Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnline CONFERENCE PAPER ARCHIVE

IFAC PapersOnLine 51-29 (2018) 238–243

Adaptive Sampling of Ocean Processes Using an AUV with a Gaussian Proxy Model

Gunhild Elisabeth Berget * **Trygve Olav Fossum** ** **Tor Arne Johansen** * **Jo Eidsvik** *** **Kanna Rajan** *

* *Department of Engineering Cybernetics, Center for Autonomous Marine Operations and Systems, Norwegian University of Science and Technology, 7491 Trondheim, Norway*
** *Department of Marine Technology, Center for Autonomous Marine Operations and Systems, Norwegian University of Science and Technology, 7491 Trondheim, Norway*
*** *Department of Mathematical Sciences, Norwegian University of Science and Technology, 7491 Trondheim, Norway*
(E-mail: gunhild.berget@ntnu.no, trygve.o.fossum@ntnu.no, tor.arne.johansen@ntnu.no, jo.eidsvik@ntnu.no, kanna.rajan@ntnu.no).

Abstract: This paper presents methods for building and exploiting compact spatial models on board an autonomous underwater vehicle (AUV) towards tracking suspended material plumes. The research is aiming to improve real-time monitoring of dispersal dynamics connected to marine industries such as oil and mine tailing. By exploiting in-situ information from sensors, the AUV is able to assimilate and adapt the mission capitalizing on all the information available. The spatial model is built using Gaussian process approximations and an objective function for path planning is suggested to maximize the value of the collected information.

Keywords: Gaussian processes, Adaptive sampling, AUV, Oceanography

1. INTRODUCTION

Creating models describing the ocean 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). Still, model verification and data assimilation continues to be a challenge that prompts the need for data sampled from the ocean. Such data is commonly obtained using either remote sensing, ships or buoys. This data is usually expensive to acquire and process. Hence, the ocean tends to be *undersampled* and strategic planning of missions are essential to retrieve as much information as possible. Planning of missions are usually based on historical data or simulation data from numerical models, but often the real world differs much from these data. Hence, being able to adapt the mission in real-time, adjusting the plan based on current observations, will likely improve the modeling efficiency.

In this paper we focus on a method using an AUV for sampling in-situ oceanographic data with a 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 model onboard the AUV is required. The numerical 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 proxy model represents the current state of the ocean at the time, and can be updated when new information are added.

In addition to the GP proxy model, this paper presents an objective function for path planning aiming to maximize the value of information from the samples. The objective function explores the area by choosing locations assumed to be information rich, and also considers the limitations of the AUV.

As a case study, an area in Frænfjorden, Norway containing a seafill for submarine mine tailings is investigated. The goal is to track the particle dispersal near this seafill, aiming to improve real-time monitoring of dispersal dynamics. Two existing numerical ocean models, SINMOD and DREAM, are used to train the GP proxy model creating a prior proxy model of the particle concentration. Having the prior model ready, the AUV can be deployed, and sensor readings can be used to update the proxy model onboard the AUV in real-time. In this paper a simulation study is done using data from the numerical models as sensor readings for the AUV. Figure 1 gives an overview

⋆ 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.

10.1016/j.ifacol.2018.09.509

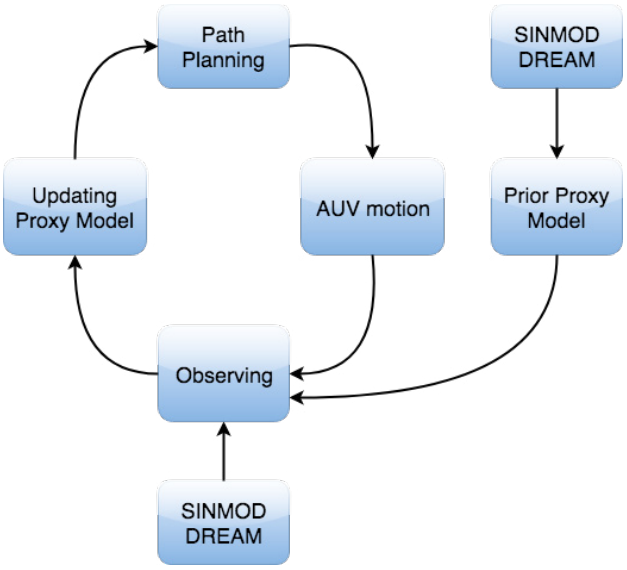

Fig. 1. Block diagram showing the simulation routine. A prior GP proxy model is built using training data from SINMOD and DREAM. Then sensor readings are used to update the proxy model real-time onboard the AUV, and the updated model is used to find the next sampling locations. In the simulation study, the sensor readings are provided by a test data set from SINMOD and DREAM.

of the method proposed in this paper, showing a block diagram of the simulation routine.

1.1 Related work and contribution

GPs (Cressie and Wikle, 2011; Eidsvik et al., 2015) are powerful for creating non-parametric, simple and time-effective models, and are widely used when creating a simplified spatial model. Using GPs for environmental sensing is among others explored in Krause et al. (2008); Zhang et al. (2012); Binney et al. (2010); Das et al. (2015).

In Krause et al. (2008), a method for static sensor placements is suggested using GPs and maximization of mutual information. Others use moving sensors attached to 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. This 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.

Zhang and Sukhatme (2007) create an adaptive sampling algorithm based on local linear regression and minimizing estimation error for a sensor network including static sensors and an AUV. Yilmaz et al. (2008); Jadaliha and Choi (2013) discuss environmental sensing and adaptive sampling using more than one robotic vehicle.

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.

Section 2 presents preliminaries explaining the spatial model and data assimilation procedures. The method used is presented in section 3, before section 4 presents simulation results.

2. MODELLING

Having a complex numerical model describing the ocean onboard the AUV is not practical due to computational limitations. Keeping an onboard representation of the environment is resolved using a simpler proxy model, based on Gaussian processes (GP). This section gives a short introduction to the spatiotemporal modelling, and the data assimilation methods used in this paper.

2.1 Numerical oceanographic models

Numerical oceanographic models are in this paper used both to train the proxy model and for simulation purposes. SINMOD and DREAM are two existing models that describe ocean dynamics, and the release and transport of drill cuttings, respectively. SINMOD is a fully coupled hydrodynamic, sea ice and ecological ocean model (Slagstad and McClimans, 2005; Lindstrøm et al., 2009). It is based on the fundamental NavierStokes equations and uses a nesting technique where high resolution models obtain their boundary conditions from larger model domains with lower resolution. This can be repeated in several steps to achieve high resolution for selected areas. SINMOD is established in configurations with horizontal resolutions ranging from 20 km to 32 m. 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, 2008). It provides time series of both 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, which will be delivered by SINMOD, as well as information about the release (amount, rate, densities, grain size distribution).

For generating the forecast data as input for the AUVs onboard model, SINMOD has been set up with 32 m resolution. Bathymetry data is based on DBM Sør-Norge, supplemented by OLEX data recorded by SINTEF Materials and Chemistry inside Frænfjorden. The atmospheric input data is produced using the Weather Research and Forecasting (WRF) (http://www.wrf-model.org/index.php) model simulated with boundary values from the ERA-Interim reanalysis, and climatologic data for freshwater run-off is used. This data is then forwarded and used as input for the DREAM model. The model area is Frænfjorden (Norway) and the data is from two consecutive days April 1st and April 2nd 2013. Data from 1st April is used

as training data for the proxy model, and data from 2nd April is used as test data in the simulation study.

2.2 Spatial Model

A GP is chosen to model the underlying spatial dependencies of the particle concentration. The 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 are assumed to be Gaussian with mean μ_i and variance σ_i^2. The random variable defining the concentration at location s_i is denoted $x(s_i)$. Hence, the joint distribution of the state at all the locations $\boldsymbol{x} = [x(s_1), \ldots, x(s_n)]$ is multivariate Gaussian

$$\boldsymbol{x} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \tag{1}$$

with mean vector $\boldsymbol{\mu} = [\mu_1, \ldots, \mu_N]^T$ and a positive definite covariance matrix $\boldsymbol{\Sigma}$. 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 known 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

$$R_{ij} = (1 + \phi h_{ij}) \exp(-\phi h_{ij}), \tag{2}$$

where $h_{ij} = |s_i - s_j|$ is the Euclidian distance between two locations and ϕ 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.

2.3 Prior distribution

The model is assumed to be updated sequentially for time steps $t = 1, \ldots, T$ adding information from observations in every time step. The initial prior belief at $t = 0$ ($\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$) is found using the training data from the numerical models (SINMOD/DREAM). The empirical mean of the training data in each location is used as the prior mean $\boldsymbol{\mu}_0$. Assuming M data $[y^*_{i,1}, \ldots, y^*_{i,M}]$ in location s_i, this is given by

$$\mu_i^* = \frac{1}{M} \sum_m y^*_{i,m}, \tag{3}$$

and the prior mean of the proxy model is obtained as the vector $\boldsymbol{\mu}_0 = [\mu_1^*, \ldots, \mu_N^*]^T$.

A common approach for GPs is to simplify and assume the same variance in each location. However, when modeling ocean processes factors such as topology, currents, wind patterns, and freshwater run-off in coastal areas imply that some locations will have elevated variability. Thus, the prior variance of the state in each location is chosen to be the empirical variance from the training data (Stein, 2005)

$$\sigma_i^{*2} = \frac{1}{M-1} \sum_m (y^*_{i,m} - \mu_i^*)^2. \tag{4}$$

The entries of the prior covariance matrix $\boldsymbol{\Sigma}_0$ are given by

$$\boldsymbol{\Sigma}_0(i,j) = \sigma_i^* \sigma_j^* R_{ij} \tag{5}$$

where R_{ij} defines the correlation between points s_i and s_j as defined by (2).

2.4 Data assimilation

To model the temporal changes, a simple 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 can be tuned to fit the modelled domain. This temporal process alone does not model the dynamics of the process, and hence, we rely on the observations from the AUV to catch the changes.

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, a matrix that contains 1 entries only at the sampled indices, otherwise it is 0. $\boldsymbol{\epsilon}_t \sim N_N(0, \boldsymbol{\Omega})$ is a normally distributed error term with zero-mean and covariance $\boldsymbol{\Omega}$, assumed to be Gaussian, describing the measurement noise.

Since a GP is fully represented by its mean and covariance matrix, these are the only thing 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)

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

3. METHOD

Having set the foundation by suggesting a proxy spatial model in section 2, we now proceed to explain the adaptive sampling method, including the path planning method using an objective function and the overall sampling algorithm.

3.1 Objective function

To obtain an information rich path for the AUV, an objective function is suggested. The function is created based on three criteria.

(1) Locations with high variance are preferred
(2) Locations close to the previous sampling location are preferred
(3) Locations with high predicted concentration are preferred

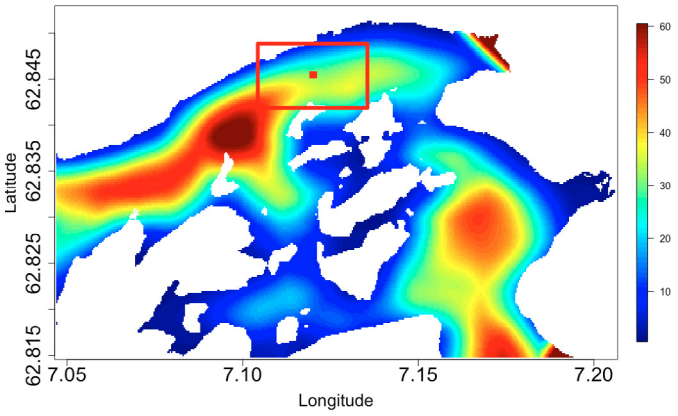

Fig. 2. Plot of the bathymetry of Frænfjorden. The red rectangle shows the selected area, and the red dot indicates the location of the seafill.

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 criteron comes from the energy and speed limitation of the AUV. When choosing the next sample location, it is essential that the AUV does not travel too far. The last 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, this last criterion is inspired by this observation, and assumes that locations with high predicted concentration will be rich with information.

The suggested objective function is then created by having a term for each of these criteria. At time step t for location s_i given the previous sampling location S_{t-1}, the objective function is given by

$$f_t(s_i) = \theta_1 \sigma^2_{i,t} - \theta_2 |S_{t-1} - s_i| + \theta_3 \mu_{i,t} \tag{9}$$

where the constant parameters $\boldsymbol{\theta} = [\theta_1, \theta_2, \theta_3]$ defines the weighting for each of the three criteria. These parameters together with the parameter V in the updating equations (8) are tuned by trial and error to obtain the desired behavior of the AUV.

3.2 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 $s_i \in [s_1, \dots, s_N]$,

$$S_t = \text{argmax}_{s_i}(f_t(s_i)). \tag{10}$$

The details of the sampling method are given in Algorithm 1.

The method is a greedy method that sequentially chooses the best sampling location. First, the spatial GP model is created using training data from simulations. Then the sampling starts by evaluating the objective function and choosing the location which maximizes it. After reaching the desired location and doing observations, the GP model is updated and the variance is increased in the unobserved locations.

Algorithm 1 Sampling method

```
procedure SAMPLING
    Initialize GP
    for t = 1, ..., T do
        for s = s_1, ..., s_N do
            Evaluate f_t(s)
        Choose S_t = argmax_s(f_t(s))
        Go to location S_t
        Retrieve observations from S_t
        Assimilate data according to (8)
```

4. SIMULATION RESULTS

Data for April 1st 2013 describing the particle concentration in the fjord obtained from DREAM was used to train the GP, and to create the prior mean $\boldsymbol{\mu}_0$ and covariance $\boldsymbol{\Sigma}_0$. For simplicity, only a small area around the seafill was considered (2560 m $\times$ 1280 m) and the area was divided into a regular grid with grid cells of size 32 m $\times$ 32 m (the same grid as for the simulation data). Also, in this initial simulation only one depth layer was considered at $\approx$ 15 m depth, but this could be expanded to 3 dimension considering multiple snapshots in different depth layers. Figure 2 shows the bathymetry of the fjord and the selected area as a red rectangle. From the training data it was observed that the distribution of particles in the location of the seafill was rapidly changing and had very little correlation with the neighboring sites. Thus, this location was disregarded in our model. When plotting the results, this was handled by setting the variance to 0 in the location, and using the true value from the test data as the mean.

The spatial model and the sampling method was implemented using the language R (R Core Team, 2017), and test data from DREAM (April 2nd) was used as sensor readings for the AUV. Hence, for this simulation we consider the test data from DREAM to be the true distribution at all time. The time step was discretized into intervals of 5 minutes. A total of 54 updates were simulated, which corresponds to monitoring the outlets for 4.5 hours (270 minutes). The values used for the tuning parameters was $\boldsymbol{\theta} = [1, 125, 100]$ and $V = 0.05$.

The results of the simulation study are shown in Figure 3, showing results at four different time steps. The particle concentration is measured in $\mu g/L$, and the color bar shows the intensity at each location. The x-and y-axis shows the distance in metres from the seafill. The first column of plots shows the true values from DREAM. The predicted particle concentration is shown in the second column, and the third column shows the prediction variance together with the path of the AUV showing the 10 most recent sampling locations as small white dots and the current position of the AUV as a large white dot.

Comparing the predicted particle concentration with the truth from the test data it can be seen that the sampling method generally gives a smooth prediction that coincides quite well with our "true" distribution. Still, many of the finer details are overlooked, and more samples are needed to model these details.

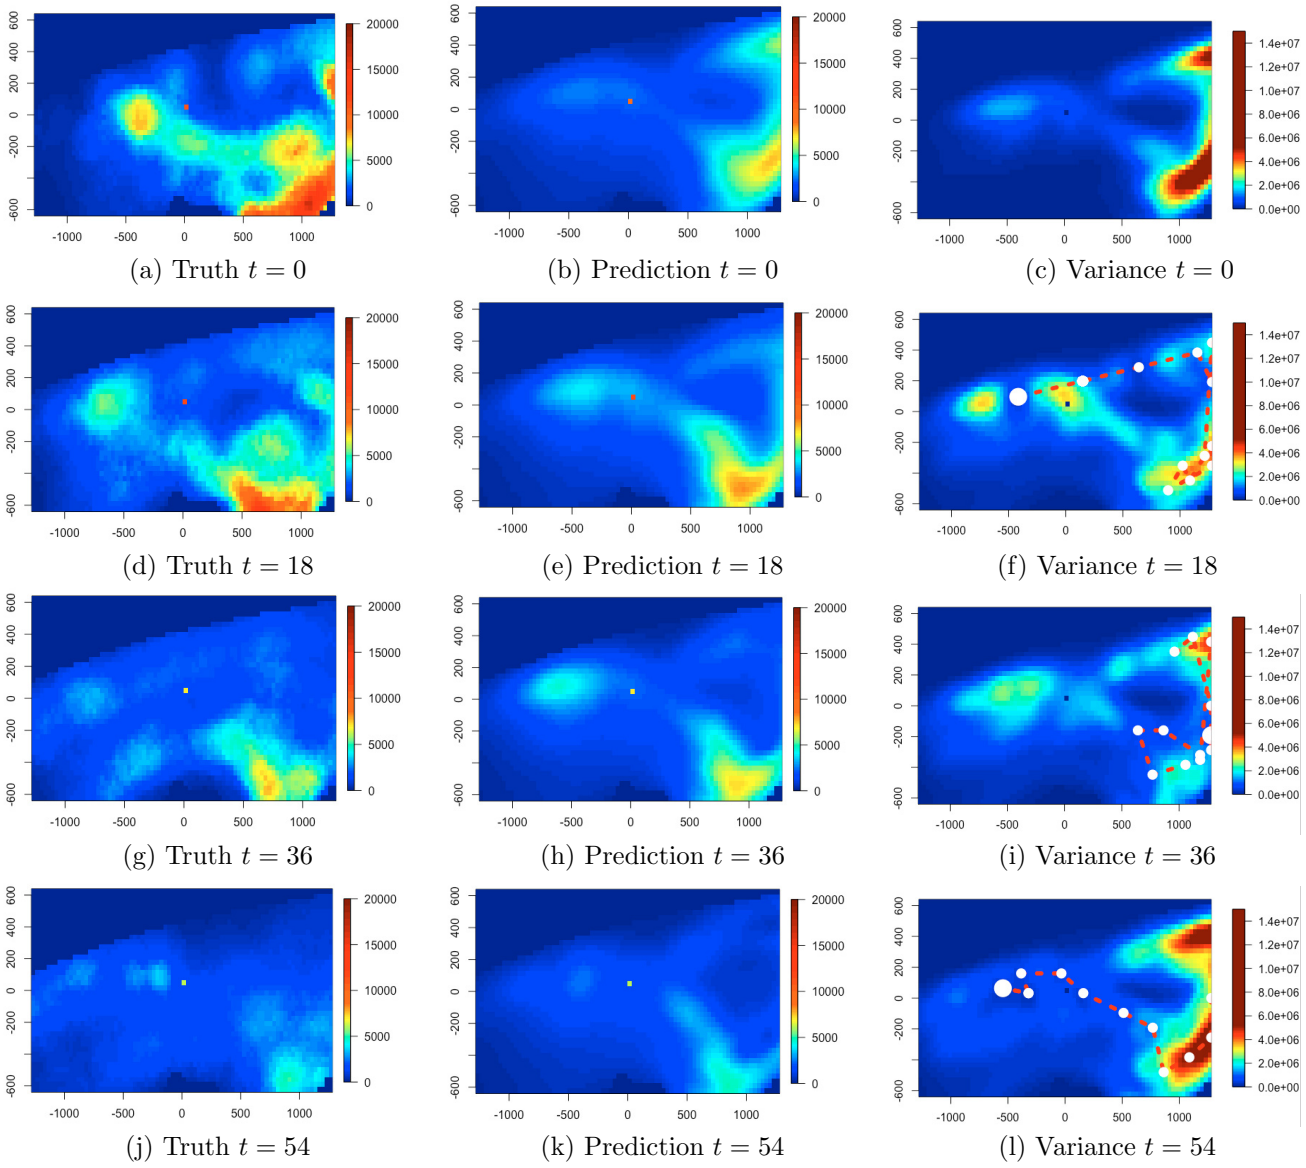

Fig. 3. Results of the simulation study at four different time steps $t = 0, 18, 36, 54$ corresponding to [0,90,180, 270] minutes. The particle concentration is measured in $\mu g/L$, and the color bar shows the intensity at each location. The x-and y-axis shows the distance in metres from the seafill. The first column of plots corresponds to the true values (a),(d),(g) and (j). The predicted mean particle concentration is shown in the second column (b), (e), (h) and (k), and the third column corresponds to the prediction variance (c), (f), (i) and(l). The path of the AUV is plotted as a red line, the small white dots shows the 10 most recent sampling locations and the large dot indicates the current position of the AUV.

Considering the path of the AUV together with the prediction variance, it can be seen that the variance is decreased near the recently sampled locations. The increase in variance proportional to the prior variance can also be clearly seen in the prediction variance plots.

The objective function controls the AUV path in an intuitive way, leading it to unexplored areas on the one hand, and in most cases assuring a reasonable travel speed for the AUV between sampling locations. The travel distance of the AUV lies between 100-400 metres for most time steps, which is a suitable distance given the time step length and the AUV speed.

An issue with the model, is that it does not seem to keep up with the rapid changing of the ocean process. Since the model relies on the observations from the AUV to catch the change in the particle concentration, the prediction results far from recently sampled locations will be inaccurate. When the dynamics of the particle transportation are fast, we will get a delay in the updating of the model. As an example, the prediction results in the upper left area can be considered. At t=18, the prediction in this area shows a smaller concentration than the true value. Then at t=36, the predicted concentration has increased in this area, but in the true model the particles have moved resulting in a low density. Finally, at t=54 the predicted and the true value is quite similar. This example shows the delay in the predicted values.

5. CONCLUSION AND FUTURE WORK

A method for adaptive sampling of ocean processes using an AUV is suggested, and tested using simulation data from numerical models of particle transportation near a seafill. The spatial model reconstructs the true field quite well, showing the same tendencies as the true field. Still, the temporal variability of the particle transporting is faster than the AUV can keep up with, indicating that more samples from multiple vehicles or buoys and/or a better temporal model is useful.

Future work includes expanding the model such that it considers the temporal variability of particle transportation, such that the the non-stationarity of the model is not driven by the collected data alone. Path planning can be improved both by considering optimizing for a sequence of points instead of only choosing one sampling location at a time. Fieldwork is also planned, enabling testing of the method in real ocean conditions. This will give insight in how the method works in the real world, and how this differs from simulation.

ACKNOWLEDGEMENTS

The authors would like to thank Finn Are Michelsen and Raymond Nepstad from SINTEF Ocean AS for running the numerical models SINMOD and DREAM, and supplying the data sets used in section 4.

REFERENCES

Binney, J., Krause, A., and Sukhatme, G.S. (2010). Informative path planning for an autonomous underwater vehicle. In *2010 IEEE International Conference on Robotics and Automation*, 4791–4796. doi: 10.1109/ROBOT.2010.5509714.

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. doi:10.1177/0278364913488427.

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

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. doi: 10.1177/0278364915587723.

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.

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. doi:10.1016/S1463-5003(00)00014-7.

Jadaliha, M. and Choi, J. (2013). Environmental monitoring using autonomous aquatic robots: Sampling algorithms and experiments. *IEEE Transactions on Control Systems Technology*, 21(3), 899–905. doi: 10.1109/TCST.2012.2190070.

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.

Lindstrøm, U., Smout, S., Howell, D., and Bogstad, B. (2009). Modelling multi-species interactions in the barents sea ecosystem with special emphasis on minke whales and their interactions with cod, herring and capelin. *Deep Sea Research Part II: Topical Studies in Oceanography*, 56(21), 2068 – 2079. doi: https://doi.org/10.1016/j.dsr2.2008.11.017. The Proceedings of the ECONORTH Symposium on Ecosystem Dynamics in the Norwegian Sea and Barents Sea.

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

R Core Team (2017). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL `https://www.R-project.org/`.

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

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. doi: https://doi.org/10.1016/S1364-8152(98)00048-6.

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, 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.

Stein, M.L. (2005). Nonstationary spatial covariance functions. *University of Chicago, CISES Technical Report 21*.

Yilmaz, N.K., Evangelinos, C., Lermusiaux, P.F.J., and Patrikalakis, N.M. (2008). Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming. *IEEE Journal of Oceanic Engineering*, 33(4), 522–537. doi: 10.1109/JOE.2008.2002105.

Zhang, B. and Sukhatme, G.S. (2007). Adaptive sampling for estimating a scalar field using a robotic boat and a sensor network. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, 3673–3680. doi:10.1109/ROBOT.2007.364041.

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. doi: 10.1109/JOE.2012.2197272.