Adaptive Underwater Robotic Sampling of Dispersal Dynamics in the Coastal Ocean

Gunhild Elisabeth Berget[1,2], Jo Eidsvik[3], Morten Omholt Alver[1], Frédéric Py[4], Esten Ingar Grøtli[4], and Tor Arne Johansen[1,2]

[1] Department of Engineering Cybernetics,
Norwegian University of Science and Technology,
Trondheim, Norway
[2] Centre for Autonomous Marine Operations and Systems,
Norwegian University of Science and Technology,
Trondheim, Norway
[3] Department of Mathematical Sciences,
Norwegian University of Science and Technology,
Trondheim, Norway
[4] Sintef Digital, Trondheim, Norway

Abstract. To get a better understanding of the highly nonlinear processes driving the ocean, efficient and informative sampling is critical. By combining robotic sampling with ocean models we are able to choose informative sampling sites and adaptively change our path based on measurements. We present models exploiting prior information from ocean models as well as real-time information from in situ measurements. The method uses compact Gaussian process modeling and objective functions to locate informative sampling sites. Our aim is to get a better understanding of ocean processes and improve real-time monitoring of dispersal dynamics. The case study focuses on a fjord located in Norway containing a seafill for mine tailings. Transportation of the deposited particles are studied, and the sampling method is tested in the area. The results from these sea trials are presented.

Keywords: Adaptive sampling, Gaussian processes, AUV, Oceanography

1 Introduction

There has been a steady increase in both knowledge and awareness concerning the environmental impact and potentially harmful effects of the discharge of drill cuttings, mud and mine tailings in the sea. Cold-water coral reefs and sponges, found in abundance along the Norwegian continental shelf, are two examples of sensitive species that may be negatively impacted by such sea floor deposition [Trannum et al., 2010]. This prompts the need for both dispersion modelling in the planning phase, and monitoring of the sedimentation, turbidity and other variables in the execution phase, and is a motivation for developing information-driven sampling methods of particle dispersal, which is the focus of this paper.

Creating models describing the ocean dynamics is challenging because of its large scale nonlinear processes and high spatio-temporal variability. Existing models continuously refine numerical methods towards improving accuracy, [Griffies et al., 2000]. But still the existing models are prone to errors; simplifications are done, errors can occur because of the numerical implementation and so on. Model verification and data assimilation continues to be a challenge that prompts the need for in situ measurements. Such data is commonly obtained using either remote sensing, ships or buoys. This data is usually expensive to acquire and process, making it hard to observe the entire environment in detail. Hence, the ocean tends to be *undersampled*, leading to large gaps in our understanding of the ocean [Stewart, 2008]. Thus, when designing an experiment or observation system, we need to address the problem of when and where to sample. By thoroughly planning a mission, using all information available, sampling locations can be chosen such that they maximize the information retrieval during a survey. Planning should therefore exploit prior information from e.g. ocean models, satellite data, stationary buoys and data from previous missions. In addition, real-time in situ measurements can be used to adapt an ongoing mission such that a best possible sampling strategy is obtained.

Mobile robotic platforms, like autonomous underwater vehicles (AUVs), have recently become more affordable, robust and viable for scientific exploration, thus providing an efficient platform for autonomous collection of in situ oceanographic data. In this paper we focus on a method using an AUV for sampling in situ turbidity data with the goal of tracking suspended material plumes, and being able to adjust the mission in real-time. To obtain real-time adaption, a faster-than-real-time particle dynamics model onboard the AUV is required. The numerical ocean models have a high computational load, making them unfit for running on embedded robotic systems with both data processing and storage constraints. Hence, a simpler, more compact model approximating the processes is built based on Gaussian processes (GP). This simplified low-complexity proxy model represents the state of the ocean at the time, and can be updated when new information is available. Two existing numerical models, SINMOD (ocean model) and DREAM (particle dynamics model), are used to train the GP proxy model creating a prior proxy model of the particle concentration. Sensor readings on the AUV can then be used to update the proxy model onboard in real-time. To maximize the value of information from the samples from the AUV, an objective function for path planning is presented. The objective function assures that the area is explored by choosing locations assumed to be information rich, and also considers the limitations of the AUV.

As a case study, the proposed algorithm is tested in Frænfjorden (Norway) which contains a seafill for submarine mine tailings. The goal is to track the particle dispersal near this seafill, aiming to improve real-time monitoring of dispersal dynamics.

The paper is organized as follows. Section 2 provides definitions and background information on modeling, methods and data assimilation. Section 3 explains our approach and the implementation of the developed methods. Section

4 shows the results from the field experiment. Lastly, section 5 concludes with a summary discussion and future work.

1.1 Related work

GPs [Cressie and Wikle, 2011],[Eidsvik et al., 2015] are powerful for creating non-parametric, simple and computationally efficient models. They are widely used when creating a data-driven spatial model, and are popular within environmental sensing. This is among others explored in [Krause et al., 2008], where a method for static sensor placement is suggested using GPs and maximization of mutual information. Others use GPs in combination with robotic vehicles as in [Zhang et al., 2012] where an AUV is used to track an upwelling front or in [Das et al., 2015] which use an AUV to collect samples for ex-situ analysis, selecting the sampling locations based on previous missions and maximizing a utility function.

When introducing robotic vehicles for sampling, path planning is required to obtain the optimal sampling path. Finding the optimal path is among others discussed in [Binney et al., 2010] which use the measure of mutual information to optimize information gain along a 2D path for a marine glider. This is further elaborated and tested with a surface vehicle in [Binney et al., 2013], where a comparison of greedy vs. recursive greedy approaches is explored for a similar problem.

A common approach when building a GP model is to assume stationary variance, but when modeling particle transportation there is reason to believe that some sites vary more than others. In this paper an approach using non-stationary variance is suggested, using empirical variance from numerical models as training data for the model variance. A similar approach is explored in [Fossum et al., 2018], which uses temperature as an information utility and trains a GP model using the ocean model SINMOD with the goal of tracking an ocean front. Our paper is based on the ideas discussed in [Berget et al., 2018], where a low-complexity GP model of particle concentration is built and the sampling algorithm is tested with simulated data. This paper presents results of this method using experimental oceanographic data from a sea trial.

2 Background

This section presents definitions and background information on modeling and data assimilation.

2.1 Ocean models

In this paper, we use numerical oceanographic models to build a prior belief of the state of the ocean. Ocean models are models that describe the state of the ocean at a given time, providing information on temperature, salinity, currents, density and pressure. The models are based on a set of thermodynamic and hydrodynamic equations, commonly called the primitive equations, and these are

solved using numerical techniques. Running of these models is a computer intensive task, since it involves solving a large set of equations. Thus, we get a trade-off between the resolution of the model and the available computer resources. Because of these limitations a high-resolution model can only be computed for small areas. A common technique is nesting, where the model is run at a large scale producing boundary conditions for smaller scale models with higher resolution.

Input to ocean models typically includes tides, sea-level pressure, wind, heat exchange, bathymetry and freshwater runoff. Errors in the input data as well as numerical errors can affect the quality of the output data. Hence, there is a need to develop enabling technology that performs efficient and targeted sampling of the ocean. Observations from different platforms such as buoys, ship-based sampling or robotic sampling from for example AUVs can be used to evaluate the performance of the models. In addition to such hindcast model validation and correction, the information can be used in real-time to improve the ocean model using data assimilation.

In the case study, we use two specific models to build our prior belief, SINMOD and DREAM. SINMOD describes the ocean dynamics and data from this model is used as input to DREAM, which is our particle transportation model. These are typical examples of ocean models.

SINMOD - The ocean model SINMOD is a numerical ocean model system that connects and simulates physical and biological processes in the ocean [Slagstad and McClimans, 2005], [Wassmann et al., 2006]. 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 includes atmospheric forcing, freshwater run-off, and boundary conditions (density, tidal forcing, currents).

In our experiments, SINMOD has been set up with 32 m horizontal resolution. For input, the bathymetry data is based on DMB Sør-Norge, supplemented by OLEX data recorded by SINTEF Materials and Chemistry. The atmospheric input data is produced using the Weather Research and Forecasting (WRF) (https://www.mmm.ucar.edu/weather-research-and-forecasting-model) model simulated with boundary values from the ERA - Interim reanalysis, and climatologic data for freshwater run-off is used.

DREAM - The particle transportation model DREAM is a Lagrangian particle transport model which can be used to simulate behaviour and fate of marine pollutants, including particulate discharges from drilling operations [Rye et al., 1998, Rye et al., 2008]. It provides time series of concentrations of released materials in the water column, as well as deposition of these materials onto the sea floor. Input to the DREAM model includes hydrodynamic data, in our case delivered by SINMOD, as well as information about the release (amount, rate, densities, grain size distribution). DREAM is often used as a decision support tool for management of operational discharges to the marine environment.

For Frænfjorden DREAM was set up to use input data from the ocean model SINMOD. The horizontal resolution was set to 38 m.

2.2 Low complexity data-driven spatial models

To be able to adjust a mission based on real-time measurements, a low complexity data-driven model that can be updated relatively quickly is needed. Since numerical ocean models, like SINMOD, are computer intensive, it is currently not feasible to run these onboard a robotic platform. Still, it is desirable to have a representation of the spatial conditions, and a common solution is to use a stochastic proxy model based on GPs.

A GP is a collection of random variables having a multivariate normal probability density function. When the variables are allocated to spatial locations, the spatial dependencies can be modeled through the covariance of the density function. A definition of a GP representing a 2-D spatial process is given by:

Consider a real-valued stochastic process $\{X(s), s \in \Omega\}$, where $\Omega \subset \mathbb{R}^2$ is a set of locations. This is a GP if, for any finite choice of N distinct locations $s_1, \ldots, s_N \in \Omega$, the random vector $\boldsymbol{x} = [x(s_1), \ldots, x(s_N)]$ has a multivariate normal probability density function:

$$p(\boldsymbol{x}) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{\frac{N}{2}} |\boldsymbol{\Sigma}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right), \qquad (1)$$

defined by the mean vector $\boldsymbol{\mu} = \mathrm{E}(\boldsymbol{x})$, and the symmetric positive definite covariance matrix $\boldsymbol{\Sigma} = \mathrm{cov}(\boldsymbol{x}, \boldsymbol{x})$.

A GP is a popular stochastic model, especially for environmental processes, and this is often attributed to two essential properties. First, as can be seen from equation (1), a GP is fully expressed by its mean and covariance. Thus, only the first- and second-order moments need to be specified when building the model. Second, the procedure for prediction and assimilation is uncomplicated since it is inherent to the fundamental equations of the model.

In this work, a GP is used to represent the spatial dependencies of the particle concentration in the ocean, and hindcast and prediction data from the numerical model DREAM is used to specify the GP model parameters.

3 Methods

In this section we present the onboard spatial GP model, the sampling algorithm used in our survey and the implementation of the algorithm.

3.1 Gaussian process specification

In this paper only one depth layer within our survey area is considered, although the concepts are general and can be extended to 3D and also temporal dimension. Our 2-dimensional domain is divided into a regular grid with N grid points

$[s_1, \ldots, s_N]$, and the particle concentration in location s_i is assumed to be Gaussian with mean μ_i and variance σ_i^2. Hindcast data and prediction data from the DREAM model is used to specify the parameters (the mean and the covariance) in our GP model. The prior mean particle concentration $\boldsymbol{\mu}_0 = [\mu_1, \ldots, \mu_N]$ is approximated by computing the empirical mean concentration μ_i^* in each location s_i, $i = 1, \ldots, N$. Assuming M data $[y_{i,1}^*, \ldots, y_{i,M}^*]$ in location s_i, this is given by

$$\mu_i^* = \frac{1}{M} \sum_{m=1}^{M} y_{i,m}^*, \tag{2}$$

and the prior mean of the proxy model is obtained as the vector $\boldsymbol{\mu}_0 = [\mu_1^*, \ldots, \mu_N^*]^T$. As the GP is specified only in two dimensions, the mean values constitute a 2D particle concentration surface.

The prior covariance matrix $\boldsymbol{\Sigma_0}$ is given as

$$\boldsymbol{\Sigma_0} = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \ldots & \Sigma_{1N} \\ \Sigma_{21} & \Sigma_{22} & \ldots & \Sigma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{N1} & \Sigma_{N2} & \ldots & \Sigma_{NN} \end{bmatrix}, \tag{3}$$

where $\Sigma_{ij} = \sigma_i \sigma_j R_{ij}$ and R_{ij} is the correlation function. Hence, the diagonal of the covariance matrix contains the variances σ_i^2, and the off-diagonal elements describe the covariance between the locations. The fundamental concept of modelling spatial correlation needs to fulfill two main properties: i) that correlation decays with distance and ii) that the covariance matrix is positive definite. To achieve this, it is common to use generic correlation functions or kernels. By comparing covariance functions with the empirical covariance of the training data, Matern (3/2) kernel [Matérn, 2013] is chosen. The function is given by

$$R_{ij} = (1 + \phi h_{ij}) \exp(-\phi h_{ij}), \tag{4}$$

where $h_{ij} = |s_i - s_j|$ is the Euclidian distance between two locations s_i and s_j and $\phi > 0$ is a constant meta-parameter regulating the correlation decay with the distance. The best value for ϕ is estimated using training data by choosing the best fit of the covariance function to the empirical covariance in the data.

When modeling ocean processes, factors such as bathymetry, currents, wind patterns, and freshwater run-off in coastal areas imply that some locations will have elevated variability. Thus, we choose a non-stationary model where the prior variance of the state in each location is chosen to be the empirical variance from the DREAM data [Jun and Stein, 2008]

$$\sigma_i^{*2} = \frac{1}{M-1} \sum_{m=1}^{M} (y_{i,m}^* - \mu_i^*)^2. \tag{5}$$

To model the temporal changes, a first order Markovian process is suggested

$$\boldsymbol{x}_t = \boldsymbol{x}_{t-1} + \boldsymbol{q}_t, \tag{6}$$

where $\boldsymbol{q}_t \sim N_N(0, V\boldsymbol{\Sigma}_0)$ is a N-dimensional normally distributed vector with zero mean and covariance matrix $V\boldsymbol{\Sigma}_0$ where $V > 0$ is a constant parameter. This temporal model assumes that the current step in time is similar to the previous with an increase in variance proportional to the prior covariance matrix $\boldsymbol{\Sigma}_0$. In this way, parts of the spatial correlation between the locations is maintained, and the increase in variance due to the dynamics of the particle transportation is modeled. The constant value V determines the size of the increase in variance, and this value was tuned to fit the modelled domain.

This is a simplified temporal model that only increases the variance in each time step, and does not consider the temporal dispersal dynamics. Hence, we rely on the observations from the AUV to catch the changes in the particle distribution.

3.2 Data assimilation

The observation model is given by

$$\boldsymbol{y}_t = \boldsymbol{G}_t\boldsymbol{x}_t + \boldsymbol{\epsilon}_t. \tag{7}$$

Here, $\boldsymbol{G}_t$ is the sampling design at time t, a matrix of size $(M \times N)$ that contains 1 at the entries that correspond to the sampled locations at time t, and 0 otherwise. $\boldsymbol{\epsilon}_t \sim N_M(0, \boldsymbol{\Omega})$ is a normally distributed error term with zero-mean and covariance $\boldsymbol{\Omega}$, assumed to be Gaussian, describing the measurement noise. In our experiments $M = 1$ and $\boldsymbol{G}_t$ will be a vector of size N.

Since a GP is fully represented by its mean and covariance matrix, these are the only parameters that needs to be updated in each time step. Exploiting the properties of the Gaussian distribution, the conditional updated mean and covariance matrix at time step t: $\boldsymbol{\mu}_t = E(\boldsymbol{x}_t|\boldsymbol{y}_1, \ldots, \boldsymbol{y}_t)$ and $\boldsymbol{\Sigma}_t = \mathrm{cov}(\boldsymbol{x}_t|\boldsymbol{y}_1, \ldots, \boldsymbol{y}_t)$ can be found by [Rasmussen and Williams, 2005], letting $\boldsymbol{\Sigma}_t^*$ represent the increased covariance after applying the temporal model in (6)

$$\begin{aligned}
\boldsymbol{K}_t &= \boldsymbol{\Sigma}_{t-1}^*\boldsymbol{G}_t^T(\boldsymbol{G}_t\boldsymbol{\Sigma}_{t-1}^*\boldsymbol{G}_t^T + \boldsymbol{\Omega})^{-1} \\
\boldsymbol{\mu}_t &= \boldsymbol{\mu}_{t-1} + \boldsymbol{K}_t(\boldsymbol{y}_t - \boldsymbol{G}_t\boldsymbol{\mu}_{t-1}) \\
\boldsymbol{\Sigma}_t &= \boldsymbol{\Sigma}_{t-1}^* - \boldsymbol{K}_t\boldsymbol{G}_t\boldsymbol{\Sigma}_{t-1}^* \\
\boldsymbol{\Sigma}_t^* &= \boldsymbol{\Sigma}_t + V\boldsymbol{\Sigma}_0.
\end{aligned} \tag{8}$$

3.3 Objective function

To obtain an informative path for the AUV, an objective function is suggested. The function is created based on three criteria, as in [Berget et al., 2018]

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

The first criterion is chosen because observing in areas with high variance leads to a reduction in total variance, hence creating a more accurate model. This criterion also ensures that the AUV travels to areas that are unexplored. 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 each of two first criteria. At time step t for location s_i, the objective function is given by

$$f_t(s_i) = \theta_1 \sigma^2_{i,t} + \theta_2 \mu_{i,t} \tag{9}$$

where the constant parameters $\boldsymbol{\theta} = [\theta_1, \theta_2]$ defines the weighting for each criteria. These parameters together with the parameter V in the updating equations (8) are tuned by trial and error in simulation to obtain the desired behavior of the AUV. Especially, the tuning of V is important to maintain the balance between the terms in the objective function (9). As seen from equation (8), the assimilation of data results in decreased covariance near the observation sites. Thus, a small V would result in low covariance values as more observations are taken, making the second term in the objective function dominating. On the other hand, a large V would result in the covariance growing out of bounds which would make the first term dominating.

3.4 The algorithm

The sampling location S_t at time step t is chosen as the location that maximizes the objective function $f_t(s_i)$ for $s_i \in [s_1, \ldots, s_N]$. So given the previous sampling location S_{t-1}, we choose the next sampling location as

$$S_t = \arg\max_{s_i} \quad f_t(s_i) \tag{10}$$

$$\begin{aligned} \text{s.t.} \quad & |S_{t-1} - s_i| \geq d_{\min} \\ & |S_{t-1} - s_i| \leq d_{\max}. \end{aligned} \tag{11}$$

The constraints on the objective function ensures 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 Euclidian distance between the previous sampling location and s_i. Details of the sampling method are given in Algorithm 1.

This is a greedy method that chooses the best next sampling location at each step. First, the low-complexity GP model is initialized using training data from the numerical particle transportation model DREAM. The sampling then begins by evaluating the objective function and choosing the location which maximizes

Algorithm 1 Sampling method

```
procedure SAMPLING
    Initialize GP according to (2) and (3)
    for t = 1, ..., T do
        for s = s_1, ..., s_N do
            Evaluate the objective function f_t(s) (9)
        Choose next sampling location S_t according to (10)
        Go to location S_t
        Retrieve observations y_t from S_t
        Assimilate data according to (8)
```

it. After reaching the desired location and observing there, data assimilation is done and the GP model is updated such that the variance is increased in the unobserved locations.

3.5 Implementation

Figure 1 shows the layout of the agent architecture used in the experimental survey. The Unified Navigation Environment (DUNE) [Pinto et al., 2012] is running 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 and 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 also receiving progress updates on the current state.

Our sampling algorithm was written as python-code, and was implemented as a reactor in T-REX. A reactor is a component of T-REX acting as an internal control loop in the framework, and is capable of producing goals that the planner integrates in to a series of actions (e.g., Goto, Arrive_at and so on). The set of all those actions forms a plan which was 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 and the parameters in the objective function was tuned in a simulated environment as similar as possible to our embedded system. The layout of the simulator is shown in Figure 1. In simulations, sensor readings from the AUV was replaced with data from SINMOD and DREAM, and an AUV simulator in DUNE was used to simulate the behaviour of the AUV.

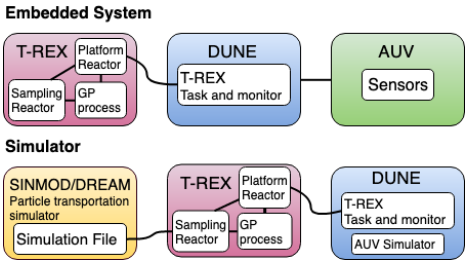

Fig. 1: Block diagram showing the layout of both the embedded system and the simulator environment.

4 Field experiments

The aim of the field experiment is to verify the sampling algorithm's ability to construct an informative adaptive mission based on in situ observations. The experiment was carried out on October 19, 2018 in Frænfjorden, Norway. Figure 2 shows the chosen operational area together with the executed adaptive AUV path. The location of the outlet and the location of a stationary buoy equipped with a CTD sensor and a turbidity sensor is also marked in the figure.

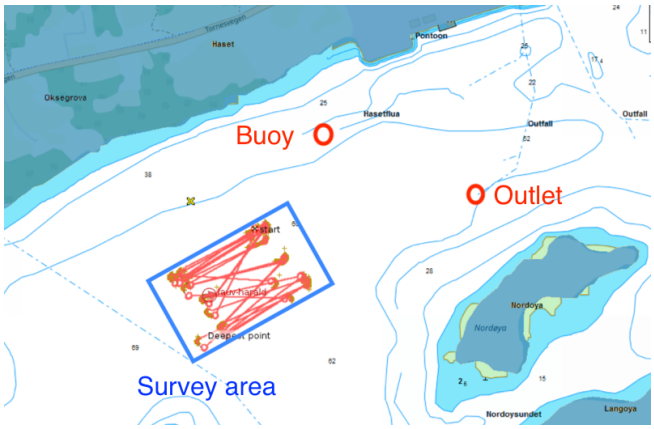

Fig. 2: Map of the chosen operational area. The location of the stationary buoy and the location of the outlet is marked in red in the figure, and the path of the AUV is drawn as a red line inside the survey area.

4.1 Experimental setup

Frænfjorden is a fjord with shallow water, fishing nets and boat traffic, and is therefore a risky place for an AUV mission. This led to limitations when choosing

the operational area. Considering the risk, and also assuring that the survey area was fairly close to the outlet of particles, we ended up with the fairly small area (550 m $\times$ 250 m) shown in Figure 2. The waypoint grid was chosen inside this area, so that the AUV would not operate outside. A drifting margin of 50 m was allowed. The grid cell size was set to 32 $\times$ 32 meters, and the parameters constraining the travel length of the AUV (see Eq. (11)) was set to $d_{\min} = 250$ m and $d_{\max} = 400$ m.

The implemented algorithm only considers one depth layer, and the layer used in the mission (approximately 22-27 meters depth) was chosen as the layer with most variation in the simulation data from DREAM. A mission was run for 110 minutes, starting 8 am on 19th October 2018.

Fig. 3: NTNU's Light AUV platform (Harald) used in our work.

Our robotic platform (see Figure 3) consisted of a Light AUV from Ocean-Scan equipped with a Wetlabs EcoPuck sensor measuring chlorophyll *a* concentration, color dissolved organic matter (cDOM) and total suspended matter (TSM). The sensor reading used for our algorithm was the TSM. Due to lack of calibration of the TSM and turbidity data with the actual sediment concentrations from the site, a direct comparison is not possible. However, we choose to compare relative increases in the signals and assume a linear relation between the values. In this way, we can compare and evaluate the model.

In order to avoid vehicle drift outside of the safe operation area, the AUV was surfaced every time it reached a waypoint. When diving, the AUV followed a Yoyo path which can be seen in Figure 4 showing the path in depth direction for the first 20 minutes of the mission. The sensor value used in the algorithm was chosen as the maximum TSM value in a span of 10 seconds near the waypoint.

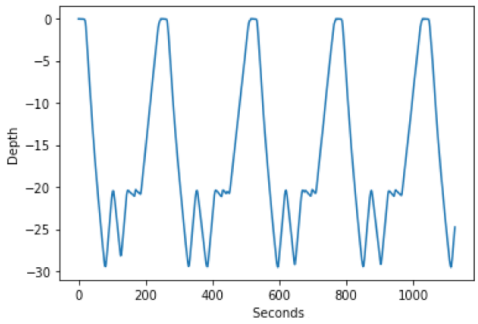

Fig. 4: Plot showing the yoyo path in depth direction of the AUV for the first 20 minutes of the survey.

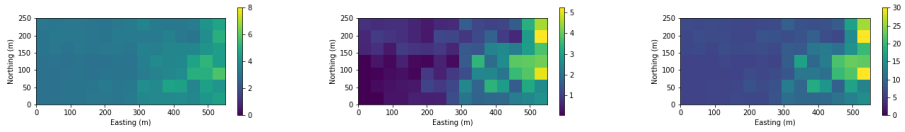

(a) Prior mean (b) Prior standard deviation (c) Prior objective

Fig. 5: The prior states of the survey area.

4.2 Results and evaluation

This section presents the results related to the proposed algorithm. Figure 5 shows the prior mean, the prior standard deviation and the prior objective function. This was found using the empirical mean and variance from the DREAM data in the area, as explained in section 4.

Figure 6 shows the updated state at selected time steps during the mission. The estimated mean, the model variance and the objective function is plotted together with a red line showing the path of the AUV for the last 5 waypoints. The sampling locations are shown as red dots, where the largest disk represents the most recent sampling location.

It can be seen that the AUV chooses sampling locations according to the objective function, but also takes into account the travel length criteria. The western region of the survey area is never explored due to its low objective value. This region is far from the outlet of particles and the model data from DREAM shows both low variance and low concentrations in this area. Thus, this is by our model considered a less interesting area to sample in. However, since our model increases the variance with time in unsampled sites, there is reason to believe that this area would have been explored if the mission ran for a longer time period.

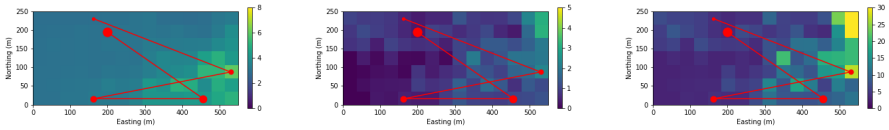

(a) Estimated mean at t = 30 min.

(b) Standard deviation at t = 30 min.

(c) Objective function at t=30 min.

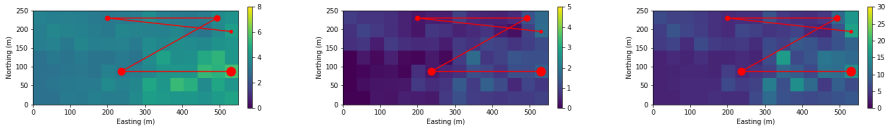

(d) Estimated mean at t = 50 min.

(e) Standard deviation at t = 50 min.

(f) Objective function at t=50 min.

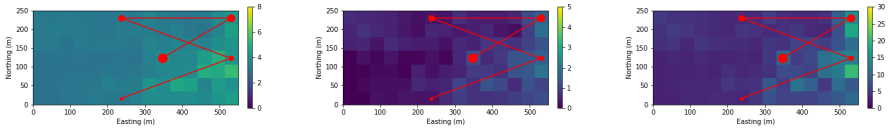

(g) Estimated mean at t = 70 min.

(h) Standard deviation at t = 70 min.

(i) Objective function at t=70 min.

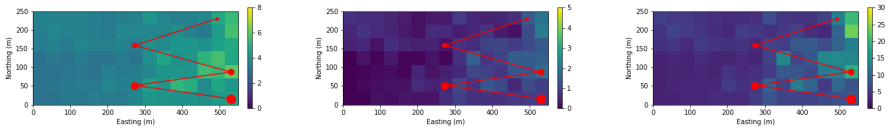

(j) Estimated mean at t = 90 min.

(k) Standard deviation at t = 90 min.

(l) Objective function at t=90 min.

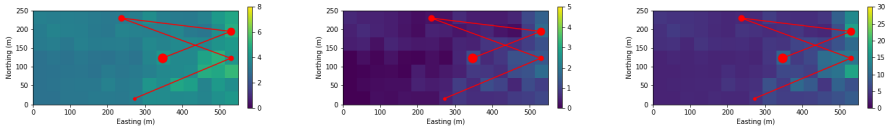

(m) Estimated mean at t = 110 min.

(n) Standard deviation at t = 110 min.

(o) Objective function at t=110 min.

Fig. 6: Updated state at selected time steps during the mission. The estimated mean, the model standard deviation and the objective function is plotted together with a red line showing the path of the AUV for the last 5 waypoints.

At the end of the mission the variance in our GP model is highly reduced compared to the prior variance. This indicates that the sampling algorithm works according to the model. The predicted mean particle concentration is generally higher near the outlet of particles, which is what we expected from theory and from our numerical model DREAM.

4.3 Buoy data and tidal effects

To get a better understanding of the area and the ground truth, data from other marine platforms are considered. In the fjord there is also a moored buoy equipped with a turbidity sensor. The buoy is located about 200 m from the survey area (see Figure 2). Data from the buoy is presented in Figure 7b, which shows the turbidity data for the whole day of the mission. A plot of the tidal effects of the same day is included in Figure 7a. Comparing the turbidity data with the tidal effects, it can be seen that the turbidity measurements seems to be highest close to a tidal low. During a tidal low, the water is flowing westward, out of the fjord. Because the outlet lies to the east of the buoy this is when the highest concentration of particles is transported towards the buoy. The differences between the measured turbidity values are quite large depending on where in the tidal cycle we are, indicating that the transportation of particles in the fjord are highly affected by the tidal effect. Thus, adding this effect to our onboard model would likely lead to an improvement.

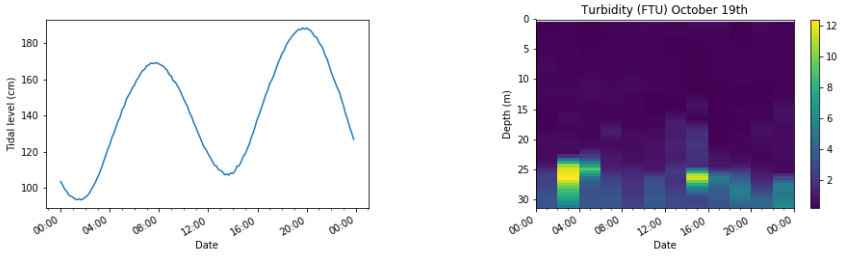

(a) Tidal levels on the day of the survey.

(b) Turbidity measurements from the day of the survey.

Fig. 7: Turbidity measurements from the moored buoy.

5 Closing remarks

This work suggests a method for adaptive sampling of ocean processes with an AUV. The method utilizes prior information from ocean models and combines this information with in-situ observations to obtain a best possible sampling

strategy. A low-complexity onboard model using GPs with a non-stationary covariance function is built based on data from numerical ocean models, and an objective function that is rewarding locations with high uncertainty and high predicted particle concentration in the model, is used to collect samples in informative locations. The method is implemented on an AUV, and tested near a seafill for mine tailings in Frænfjorden, where the objective is to map the concentration of particles around the seafill.

The field trial shows that the variance in the model is reduced and that the locations are chosen according to the objective function. The sampling seems to cover the area such that more samples are collected near the outlet than further away from the outlet. Verification of the predicted particle distribution is hard since there is no easy way to know the ground truth. Still, the trends are as expected from common sense and from the numerical ocean and particle transportation models.

Future work includes expanding the model to 3D, such that it considers more than one depth layer. Also by considering our numerical models and the data from the stationary buoys it is clear that the temporal variability is large, and that our onboard model should include this. Especially, the tidal effects seem to have a high effect on the dynamics, and the tidal cycle can be fairly reliably predicted by a prior ocean model run. Other improvements include an extended path planning method which optimizes for a sequence of points instead of one single sampling location at a time.

Acknowledgements

This work was supported by the Research Council of Norway through 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 Tor Nordam from SINTEF Ocean AS for running the numerical model DREAM, and supplying the data sets.

References

[Berget et al., 2018] Berget, G. E., Fossum, T., Johansen, T. A., Eidsvik, J., and Rajan, K. (2018). Adaptive sampling of ocean processes using an auv with a gaussian proxy model. *IFAC-PapersOnLine*, 51:238–243.

[Binney et al., 2010] Binney, J., Krause, A., and Sukhatme, G. S. (2010). Informative path planning for an autonomous underwater vehicle. In *2010 IEEE International Conference on Robotics and Automation*, pages 4791–4796.

[Binney et al., 2013] Binney, J., Krause, A., and Sukhatme, G. S. (2013). Optimizing waypoints for monitoring spatiotemporal phenomena. *The International Journal of Robotics Research*, 32(8):873–888.

[Cressie and Wikle, 2011] Cressie, N. and Wikle, C. (2011). *Statistics for Spatio-Temporal Data*. CourseSmart Series. Wiley.

[Das et al., 2015] Das, J., Py, F., Harvey, J. B., Ryan, J. P., Gellene, A., Graham, R., Caron, D. A., Rajan, K., and Sukhatme, G. S. (2015). Data-driven robotic sampling

for marine ecosystem monitoring. *The International Journal of Robotics Research*, 34(12):1435–1452.

[Eidsvik et al., 2015] Eidsvik, J., Mukerji, T., and Bhattacharjya, D. (2015). *Value of information in the earth sciences : integrating spatial modeling and decision analysis*. Cambridge University Press, Cambridge.

[Fossum et al., 2018] Fossum, T., Eidsvik, J., Ellingsen, I., Alver, M., Fragoso, G., Johnsen, G., Mendes, R., Ludvigsen, M., and Rajan, K. (2018). Information-driven robotic sampling in the coastal ocean. *Journal of Field Robotics*.

[Griffies et al., 2000] Griffies, S. M., Böning, C., Bryan, F. O., Chassignet, E. P., Gerdes, R., Hasumi, H., Hirst, A., Treguier, A.-M., and Webb, D. (2000). Developments in ocean climate modelling. *Ocean Modelling*, 2:123–192.

[Jun and Stein, 2008] Jun, M. and Stein, M. L. (2008). Nonstationary covariance models for global data. *The Annals of Applied Statistics*, 2(4):1271–1289.

[Krause et al., 2008] Krause, A., Singh, A., and Guestrin, C. (2008). Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Mach. Learn. Res.*, 9:235–284.

[Matérn, 2013] Matérn, B. (2013). Spatial variation. *Meddelanden från Statens Skogsforskningsinstitut*, 36(5):1–144.

[Pinto et al., 2012] Pinto, J., Calado, P., Braga, J., Dias, P., Martins, R., Marques, E., and Sousa, J. (2012). Implementation of a control architecture for networked vehicle systems. *IFAC Proceedings Volumes*, 45(5):100 – 105. 3rd IFAC Workshop on Navigation, Guidance and Control of Underwater Vehicles.

[Rajan and Py, 2012] Rajan, K. and Py, F. (2012). T-rex: Partitioned inference for auv mission control. In Roberts, G. and Sutton, R., editors, *Further advances in unmanned marine vehicles*, chapter 9, page 171 –199.

[Rasmussen and Williams, 2005] Rasmussen, C. E. and Williams, C. K. I. (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.

[Rye et al., 1998] Rye, H., Reed, M., and 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 et al., 2008] Rye, H., Reed, M., Frost, T., Smit, M., Durgut, I., Johansen, O., and 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.

[Slagstad and McClimans, 2005] Slagstad, D. and McClimans, T. (2005). Modeling the ecosystem dynamics of the barents sea including the marginal ice zone: I. physical and chemical oceanography. 58:1–18.

[Stewart, 2008] Stewart, R. H. (2008). *Introduction To Physical Oceanography*. Texas A & M University.

[Trannum et al., 2010] Trannum, H. C., Nilsson, H. C., Schaanning, M. T., and 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 et al., 2006] Wassmann, P., Slagstad, D., Riser, C. W., and 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 et al., 2012] Zhang, Y., Godin, M. A., Bellingham, J. G., and Ryan, J. P. (2012). Using an autonomous underwater vehicle to track a coastal upwelling front. *IEEE Journal of Oceanic Engineering*, 37(3):338–347.