AUVs can be used to evaluate or, through data assimilation, improve the model estimates.

In this paper, we use data from two specific models to build the stochastic onboard model, SINMOD (ocean dynamics model) (Slagstad & McClimans, 2005; Wassmann et al., 2006) and DREAM (particle transportation model) (Rye et al., 1998, 2008). Nepstad et al. (2020) described the SINMOD and DREAM setups for Frønfjorden and evaluated the models against field measurements. The relation between SINMOD, DREAM, and the onboard model can be seen in Fig. 2. More details on the two ocean models are given below.

2.2.1 SINMOD: ocean dynamics model

SINMOD is a numerical ocean model system that connects and simulates physical and biological processes in the ocean (Slagstad & McClimans, 2005). The model is based on the Navier–Stokes equations and uses a nesting technique where high resolution models obtain their boundary conditions from larger model domains with lower resolution. SINMOD is established in configurations with horizontal resolutions ranging from 20 km to 32 m. Input to the model

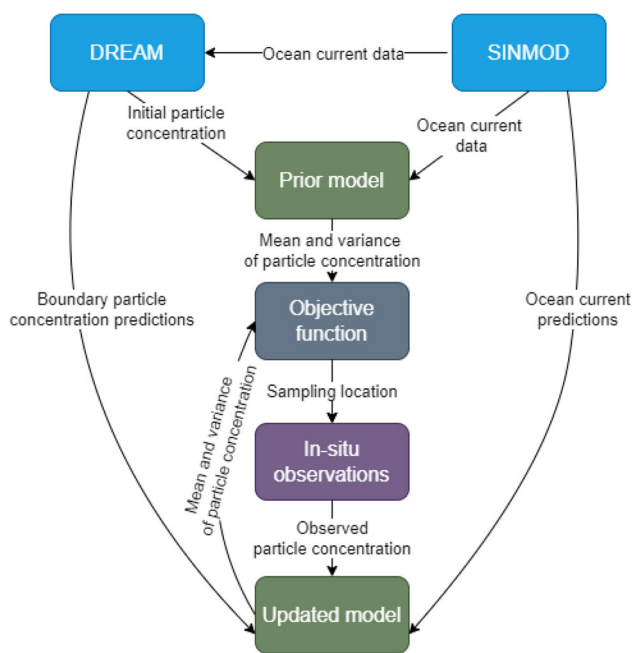


Fig. 2 The relation between SINMOD, DREAM, and the adaptive sampling framework. SINMOD provides ocean currents data as input to the particle transportation model DREAM. Both SINMOD and DREAM provides data used to build the prior belief of the onboard model. Ocean current predictions are used in the temporal part of the model and particle concentration predictions on the boundary are used in the updating of the model

includes atmospheric forcing, freshwater runoff, and boundary conditions (temperature, salinity, tidal forcing, currents). Ocean dynamics data from SINMOD is used in this case study as input to DREAM and the onboard proxy model (see Fig. 2). Atmospheric forcing for daily forecast simulations was obtained from NOAA's Global Forecast System with 0.5° resolution (for large-scale model domains) and from the Norwegian Meteorological Institute's MetCoOp Ensemble Prediction System with 2.5 km resolution (for local model domains).

2.2.2 DREAM: particle transportation model

DREAM is a Lagrangian particle transport model which can be used to simulate the behavior and fate of marine pollutants, including particulate discharges from drilling operations (Rye et al., 1998, 2008). It provides time series of concentrations of released materials in the water column and deposition of these materials onto the sea floor. Input to the DREAM model includes hydrodynamic data delivered by SINMOD and information about the release (amount, rate, densities, grain size distribution). DREAM is often used as a decision support tool for managing operational discharges to the marine environment. This paper uses data from DREAM to set the initial belief for the mine tailings distribution on

the onboard proxy model. This data is also used to estimate parameters in the SPDE model, including the GP kernel.

2.3 Adaptive sampling

Observing the ocean by taking measurements or samples is crucial to increasing our understanding of ocean processes. Ocean sampling can be seen as the action of sensing or measuring one or more physical properties in the ocean at a specific point in space and time. Traditional methods include mounting sensors to buoys which limits the sampling to only one location, or immersing instruments in the water from boats, commonly using a predetermined plan of sampling locations. Exploiting robotic platforms and especially AUVs can intensify ocean sampling, giving more data and better insight.

We need more sophisticated sampling strategies to exploit the benefits that robotic platforms have regarding autonomy and increased mobility. To obtain the most scientifically relevant measurements effectively, adaptive approaches need to be developed. In contrast to static/pre-scripted schemes, adaptive/data-driven strategies can react to measured conditions, having access to both prior and in situ information; selecting sampling locations depending on past observations taken during exploration.

The benefits of adaptive sampling will depend on the situation. There are numerous studies on the value of additional (sequential) sampling efforts in different contexts, say clinical trials in medical research (Jennison & Turnbull, 1999) or geoscience data (Eidsvik et al., 2015). In our case, there would be limited to gain by adaptation if a priori information makes it possible to pre-script a high-value AUV plan before deployment. However, if there is high prior variability, the ability to adapt the initial plan one way or the other, based on onboard measurements, would bring added value by guiding the AUV to more relevant places.

3 Modeling

This section describes the advection–diffusion SPDE model and its associated spatio-temporal GP representation. We also outline parameters in the model.

3.1 Spatio-temporal onboard model

To select optimal sampling locations in a dynamic domain, information about the spatial conditions at all times is essential. Since the computational demands using ocean models such as SINMOD and DREAM are too high for real-time onboard modeling, we suggest a stochastic spatio-temporal proxy model based on an advection–diffusion SPDE. When this SPDE is discretized, it represents a spatio-temporal GP

model. Apart from having a smaller computational footprint, GPs are conventional tools for statistical modeling of spatial data and have been widely adopted in oceanographic applications (Binney et al., 2013). In this paper, the proxy model is used to approximate the distribution of the particle concentration near the seafill in Frænfjorden.

There exist various types of spatio-temporal GP models, see e.g. Chapter 6 of Cressie and Wikle (2011) for a review. For statisticians, it has been important to understand the spatio-temporal covariance function of such processes, which is crucial in spatio-temporal Kriging. In the simplest case, the covariance function is separable in space and time, but there has also been much work on non-separable covariance functions (Gneiting, 2002). Stochastic differential and difference equations in the spatial domain fit naturally into state-space models, which enable spatio-temporal Kalman filtering using in-situ observations. The spatial version of the autoregressive process is one such example of a linear difference equation. So are integro-difference equations which weight spatial variables from the previous time step to predict the spatial variables at the current time step (Storvik et al., 2002). Recently, extensions of the advection–diffusion equation have been developed (Sigrist et al., 2015; Richardson, 2017). These models provide more interpretability in the spatial weighting as well as opportunities for embedding time-dependent drift. We use one of these models in the current paper. The advection–diffusion equation is also endorsed in the physical modeling community, and in our setting the advection can be specified from the tidal currents and the diffusion coefficients can be set from the ocean models SINMOD and DREAM.

We consider a real-valued stochastic process $\{X(t, s), t \geq 0, s \in \mathcal{D}\}$, representing the uncertain particle concentration on log scale in (east,north) location $s = (s_e, s_n) \in \mathcal{D} \subset \mathbb{R}^2$ at time t .

The spatio-temporal process of transport of a medium in a fluid can be expressed by the advection–diffusion SPDE (Sigrist et al., 2015)

$$\frac{\partial X(t, s)}{\partial t} = \lambda X(t, s) - \mathbf{c}(t, s)^T \nabla X(t, s) + \nabla \mathbf{D} \nabla X(t, s) + B(t, s) + \epsilon(t, s), \quad (1)$$

with $\mathbf{c}(t, s) = (c_e(t, s), c_n(t, s))^T$ is the drift vector field for advection, \mathbf{D} the diffusion matrix, $\epsilon(t, s)$ a noise term, ∇ the gradient operator $\left(\frac{\partial}{\partial s_e}, \frac{\partial}{\partial s_n}\right)$ and $\lambda \in [-1, 0]$ a damping constant controlling the autoregressive relationship between state vectors (Richardson, 2017). The noise term $\epsilon(t, s)$ is assumed to be a zero-mean Gaussian noise term that is uncorrelated in the temporal domain but spatially correlated, $\text{Corr}(\epsilon(t, s), \epsilon(t, s')) = R(|s - s'|)$, where $R(\cdot)$ represents a correlation function for distance input. The source term

$B(t, s)$ describes external information, e.g. boundary information.

This paper assumes horizontal, isotropic diffusion leading to a constant diffusion parameter D . We denote the 2-dimensional drift vector as $\mathbf{c}(t, s) = (c_e^t, c_n^t)^T$ at time step t . Then we get

$$\frac{\partial X(t, s)}{\partial t} = \lambda X(t, s) - \left(c_e^t \frac{\partial X(t, s)}{\partial s_e} + c_n^t \frac{\partial X(t, s)}{\partial s_n} \right) + D \left(\frac{\partial^2 X(t, s)}{\partial s_e^2} + \frac{\partial^2 X(t, s)}{\partial s_n^2} \right) + B(t, s) + \epsilon(t, s). \quad (2)$$

The SPDE is discretized using the upwind differencing scheme to find an approximate solution. We define the length of a time step as dt . Forward differences are used in time, giving

$$\frac{\partial}{\partial t} X(t, s) \approx \frac{X(t + dt, s) - X(t, s)}{dt}. \quad (3)$$

We define a regular grid D_s with a total of N grid points with resolution ds_e in the east direction and ds_n in the north direction. The set of spatial locations (east, north) is then $s_i = (s_{e,i}, s_{n,i}) \in D_s \subset \mathbb{R}^2, i = 1, \dots, N$. For the first derivative in space, we alternate between forward and backward differences based on the direction of the drift vector \mathbf{c}_t . In the eastward direction, we have

$$\frac{\partial}{\partial s_e} X(t, s) \approx \frac{X(t, s \pm (ds_e, 0)^T) - X(t, s)}{ds_e}, \quad (4)$$

where the upper sign is used for $c_e^t < 0$ and the lower is used when $c_e^t \geq 0$. Similarly, in the northward direction, we have

$$\frac{\partial}{\partial s_n} X(t, s) \approx \frac{X(t, s \pm (0, ds_n)^T) - X(t, s)}{ds_n}, \quad (5)$$

where the upper version is used for $c_n^t < 0$ and the lower is used when $c_n^t \geq 0$. The second derivatives are discretized as follows,

$$\frac{\partial^2}{\partial s_e^2} X(t, s) \approx \frac{X(t, s + (ds_e, 0)^T) - 2X(t, s) + X(t, s - (ds_e, 0)^T)}{ds_e^2} \quad (6)$$

$$\frac{\partial^2}{\partial s_n^2} X(t, s) \approx \frac{X(t, s + (0, ds_n)^T) - 2X(t, s) + X(t, s - (0, ds_n)^T)}{ds_n^2}. \quad (7)$$

By scaling the time discretization, a change in time can be denoted $t + 1$ instead of $t + dt$. We have $\{t = t_0 + j \times dt; j =$

$0, 1, 2, \dots, T$). When inserting all the discretized terms into eq. (1), the scheme can be written as a linear process on the form

$$\mathbf{X}_t = \mathbf{A}_t \mathbf{X}_{t-1} + \mathbf{B}_t + \boldsymbol{\epsilon}_t, \quad \boldsymbol{\epsilon}_t \sim N(0, \mathbf{W}_t), \quad t = 1, \dots, T, \quad (8)$$

where \mathbf{X}_t is the size $N \times 1$ vector of particle concentrations at time step t for all spatial locations, \mathbf{A}_t is the $N \times N$ propagator matrix including information about advection, diffusion and damping, \mathbf{B}_t is a matrix adding external information, in our case boundary conditions, and \mathbf{W}_t is a positive definite variance-covariance matrix for the noise term $\boldsymbol{\epsilon}_t$.

The particle concentrations at time $t = 0$ are assumed to be Gaussian distributed with mean vector $\boldsymbol{\mu}_0 = [\mu_1, \dots, \mu_N]$ and $N \times N$ variance-covariance matrix $\boldsymbol{\Sigma}_0$.

For the boundary conditions in the SPDE, we use Dirichlet BCs on the east side of our survey area. This can be considered as a constant forcing along the east side and is chosen because of the constant flow of mine tailings that are let out on this side of the rectangle. We assume no drift across the boundary on the remaining sides of the survey area and therefore use Neumann BCs. The boundary conditions are set using predicted particle concentration data from DREAM.

3.2 Parameter specification in the SPDE model

The advection and diffusion parameters in the SPDE (1) must be specified. The velocity field is set from physical assumptions, where $\mathbf{c}_t(\mathbf{s}) = (c_e^t(\mathbf{s}), c_n^t(\mathbf{s}))^T$ as the forcing of the convection in our case is treated as the currents in the ocean. The values are set using predicted currents from the ocean model SINMOD. The isotropic diffusion D is set to be the same as for the computations in the DREAM model $D = 0.1 \text{ m}^2/\text{s}$.

The initial distribution is mainly estimated from the DREAM data. Since ocean processes are affected by currents, wind patterns, bathymetry, and freshwater runoff, we will have elevated variability in some locations. Thus, we choose a non-stationary prior model (Jun & Stein, 2008) where the mean and variance in each location are specified empirically from the training data available by DREAM:

$$\sigma_i^2 = \frac{1}{M-1} \sum_{m=1}^M (y_{i,m} - \mu_i)^2. \quad (9)$$

Here, $\mu_i = \frac{1}{M} \sum_{m=1}^M y_{i,m}$ is the empirical mean concentration in location \mathbf{s}_i over all M realizations of training data.

The prior covariance matrix $\boldsymbol{\Sigma}_0$ is then given by

$$\boldsymbol{\Sigma}_0 = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \dots & \Sigma_{1N} \\ \Sigma_{21} & \Sigma_{22} & \dots & \Sigma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{N1} & \Sigma_{N2} & \dots & \Sigma_{NN} \end{bmatrix}, \quad (10)$$

where $\Sigma_{ij} = \sigma_i \sigma_j R_{ij}$ and R_{ij} represents the spatial correlation. Hence, the diagonal of the covariance matrix in equation (10) contains the location-wise variances σ_i^2 , and the off-diagonal elements describe the covariance between the locations. The covariance function is chosen by comparing the empirical covariance of the ocean model training data with known valid parametric covariance functions. The Matern (3/2) kernel (Matérn, 2013) is fitted to such covariance function here;

$$R_{ij} = (1 + \phi h_{ij}) \exp(-\phi h_{ij}), \quad (11)$$

where $h_{ij} = |\mathbf{s}_i - \mathbf{s}_j|$ is the Euclidean distance between two locations \mathbf{s}_i and \mathbf{s}_j and $\phi > 0$ is a constant meta-parameter regulating the correlation decay with the distance. The best value for ϕ is specified using training data.

To specify covariance model parameters and visualize how the spatial correlation decays as a function of the distance between points in the domain, it is common in geostatistics to use the variogram of the data, see e.g. Cressie and Wikle (2011). The variogram works on difference $X_t(\mathbf{s}_j) - X_t(\mathbf{s}_i)$ between spatial data. Doing so, it tends to be more robust to background trends than the covariance function. The theoretical relation between the variogram $\gamma(\cdot)$ and the covariance is

$$\begin{aligned} \gamma(\mathbf{s}_i - \mathbf{s}_j) &= \frac{1}{2} \text{Var}(X_t(\mathbf{s}_i) - X_t(\mathbf{s}_j)) = \\ &= \frac{1}{2} (\text{Var}(X_t(\mathbf{s}_i)) + \text{Var}(X_t(\mathbf{s}_j)) - 2\text{Cov}(X_t(\mathbf{s}_i), X_t(\mathbf{s}_j))). \end{aligned} \quad (12)$$

Figure 3 shows results from the variogram analysis. The empirical variogram (blue) is computed by averaging the square differences of variable pairs having spatial distance within bins. We use regular bins of distances (first axis). Three different theoretical variograms were fitted to the training data. All of them fit the empirical variogram adequately, but the Matern variogram (green) was chosen here because it is closer to the empirical variogram for very small spatial distances. This indicates that a Matern spatial field has a smoothness that fits the DREAM data better. We see that the variogram increases with distance and then flattens out. Here, the effective spatial correlation length is about 500 m. For the Matern correlation function in equation (11) this means that

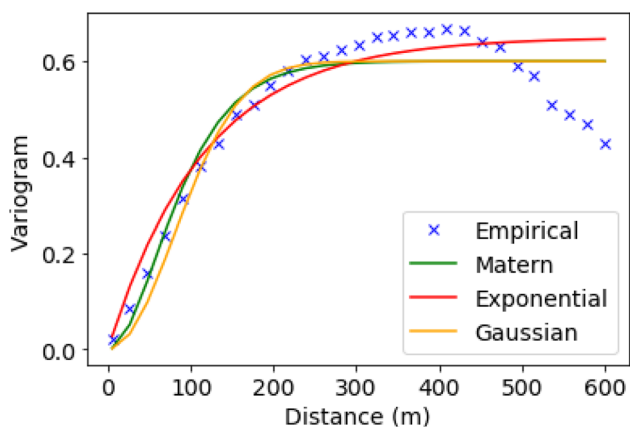


Fig. 3 The empirical variogram is shown in blue, and a Matern(3/2), an exponential, and a Gaussian variogram is fitted to the data. We find an average variance of 0.6 and an effective spatial correlation length of about 500 m (Color figure online)

$\phi \approx 0.01$. The variogram converges to the average variance which is about 0.6.

The advection–diffusion model parameters are further set from the physical assumptions in the ocean model. From visual inspection, this GP model appears to give reasonable simulations of the field but with more patchy randomness than what is seen in the ocean model. As commonly seen in physical models data, the variogram computed from the ocean models data starts to fluctuate after a while. In many cases this can be related to physical principles. Notably, the GP model is a statistical proxy model that can be run onboard but is potentially missing some large-scale oceanographic features. The ocean model is a numerical description with much skill but potentially getting systematic bias due to misspecified inputs growing with time. They both have their merits.

The covariance matrix \mathbf{W}_t in the noise term of the equation is set to have the same correlation structure as the covariance matrix for the initial state $\mathbf{\Sigma}_0$, giving $\mathbf{W}_t = V \mathbf{\Sigma}_0$, where V is a constant scaling parameter. Considering this in relation to the original continuous-time SPDE, the discretization in time would also impact this variance, but in our work, we use a fixed time discretization identical to the one used onboard. There does not seem to be any damping over time in the ocean model data, but for stability reasons (Richardson, 2017) we specify a minimal level $\lambda = -0.0010$.

4 State space model and data assimilation

The state of concentrations is assumed to be modeled by equation (8). The aim of data assimilation is to combine observations with this *prior* process knowledge to obtain an optimal sequential posterior or filtering estimate of the current state (Sar, 2013). In this study, the prior knowledge is

represented by the initial onboard model, and observations are taken from sensors on the AUV.

We define our observation model as

$$\mathbf{Y}_t = \mathbf{G}_t \mathbf{X}_t + \boldsymbol{\omega}_t, \quad \boldsymbol{\omega}_t \sim N(0, \tau^2 \mathbf{I}), \quad (13)$$

where \mathbf{Y}_t is a length m vector of observations at time t , \mathbf{G}_t is a time-varying observation matrix of size $m \times N$ that maps the state variables to the observations, and $\boldsymbol{\omega}_t$ represent the observational error which is modeled as a zero-mean Gaussian with standard deviation τ , and \mathbf{I} representing the identity matrix. Each row of \mathbf{G}_t has a one-entry on the vector position associated with an observed grid cell, and the other entries are zero. This method constraints the sampling locations to be at the predefined grid cells. Continuous path sampling, as in Choi and How (2010), could improve the sampling method and should be considered in future work.

Since the SPDE model defines a GP with a linear process representation, the spatial Kalman filter can be used to update our state at every time step. Letting $\mathbf{Z}_t = (\mathbf{Y}_1, \dots, \mathbf{Y}_t)$ represent all observations and prior information up to time t , the filtering distribution is $\mathbf{X}_t | \mathbf{Z}_t \sim N(\boldsymbol{\mu}_{t|t}, \boldsymbol{\Sigma}_{t|t})$ with mean $\boldsymbol{\mu}_{t|t}$ and variance-covariance $\boldsymbol{\Sigma}_{t|t}$ defined by

$$\begin{aligned} \boldsymbol{\Sigma}_{t|t-1} &= \mathbf{A}_t \boldsymbol{\Sigma}_{t-1|t-1} \mathbf{A}_t^T + \mathbf{W}_t \\ \mathbf{Q}_t &= \mathbf{G}_t \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t^T + \tau^2 \mathbf{I} \\ \mathbf{K}_t &= \boldsymbol{\Sigma}_{t|t-1} \mathbf{G}_t^T \mathbf{Q}_t^{-1} \\ \boldsymbol{\mu}_{t|t-1} &= \mathbf{A}_t \boldsymbol{\mu}_{t-1|t-1} + \mathbf{B}_t \\ \boldsymbol{\mu}_{t|t} &= \boldsymbol{\mu}_{t|t-1} + \mathbf{K}_t (\mathbf{Y}_t - \mathbf{G}_t \boldsymbol{\mu}_{t|t-1}) \\ \boldsymbol{\Sigma}_{t|t} &= \boldsymbol{\Sigma}_{t|t-1} - \mathbf{K}_t \mathbf{Q}_t \mathbf{K}_t^T \boldsymbol{\Sigma}_{t|t-1}. \end{aligned} \quad (14)$$

This recursion starts by the initial distribution with mean $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$, where $\mathbf{X}_0 \sim N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$. The boundary conditions and the model parameters going into the matrices and vectors in equation (14) are all specified as described in Sect. 3.2.

5 Sampling strategy

An objective function is suggested to obtain an informative path for the AUV. The function is created based on three criteria, as in Berget et al. (2019):

1. Locations with high variance are preferred
2. Locations with high predicted concentration are preferred
3. Locations leading to a suitable travel length within the next time interval for the AUV are preferred.

The first criterion is chosen because sampling in areas with a high model variance will reduce the total variance,

hence creating a more accurate model. This criterion also ensures that the AUV travels to unexplored areas. The second criterion makes the method adaptive. When studying the simulation results of the particle transport from the complex model, it is clear that the variability is highest where there is a high concentration of particles. Hence, the criterion is inspired by this observation and assumes that locations with high predicted concentration will be rich with information. The last criterion comes from the travel length limitations of the AUV. When choosing the next sample location, it is also essential that the AUV does not travel too far, assuring that the risk of the vehicle drifting outside the survey area is kept low.

The suggested objective function is then created by having a term for the first two criteria. At time step t for location s_i , the objective function is given by

$$f_t(s_i) = \theta_1 \sigma_{i,t}^2 + \theta_2 \mu_{i,t}, \quad (15)$$

where the constant parameters $\theta = [\theta_1, \theta_2]$ define the weighting for each criterion. Deciding the weighting of criteria 1 and 2 depends on design choices. When the goal is to obtain a good map of the overall region, then criterion 1 should be dominating. However, if hotspots are the main focus, criterion 2 should dominate. In our case, the parameters θ are tuned in simulation to obtain a good balance between the two criteria.

The chosen sampling location s_t^* at time step t is determined as the location that maximizes the objective function $f_t(s_i)$ for $s_i \in [s_1, \dots, s_N]$. Given the previous sampling location s_{t-1}^* , we choose the next sampling location as

$$\begin{aligned} s_t^* &= \arg \max_{s_i} f_t(s_i) \\ \text{s.t. } & |s_{t-1}^* - s_i| \geq d_{\min} \\ & |s_{t-1}^* - s_i| \leq d_{\max}. \end{aligned} \quad (16)$$

The constraints on the objective function ensure that the third criterion in the objective function is adhered to by choosing d_{\min} and d_{\max} such that a suitable travel length is obtained. $|s_{t-1}^* - s_i|$ is the Euclidean distance between the previous sampling location and s_i . Equation (16) is solved using brute-force enumeration; the value of the objective function (15) is found for all the grid cells within the region defined by the constraints in (16). The next sampling waypoint s_t^* is chosen as the grid cell with the highest objective value. In this paper, ocean currents are included in the spatio-temporal proxy model, but they are not considered in the AUV navigation plans for selecting sampling locations. Rather, the AUV is set to surface often to avoid risk and reduce navigation uncertainty, as described in the experimental section. In a larger-scale mission, it would certainly be beneficial to include more of the ocean currents in the sampling strat-

egy. Making the AUV fly with the currents may increase endurance and reduce travel time between waypoints. This topic is amongst others explored in Pereira et al. (2013) and should be considered in future work.

6 Summary of the suggested approach

We now provide an overview of the suggested approach. At the beginning of the deployment, the GP model is initialized. The following steps are executed:

- Kernel hyperparameters are estimated using prediction and hindcast data from the particle transportation model DREAM. (Described in Sect. 3.2.)
- The initial GP state defined by μ_0 and Σ_0 are found. (Described in Sect. 3.2.)
- Ocean current prediction data from SINMOD is used to define the advection fields c_t , which is used together with the diffusion constant D to define transition matrices A_t . (Described in Sect. 3.1.)
- Particle transportation data from DREAM defines the boundaries that are incorporated in the vector B_t . (Described in Sect. 3.1.)

The sampling method starts out either with the initial GP model or with the latest updated GP model. The AUV performs adaptive sampling choosing the next waypoint as described in Sect. 5. The onboard model of the AUV is then updated; the field is moved with the currents, and new information from observations is added. Updating is executed whenever a waypoint is reached, using data from observations taken along the AUV transect in the assimilation. An overview of the adaptive sampling method is given in Algorithm 1.

Algorithm 1 Adaptive Sampling method

```

Initialize GP model
while running do
  Select new waypoint using (16)
  Travel to new waypoint and collect observations
  if waypoint is reached then
    Update GP model with observations using (14)

```

7 Simulation study

A simulation study compares the proposed model with four simpler strategies. We vary both the process model and the sampling strategy. The temporal model based on the advection–diffusion SPDE is compared to a stationary pro-

Table 1 Overview of the five methods used for comparison in the simulation study

Method	Process model	Sampling strategy
SPDEObj	SPDE	Objective
SPDELawn	SPDE	Lawnmower
SPDE	SPDE	No sampling
Obj	Stationary	Objective
Lawn	Stationary	Lawnmower

Table 2 Performance measures for the five methods

Method	MAE	MSD	MO
SPDEObj	0.76	2.38	209.3
SPDELawn	0.87	2.67	272.8
SPDE	1.11	4.70	848.1
Obj	0.85	2.52	230.2
Lawn	0.99	2.67	280.5

The table shows the mean absolute error (MAE), the mean standard deviation (MSD), and the mean objective function value (MO)

process model where the next time step is equal to the previous time step, as in Berget et al. (2018). Different sampling strategies are tested with sampling strategies based on the suggested objective function, a predefined sampling strategy following a lawnmower pattern, and a method that uses only the SPDE model with no sampling (no observations are taken). The SPDE and Stationary process model are the same for all sampling strategies using the same parameter values given in the text. For the strategy using the objective function we choose the travel distance restrictions $d_{\min} = 180$ meters and $d_{\max} = 220$ meters and set the time between each sample to be 5 min. This gives a suitable travel speed for the AUV, and a constant time step simplifies the updating of the temporal model. When using the lawnmower pattern, the distance between each sampling location is set to 200 m, and the same time step of 5 min is used. The parameters in the objective function are set to $\theta_1 = 15$ and $\theta_2 = 23$. The different strategies are shown in Table 1.

Ocean model data from two consecutive days, 1st and 2nd of April 2013, is used in the simulation study. The data from April 1st is used to train the SPDE parameters as described in sections 3 and 4. Data from April 2nd is used as the test data, meaning that the predicted field on April 2nd is our ground truth and the field from which observations (samples) are drawn. We use realizations of the GP prior as the initial belief and run 100 replicate runs with the presented sampling methods in Table 1. We set the measurement error standard deviation $\tau = 0.5 \log(\text{ppb})$. We simulate for a total of 6 h, aiming to map the field on April 2nd from 15:00 to 21:00.

7.1 Results and discussion

In this section, we present results from the simulation study and discuss the performance of the methods.

7.1.1 Performance measures

To evaluate the different methods, three performance measures are studied. The mean absolute error (MAE) is computed as the mean of the absolute error between the predicted concentration $\mu_{t|t}$ and the true field from DREAM (April 2nd) for all the grid cells and time steps. The mean standard deviation (MSD) is computed similarly but then focuses on the standard deviation in the proxy model in all grid cells at all time steps. Lastly, we compute the mean objective value (see Eq. (15)) (MO) in all grid cells at every time step. Table 2 shows these performance measures computed as the mean over the total time steps. Figure 4 shows how the mean objective value (MO) (see Eq. (15)) evolves through time, plotting the mean value for every time step.

Considering the MAE values in Table 2, we see that SPDEObj performs best in this case. Including a non-stationary process model increases the method's performance, as this is the case using both the Obj sampling strategy and the Lawn sampling strategy. The objective function also increases the method's performance compared to the pre-planned lawnmower strategy. Considering the MSD and MO measures, the SPDE method stands out with significantly higher values than the other methods. This is due to the

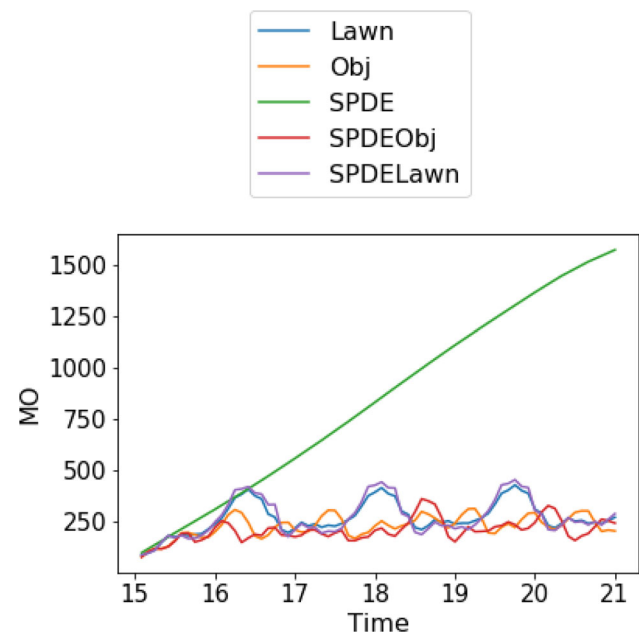


Fig. 4 The mean objective value (MO) (see Eq. (15)) computed over all the grid cells for each of the methods. Results are shown for the whole simulation time period spanning from 15:00 to 21:00




Fig. 5 Hourly plot of the model standard deviation at the given time together with the path for the different sampling strategies up until that time. The blue dot shows the path's starting point, and the yellow is the

sampling location at the current time. The arrow indicates the direction of the AUV (Color figure online)

method doing no sampling at all, and hence there is no way for this method to decrease the proxy models' variance, affecting the standard deviation and the objective values.

The effect of increasing variance with no sampling can also be seen in Fig. 4. Again, the SPDE method stands out with steadily increasing values while the other methods sta-

bilize at a mean objective value between 200 and 300. For both lawnmower methods (SPDELawn and Lawn), periodical fluctuations in the objective function can be observed. This occurs when the sampling is done close to the east end of the survey area, where the assimilation with new data affects a minor part of the modeled field compared to adding information in the center. Gener-

ally, the Obj and SPDEObj methods have a lower overall mean objective function compared with the other methods, indicating an intelligent sampling design.

In Fig. 4, we see that the Lawn method has a lower MO value at some time steps compared to the SPDELawn method. This occurs at times when both Lawn approaches get high MO. Using the lawnmower pattern for sampling (see Fig. 5a) is the reason for this unexpected behavior. Sampling is performed in boxes, and some locations in the middle of the boxes get little influence from measurements. The SPDELawn method solves this by having the temporal model move neighboring sample values into this box. The Lawn method does not have a solution for this, and the values inside this box will be similar to the initial values. Since the initial mine tailings concentration in our example (see Fig. 7a) has low values all over the area, the tailing concentration inside this box will stay low throughout the whole simulation for the Lawn method. This will affect the MO value, leading to artificially low values, even though the performance (here, better illustrated by the MAE) can be poor.

7.1.2 Sampling design and field predictions

Figure 5 shows the AUV path and waypoints for the different sampling strategies plotted on top of the standard deviation at every hour of the simulation study. The same lawnmower pattern is used for SPDELawn and Lawn and can be seen in the first column of the figure. The path chosen by the Obj and the SPDEObj method can be seen in respectively the second and third columns. All the methods cover the area quite well, but both Obj and SPDEObj focus on the eastern area. This region is closer to the outlet, and the true mine tailings distribution (seen in the first column of Fig. 7) shows elevated concentrations in this area. Hence, more frequent sampling here is expected because of the proposed objective function (see Eq. (15)).

Figure 6 gives insight into how the GP methods behave, showing the distribution in one single grid cell over time using SPDEObj. The blue line shows the method's mean value, the green line is the true value, and the two red dotted lines are the 95% confidence interval. The point we are considering is 340 ms east and 200 ms north in the survey area. When sampling and assimilation are done close to this point, we see a reduction in the confidence interval, and the mean value tends to be corrected towards the actual value.

In Fig. 7 we focus on the two methods Obj and SPDEObj. The updated model mean for every hour is plotted in columns 2 and 3, and the true field (the field where the observations are drawn from) can be seen in the first column. Comparing the methods with the true field, they capture the trend quite well. The display illustrates an important difference between the two methods; because of its stationary process model, the Obj method inherits the structure from the GP through-

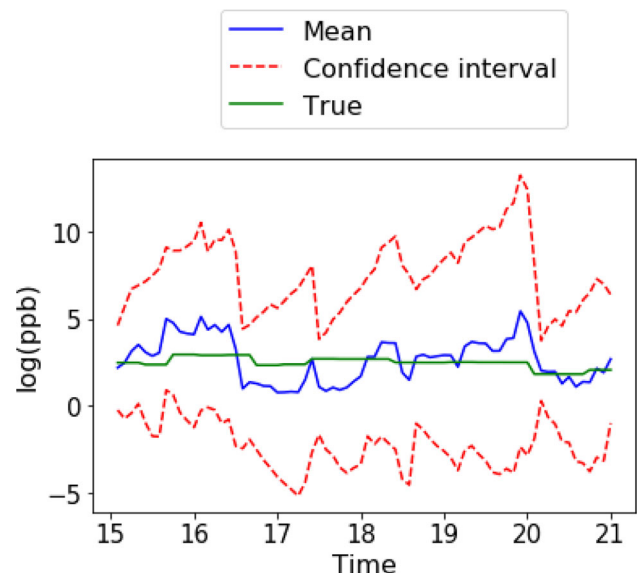


Fig. 6 The distribution in one single grid cell using the SPDEObj method. Showing mean (blue line) and 95% confidence interval (red dotted lines) together with the true value (green line). The selected grid cell is located 340ms east and 200ms north in the survey area (Color figure online)

out the whole mission, relying on observations to change its predicted field. This is not the case with SPDEObj, which has a non-stationary process model that moves the field with the currents. Although data from two different days are used as training and test data, both datasets are from the same ocean models. This may cause the comparison with ground-truth to be optimistic. Still, based on the simulation results, it is reasonable to assume that SPDEObj can map the field quite well. The MAE scores of the two methods are indicated for different times (Fig. 7). With the relatively high dynamical uncertainty, the MAE increases with time. The MAE is generally lower for the SPDEObj approach.

Simulation studies with different AUV velocities have shown that the SPDEObj method's advantages are higher when the AUV takes fewer samples in the survey area over the same time period. Naturally, lower AUV velocities led to fewer samples (comparable to having a bigger survey area), while higher velocities led to an increase in the number of samples. When doubling the number of samples, the difference in the performance between the approaches was minimal, making it hard to decide which method was best. However, when operating with half the number of samples (compared to the simulation study above), the differences became more evident, and the SPDEObj method showed improved results. Hence, it is reasonable to think that the difference between the methods will increase with a larger survey area.





Fig. 7 Hourly plot of the “true” value and the model mean value for the Obj and SPDEObj method. The MAE scores are included in the captions

7.1.3 Adaptive versus pre-planned designs


We now give an example comparing an the adaptive approach (SPDEObj) with a pre-planned method (SPDELawn) (see Fig. 8). The example shows the predicted results at the end of a mission run of total 5 h. The true field is seen in Fig. 8a. Figure 8b, c show the model mean and variance for the adaptive approach, and Fig. 8d, e present the model state for the pre-planned approach. The AUV path describing the sam-

pling design for one hour before the end time is shown in the variance plots.


In the example, the adaptive approach gives a better estimate of the true field than the pre-planned method. From the true field, we observe increased concentrations on the eastern side. The adaptive approach captures this information, and because of elevated values in the area, continues to investigate the eastern side. The pre-planned lawnmower approach, on the other hand, is set to sample in the western area at the




(a) True field after 5 hours of sampling.




(b) Predicted model mean using the adaptive approach SPDEObj.



(d) Predicted model mean using the pre-planned approach SPDELawn.



(c) Model variance of the adaptive approach SPDEObj. The AUV path one hour before the current state are shown in red.



(e) Model variance of the pre-planned approach SPDEObj. The AUV path one hour before the current state are shown in red.

Fig. 8 Example run comparing a pre-planned approach with an adaptive approach. The true field after five hours of sampling can be seen in Fig. 8a. Figure 8b, c show the model mean and variance for the adaptive

approach, and Fig. 8d, e present the model state for the pre-planned approach. The AUV path showing the sampling design for one hour before the current time is shown in the variance plots

time, and the samples taken are not as informative. Adaptive approaches excel when conditions are not as expected. In such cases, the flexibility of an adaptive approach is beneficial since the sampling design is automatically adjusted. A pre-planned approach will always stick to the pre-scripted sampling design.

8 Experimental results

As a proof of concept, field experiments were conducted in Frænfjorden. Two experiments were carried out on December 17th and 18th, 2020. The operational discharge outlet at the time of the experiments is marked as 2.1 in Fig. 1, and a rectangular operational area of size 1150m × 250m (also marked in the figure) was used. We focus on the depth layer at around 22–25 ms since this is a layer with a lot of variation and a layer where we, according to the DREAM and SIN-

MOD models, will find high concentrations of mine tailings. This is verified by our sensor data as well. Figure 11 shows one transect (waypoint-to-waypoint) for the AUV including sensor data values. It is clear from the figure that a layer with increased turbidity is reached.

8.1 Robotic platform and implementation

The robotic platform in our experiments consisted of a Light AUV from OceanScan (see Fig. 9) equipped with a Wetlabs EcoPuck sensor measuring total suspended matter (TSM), which is a measure of optical back-scattering in the form of an attenuation coefficient. The measured sensor values of TSM indicate turbidity, and this observation value was used to assimilate our model. It was converted to a measure of (log) concentration to adjust the value to our onboard model. This conversion is not straightforward, and we need to determine the uncertainty τ . The Wetlabs EcoPuck sen-



Fig. 9 The light AUV (Harald) produced by OceanScan was used in our experiments. The length of the AUV is 240 cm, and it is equipped with multiple sensors like CTD, Dissolved Oxygen, and Fluorometer

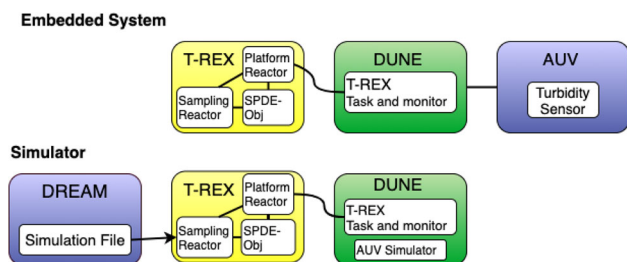


Fig. 10 The architecture of the embedded system (on top) and the simulator (bottom). The Unified Navigation Environment (DUNE) runs onboard the AUV, and the agent architecture T-REX sits on top of this, allowing the embedding of multiple decision processes such as planning

sensor user manual suggests an attenuation error of about 4 % at attenuation coefficient 1 m^{-1} . With concentrations varying between approximately 3 and 8 log(ppb), assuming measurement standard deviation of 4% gives τ between 0.12 and 0.32 log(ppb). However, both navigation errors and vehicle drift can lead to uncertainty in the sampling location. Hence, we set $\tau = 0.5 \text{ log(ppb)}$ for the experiments.

Figure 10 shows the layout of the agent architecture used in the experimental survey. The Unified Navigation Environment (DUNE) (Pinto et al., 2012) runs onboard the AUV and is used for control, navigation, vehicle supervision, communication, and interaction with actuators. On top of this sits the agent architecture T-REX (Teleo-Reactive EXecutive) (Rajan & Py, 2012), which enables an adaptive mission. T-REX allows the embedding of multiple complex decision processes (including planning). The communication between DUNE and T-REX was handled by the LSTS toolchain (Pinto et al., 2013), which provides the back-seat driver API to DUNE. This allows external controllers, such as T-REX, to provide desired tasks for our platform while receiving progress updates on the current state. Our sampling algorithm was written as a python script and was implemented as a reactor in the T-REX framework. A reactor is a component of T-REX acting as an internal control loop and is capable of producing goals that the planner integrates into

a series of actions (e.g., Goto and Arrive_at). The set of all those actions forms a plan built while ensuring operational constraints of the mission (e.g., the vehicle should never dive deeper than a certain depth or leave a defined area.) The plan is then sent to DUNE, which handles low-level control and execution. Before deploying the AUV, our method was tested in a simulated environment similar to our embedded system. In simulations, sensor readings from the AUV were replaced with data from SINMOD and DREAM, and an AUV simulator in DUNE was used to simulate the behavior of the AUV. The layout of the simulator is shown in Fig. 10.

8.2 AUV path

Inside the operational area, a 2-dimensional grid of size 58×12 was defined, giving $N = 684$. The distance between each node was approximately 20 m in both grid directions, and each grid point was a possible waypoint for the AUV. The AUV was set to surface whenever reaching a new waypoint to reduce the uncertainty caused by navigation errors and vehicle drift and for safety reasons. The AUV did a yoyo maneuver between 20 and 26 ms between each waypoint. As in the simulation study, we set $d_{\min} = 180$ meters and $d_{\max} = 220$ meters. The AUV speed is set to 1.2m/s, and as opposed to the simulation study, where a constant time of 5 min between two waypoints was used, we now let the time vary. The time used between two waypoints is found using the start and end times of the AUV transect. The temporal model and the updating are done when a waypoint is reached, updating the model with the amount of time used for that particular transect. Figure 11 shows one transect (waypoint-to-waypoint) from our survey, starting with a dive until reaching the selected depth layer, doing the yoyo maneuver, and then surfacing in a corkscrew maneuver. The corkscrew maneuver was chosen to obtain a sampling value as close as possible to the waypoint. The yoyo maneuver was chosen to increase the possibility of finding the layer of elevated mine tailings concentration. The sensor reading used for updating the model was chosen as the maximum value in the layer 22–25 ms (the layer marked with red surfaces) for each of the separate yoyos. For the transect in Fig. 11, this would result in 6 updating values. Because the transect length could vary for different waypoints, the number of yoyos was not fixed but mostly alternated between 5 or 6 yoyos per transect, giving 5 or 6 updating values per transect.

8.3 Results

The survey area is located near the coast in a fjord arm, so the currents are highly influenced by the tidal cycle. To get a better insight into how the ocean currents behave at the time of the missions, we consider the tidal cycle at the time in Fig. 12. Our temporal model considers the tides' varia-




Fig. 11 Showing the AUV transect from waypoint to waypoint, where the distance between the waypoints was 218 ms. The color of the transect indicates normalized mine tailings concentration values, and the two red surfaces show the depth layer (22–25 ms) where samples are collected (Color figure online)




Fig. 12 The blue line shows the tidal cycle on December 17th, showing low tides at around 07:00 and high tides around 13:00. The red boxes indicate the time intervals for the two missions

tions, and we choose to focus on two missions occurring at different intervals in the tidal cycle. The first mission is

before a tidal high, and the second is when the tides turn at the top. The period of the two missions is marked in red in the figure. Hourly predicted ocean currents between 09:00 and 13:00 from the SINMOD model can be seen in Fig. 13. As expected, the currents are highly influenced by the tides showing an inward flow (eastward) before the tidal high (09:00–11:00) and weaker ocean currents during the tidal high (11:00–13:00). These predictions are used as input to our onboard model as the forcing in the SPDE model.

8.3.1 Mission 1: 08:53–10:44

The first mission was conducted between 08:53 and 10:44 on December 17th, before a tidal high. The estimated current from the SINMOD model can be seen in Figs. 13a, 13b, 13c. Figure 14 presents the onboard model results. The left column shows the updated mean of the onboard model, representing the mine tailings log(concentration) for selected time steps. The right column shows the onboard model standard deviation at selected time steps and the AUV path up until the given time. The red line represents the AUV path, and the red dots are the waypoints selected by our sampling strategy. The results show increased tailing concentration in the eastern area, especially at the beginning of the mission (see Fig. 14a, 14c and 14e). We observe a decay in the concentration in this area after 1.5 h (see Fig. 14g). This is




Fig. 13 Hourly ocean current predictions from SINMOD between 09:00 and 13:00. The currents are highly influenced by the tides showing an inward flow (eastward) before the tidal high (09:00–11:00) and weaker ocean currents during the tidal high (11:00–13:00)




Fig. 14 Results from mission 1. The left column shows the updated mean of the onboard model, representing the mine tailings $\log(\text{concentration})$ for selected time steps. The right column shows the onboard model standard deviation at selected time steps and the AUV

path up until the given time. The red line represents the AUV path, and the red dots are the waypoints selected by our sampling strategy (Color figure online)

when the ocean currents get weaker due to the tidal cycle. In the beginning, the chosen sampling strategy seems to focus on exploring the area, covering large parts of the whole rectangle. After the exploration, the focus switches to the eastern area, where high mine tailing concentrations have been observed.

8.3.2 Mission 2: 11:52–13:24

The second mission was conducted during a tidal high. Figure 15 shows results from this survey. The predicted SINMOD ocean currents are much weaker than for mission 1

(see Fig. 13d and 13e). Figure 15 presents the onboard model results. The left column shows the updated mean of the onboard model, and the right column shows the standard deviation together with the selected AUV path. This mission does not contain areas of clearly increased tailings concentration, and we see that the AUV chooses a different sampling strategy. Since no area shows increased concentration, there is no extra focus on any special area. Instead, the waypoints are chosen to cover as much as possible of the whole survey area during the whole mission.

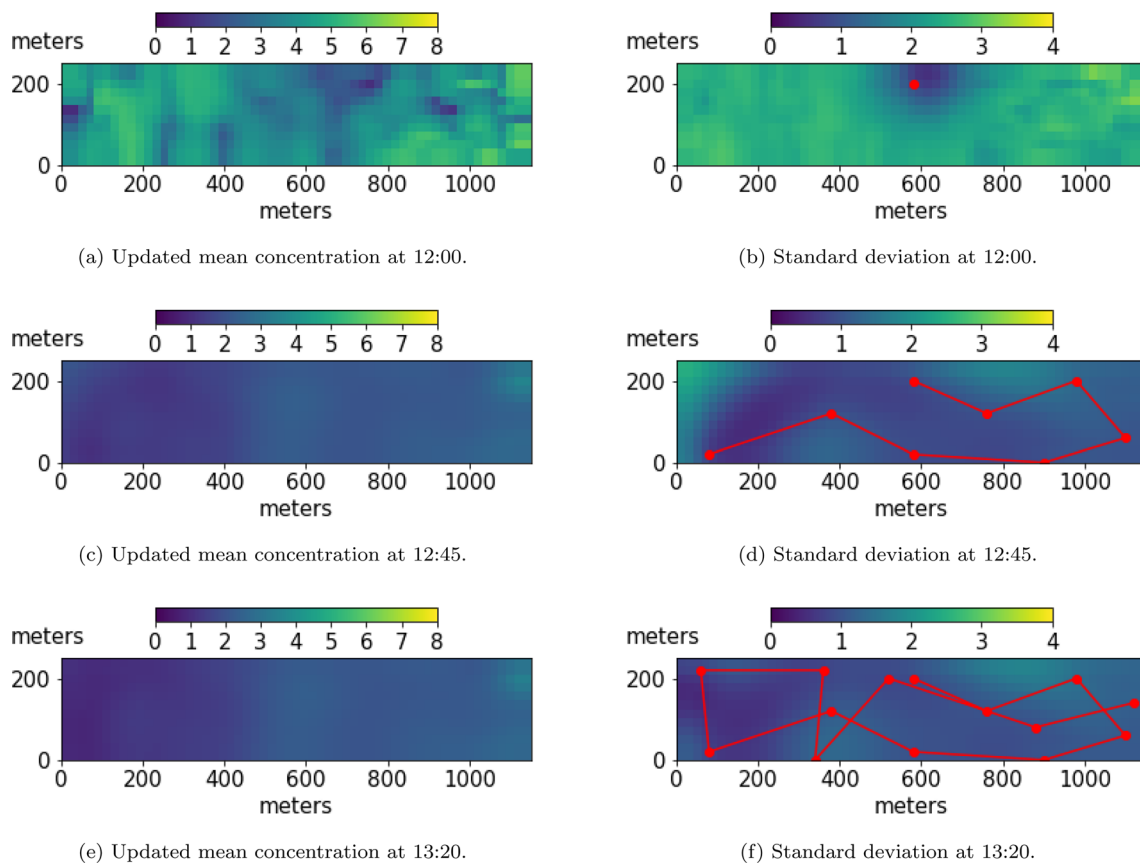


Fig. 15 Results from mission 2. The left column shows the updated mean of the onboard model, representing the mine tailings $\log(\text{concentration})$ for selected time steps. The right column shows the onboard model standard deviation at selected time steps and the AUV

path up until the given time. The red line represents the AUV path, and the red dots are the waypoints selected by our sampling strategy (Color figure online)

9 Discussion

The chosen sampling paths (see second column of Fig. 14 and 15) follow two different strategies. In the beginning, they are quite similar, focusing on exploring the area, but after the exploration, we see differences in waypoint selection. The sampling path in mission 1 (Fig. 14) focuses on the eastern area where increased concentration is observed, while the sampling path in mission 2, where there are no areas with increased tailings concentration, continues to cover the whole region also after the initial exploration. For both surveys, the standard deviation is reduced at the end of the mission, indicating that the method manages to explore the area. The selection of different strategies shows that the method is adaptive and will change depending on prior data and observations and that the sampling is focused on areas with high tailings concentration. The AUV data are undoubtedly useful for refining the DREAM results.

It is hard to rate the field experiments' performance since there is no ground truth for this case. Still, the model mean concentration (first column of Fig. 14 and 15 coincide well

with our assumptions on how the tailings concentration behave. Since the outlet is located on the east side of our area, we assume elevated concentration in the east. It is also reasonable to assume that the concentration of mine tailings in the survey area will be higher when the currents flow westward, as is the case for mission 1. We have also compared the results with the ocean model DREAM data and observed the same tendencies. However, this is not an accurate comparison since the proxy model is trained with data from this model.

10 Closing remarks

In this work, we propose a method for adaptive sampling of ocean processes with an AUV. Prior information from ocean models is combined with information from in situ measurements to obtain the best possible sampling strategy. A spatio-temporal GP proxy model with a non-stationary covariance function and a process model based on finite difference solutions of an advection–diffusion SPDE are built.

This is a low-complexity model which can run onboard the AUV. We can choose informative sampling locations by using an objective function favoring locations with high uncertainty and high predicted particle concentration. The main contribution of the paper is to have a dynamic (spatio-temporal) model onboard a dynamic agent. This enables reliable uncertainty quantification in spatial predictions while actively searching for 'hot-spots' in space and time. The method is tested in simulations and compared with simpler models and sampling strategies. A simulation study shows that the SPDE process model gives improved prediction results in a dynamic environment, and the sampling strategy using the suggested objective function performs better than the sampling strategy using a simpler lawn mower design. Since there is no easy way to obtain a ground truth for field experiments, it is difficult to determine how well the method performs in the field. Still, our results agree well with our assumptions of the mine tailings distribution at the time, and the selection of different sampling strategies for two separate missions illustrates the method's adaptivity. Future work includes expanding the model to 3D (with depth coordinates), adding current estimation from measurements onboard the AUV, including ocean currents in the waypoint selection method, and exploration using multiple AUVs. It is also interesting to test such methods for extended time periods, potentially combining the AUV measurements with satellite data or data from drones.

Funding Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital). This work is supported by the norges forskningsråd (Grand Nos. 223254, 267793)

Data availability The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Berget, G. E., Eidsvik, J., Alver, M., Py, F., Grøtli, E. I., & Johansen, T. A. (2019). Adaptive underwater robotic sampling of dispersal dynamics in the coastal ocean. In *Proceedings of the international symposium on robotics research* Accepted.
- Berget, G. E., Fossum, T., Johansen, T. A., Eidsvik, J., & Rajan, K. (2018). Adaptive sampling of ocean processes using an AUV with a gaussian proxy model. *IFAC-PapersOnLine*, 51, 238–243.
- Binney, J., Krause, A., & Sukhatme, G. S. (2013). Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8), 873–888.
- Choi, H.-L., & How, J. P. (2010). Continuous trajectory planning of mobile sensors for informative forecasting. *Automatica*, 46(8), 1266–1275.
- Cressie, N., & Wikle, C. (2011). *Statistics for Spatio-Temporal Data. CourseSmart Series*. Hoboken: Wiley.
- Das, J., Py, F., Harvey, J. B., Ryan, J. P., Gellene, A., Graham, R., Caron, D. A., Rajan, K., & Sukhatme, G. S. (2015). Data-driven robotic sampling for marine ecosystem monitoring. *The International Journal of Robotics Research*, 34(12), 1435–1452.
- Eidsvik, J., Mukerji, T., & Bhattacharjya, D. (2015). *Value of information in the earth sciences: Integrating spatial modeling and decision analysis*. Cambridge: Cambridge University Press.
- Foss, K. H., Berget, G. E., & Eidsvik, J. (2022). Using an autonomous underwater vehicle with onboard stochastic advection-diffusion models to map excursion sets of environmental variables. *Environmetrics*, e2702.
- Fossum, T., Eidsvik, J., Ellingsen, I., Alver, M., Fragoso, G., Johnsen, G., Mendes, R., Ludvigsen, M., & Rajan, K. (2018). Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics*, 35, 1101–1121.
- Gneiting, T. (2002). Nonseparable, stationary covariance functions for space-time data. *Journal of the American Statistical Association*, 97(458), 590–600.
- Griffies, S. M., Böning, C., Bryan, F. O., Chassignet, E. P., Gerdes, R., Hasumi, H., Hirst, A., Treguer, A.-M., & Webb, D. (2000). Developments in ocean climate modelling. *Ocean Modelling*, 2, 123–192.
- Halpern, B. S., Walbridge, S., Selkoe, K. A., Kappel, C. V., Micheli, F., D'Agrosa, C., Bruno, J. F., Casey, K. S., Ebert, C., Fox, H. E., Fujita, R., Heinemann, D., Lenihan, H. S., Madin, E. M. P., Perry, M. T., Selig, E. R., Spalding, M., Steneck, R., & Watson, R. (2008). A global map of human impact on marine ecosystems. *Science*, 319(5865), 948–952.
- Hwang, J., Bose, N., & Fan, S. (2019). AUV adaptive sampling methods: A review. *Applied Sciences*, 9(15).
- Hyun, J. W., Li, Y., Huang, C., Styner, M., Lin, W., & Zhu, H. (2016). STGP: Spatio-temporal gaussian process models for longitudinal neuroimaging data. *NeuroImage*, 134, 550–562.
- Jennison, C., & Turnbull, B. W. (1999). *Group sequential methods with applications to clinical trials*. Boca Raton: CRC Press.
- Jun, M., & Stein, M. L. (2008). Nonstationary covariance models for global data. *The Annals of Applied Statistics*, 2(4), 1271–1289.
- Krause, A., Singh, A., & Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning and Research*, 9, 235–284.
- Kvassnes, A., & Iversen, E. (2013). Waste sites from mines in Norwegian fjords. *Mineralproduksjon*, 3, A27–A38.
- Lermusiaux, P., Chiu, C.-S., Gawarkiewicz, G., Abbot, P., Robinson, A., Miller, R., Haley, P., Leslie, W., Majumdar, S., Pang, A., & Lekien, F. (2006). Quantifying uncertainties in ocean predictions. *Oceanography*, 19, 92–105.

- Luo, W. & Sycara, K. (2018). Adaptive sampling and online learning in multi-robot sensor coverage with mixture of gaussian processes. In *2018 IEEE international conference on robotics and automation (ICRA)* (pp. 6359–6364).
- Ma, K.-C., Liu, L., Heidarsson, H., & Sukhatme, G. (2018). Data-driven learning and planning for environmental sampling. *Journal of Field Robotics*, *35*(5), 643–661.
- Matérn, B. (2013). Spatial variation. Meddelanden från Statens Skogsforskningsinstitut, *36*(5), 1–144.
- Morello, E., Haywood, M., Brewer, D., Apte, S., Asmund, G., Kwong, Y., & Dennis, D. (2016). The ecological impacts of submarine tailings placement. In Hughes, R., Hughes, D., Smith, I. P., & Dale, A. (eds.), *Oceanography and Marine Biology: An Annual Review*, volume 54 (pp. 315–366). CRC Press, 1 edition.
- Nepstad, R., Liste, M., Alver, M. O., Nordam, T., Davies, E., & Glette, T. (2020). High-resolution numerical modelling of a marine mine tailings discharge in Western Norway. *Regional Studies in Marine Science*, *39*, 101404.
- Pereira, A., Binney, J., Hollinger, G., & Sukhatme, G. (2013). Risk-aware path planning for autonomous underwater vehicles using predictive ocean models. *Journal of Field Robotics*, *30*(5), 741–762.
- Pinto, J., Calado, P., Braga, J., Dias, P., Martins, R., Marques, E., & Sousa, J. (2012). Implementation of a control architecture for networked vehicle systems. *IFAC Proceedings Volumes*, *45*(5), 100–105.
- Pinto, J., Dias, P., Martins, R., Fortuna, J., Marques, E., & Sousa, J. (2013). The LSTS toolchain for networked vehicle systems. In *OCEANS 2013 MTS/IEEE Bergen: The Challenges of the Northern Dimension*.
- Rajan, K. & Py, F. (2012). T-REX: Partitioned inference for AUV mission control. *Further advances in unmanned marine vehicles* (pp. 171–199).
- Ramírez-Llodra, E., Trannum, H., Evenset, A., Levin, L., Andersson, M., Finne, T., Hilário, A., Flem, B., Christensen, G., Schaanning, M., & Vanreusel, A. (2015). Submarine and deep-sea mine tailing placements: A review of current practices, environmental issues, natural analogs and knowledge gaps in Norway and internationally. *Marine Pollution Bulletin*, *97*(1), 13–35.
- Rasmussen, C. E., & Williams, C. K. I. (2005). *Gaussian processes for machine learning (adaptive computation and machine learning)*. Cambridge: The MIT Press.
- Richardson, R. A. (2017). Sparsity in nonlinear dynamic spatiotemporal models using implied advection. *Environmetrics*, *28*(6), e2456.
- Rye, H., Reed, M., & Ekrol, N. (1998). The PARTRACK model for calculation of the spreading and deposition of drilling mud, chemicals and drill cuttings. *Environmental Modelling Software*, *13*(5), 431–441.
- Rye, H., Reed, M., Frost, T., Smit, M., Durgut, I., Johansen, O., & Ditlevsen, M. (2008). Development of a numerical model for calculating exposure to toxic and nontoxic stressors in the water column and sediment from drilling discharges. *Integrated Environmental Assessment and Management*, *4*, 194–203.
- Särkkä, S. (2013). *Bayesian filtering and smoothing*. Cambridge University Press: Cambridge.
- Sigrist, F., Künsch, H., & Stahel, W. (2015). Stochastic partial differential equation based modeling of large space-time data sets. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *77*, 3–33.
- Slagstad, D., & McClimans, T. (2005). Modeling the ecosystem dynamics of the barents sea including the marginal ice zone: I. Physical and chemical oceanography. *Journal of Marine Systems*, *58*, 1–18.
- Stewart, R. H. (2008). *Introduction To physical oceanography*. Texas A & M University.
- Storvik, G., Frigessi, A., & Hirst, D. (2002). Stationary space-time Gaussian fields and their time autoregressive representation. *Statistical Modelling*, *2*(2), 139–161.
- Trannum, H. C., Nilsson, H. C., Schaanning, M. T., & Øxnevad, S. (2010). Effects of sedimentation from water-based drill cuttings and natural sediment on benthic macrofaunal community structure and ecosystem processes. *Journal of Experimental Marine Biology and Ecology*, *383*(2), 111–121.
- Wassmann, P., Slagstad, D., Riser, C. W., & Reigstad, M. (2006). Modelling the ecosystem dynamics of the barents sea including the marginal ice zone: II. Carbon flux and interannual variability. *Journal of Marine Systems*, *59*(1): 1–24.
- Zhang, Y., Godin, M. A., Bellingham, J. G., & Ryan, J. P. (2012). Using an autonomous underwater vehicle to track a coastal upwelling front. *IEEE Journal of Oceanic Engineering*, *37*(3), 338–347.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Gunhild Elisabeth Berget is a Ph.D. Student in Engineering Cybernetics at The Norwegian University of Science and Technology (NTNU). She received her M.Sc. degree in Applied Mathematics in 2017 from NTNU. Her Ph.D. project is related to adaptive sampling using AUVs focusing on modeling, data assimilation and path planning.



Jo Eidsvik is Professor of Statistics in the Department of Mathematical Sciences at the Norwegian University of Science and Technology (NTNU). He got his M.Sc. (1997) in Applied Mathematics from the University of Oslo and his Ph.D. (2003) in Statistics from NTNU. His research interest are spatiotemporal statistics and computational statistics, often applied to the Earth sciences. Along with methods for parameter estimation, prediction and data assimilation in these settings, he also works with questions related to the design of experiments, sequential uncertainty reduction and value of information analysis in this context of spatiotemporal applications. He co-authored a book on 'Value of information in the Earth Sciences', Cambridge University press, 2015. He is Associate Editor of Statistics and Computing and Mathematical Geosciences.



Morten Omholt Alver received the M.Sc. and Ph.D. in Engineering Cybernetics from The Norwegian University of Science and Technology (NTNU) in 2002 and 2007, respectively. He is currently an Associate Professor at the Department of Engineering Cybernetics at NTNU. His research interests include ocean modelling and data assimilation, modelling of biological systems, and instrumentation and automation in marine aquaculture.



Tor Arne Johansen received the M.Sc. degree in 1989 and the Ph.D. degree in 1994, both in electrical and computer engineering, from the Norwegian University of Science and Technology, Trondheim, Norway. From 1995 to 1997, he worked at SINTEF as a researcher before he was appointed Associated Professor at the Norwegian University of Science and Technology in Trondheim in 1997 and Professor in 2001. He has published several hundred articles in the areas of

control, estimation and optimization with applications in the marine, aerospace, automotive, biomedical and process industries. In 2002 Johansen co-founded the company Marine Cybernetics AS where he was Vice President until 2008. Prof. Johansen received the 2006 Arch T. Colwell Merit Award of the SAE, and is currently a principal researcher within the Center of Excellence on Autonomous Marine Operations and Systems (NTNU-AMOS) and director of the Unmanned Aerial Vehicle Laboratory at NTNU and the SmallSat Laboratory at NTNU. He recently co-founded the spin-off companies Scout Drone Inspection, UBIQ Aerospace, Zeabuz and SentiSystems.