Hierarchical Probabilistic Regression for AUV-based Adaptive Sampling of Marine Phenomena

Jnaneshwar Das*, Julio Harvey†, Frédéric Py†, Harshvardhan Vathsangam*, Rishi Graham†,
Kanna Rajan†, Gaurav S. Sukhatme*
*Robotic Embedded Systems Laboratory, Dept. of Computer Science, University of Southern California
†Monterey Bay Aquarium Research Institute, Moss Landing, California
{jnaneshd,vathsang,gaurav}@usc.edu, {jharvey,fpy,rishi,kanna.rajan}@mbari.org

***Abstract*— Marine phenomena such as algal blooms can be detected using *in situ* measurements onboard autonomous underwater vehicles (AUVs), but understanding plankton ecology and community structure requires retrieval and analysis of water specimens. This process requires shipboard or manual sample collection, followed by onshore lab analysis which is time-consuming. Better understanding of the relationship between the observable environmental features and organism abundance would allow more precisely targeted sampling and thereby save time. In this work, we present an approach to learn and improve models that predict this relationship. Coupled with recent advances in AUV technology allowing selective retrieval of water samples, this constitutes a new paradigm in biological sampling. We use organism abundance models along with spatial models of environmental features learned immediately after AUV deployments to compute spatial distributions of organisms in the coastal ocean purely from *in situ* AUV data. We use Gaussian process regression along with the unscented transform to fuse the two models, obtaining both the mean and variance of the organism abundance estimates. The uncertainty in organism abundance predictions is used in a sampling strategy to selectively acquire new water specimens that improves the organism abundance models. Simulation results are presented demonstrating the advantage of performing hierarchical probabilistic regression. After the validation through simulation, we show predictions of organism abundance from models learned on lab-analyzed water sample data, and AUV survey data.**

I. Introduction

Marine phenomena are often observed using proxy measurements. For example, Autonomous Underwater Vehicles (AUVs) measure phytoplankton biomass using specialized optical sensors called fluorometers. Intermediate nepheloid layers (INLs) [1] and chlorophyll thin layers have been sampled using chlorophyll fluorescence and backscatter [1]–[3]. During the Gulf of Mexico oil spill in 2010, hydrocarbon plumes were studied using measurements of colored organic dissolved matter (CDOM) fluorescence onboard AUVs [4]. Observing such marine phenomena directly requires expensive instruments or lab analysis or both, making it infeasible onboard AUVs. By detecting and learning relationships between the phenomena of interest and data measurable *in situ*, we can predict beyond what AUVs can directly sense, expanding their role in oceanography dramatically.

In this work, we explore the learning and improvement of models to predict unobservable marine features from observable proxy features. We focus on detection of ribosomal RNA abundance for various phytoplankton and zooplankton taxa, although the techniques being developed can be applied to other organisms as well. Marine scientists have recently begun to develop automated sensors to measure plankton *in situ* although they are currently not feasible for deployment onboard AUVs [5]. One AUV design (Fig. 1) allows water samples to be collected on demand along with existing *in situ* measurements of environmental parameters such as temperature and salinity (among others). Studies are ordinarily carried out by collecting seawater samples from various locations, followed by analysis on shore. Routine deployment of such AUVs since 2009 followed by lab analysis of collected water samples has resulted in a dataset of organism abundance and corresponding environmental parameters. In this work, we use these data to learn the relationship between environmental parameters and organism abundance. The model thus learned is then used in conjunction with spatial models of environmental parameters to estimate spatial distributions of organisms. We do this in a purely Bayesian setting, allowing us to estimate not only the spatial distribution of organisms, but also associated uncertainty in the predictions. Finally, we develop sampling strategies that utilize uncertainty maps from the combined model to collect new water samples adaptively in future missions with the goal of improving model accuracy.

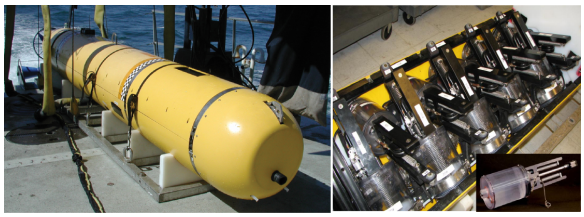

Fig. 1: The Dorado-class *Gulper* AUV (left) at the Monterey Bay Aquarium Research Institute (MBARI). It is equipped with an onboard water sample collection system (right), consisting of ten 1.85 L 'gulpers' that can be individually triggered by the onboard computer.

The relationship between environmental parameters and biological phenomena has been studied in various recent works. Clarissa et al. [6] estimate phytoplankton abundance in the coastal ocean using environmental data such as temperature and salinity from regional ocean modeling system (ROMS) and remote sensing satellites. In the machine learning literature, Phillips et al. [7] use the relationship between temperature, precipitation, and occurrence of a plant species

[1] plumes or clouds of suspended particulate matter

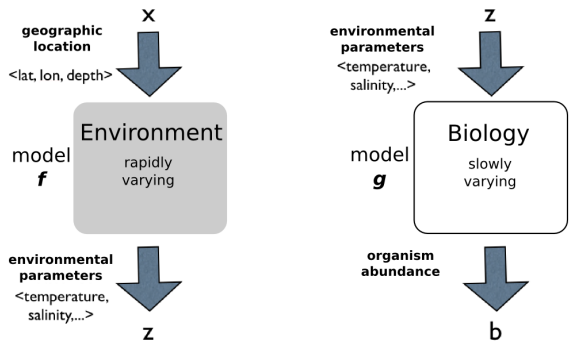

Fig. 2: Two stage regression is suitable for scenarios where relationships between observed and latent variables can be described with two models that can be learned separately.

to estimate its spatial distribution. A challenge in modeling spatial distribution of marine organisms from environmental parameters is the mobility of water due to currents and tidal cycles. Such physical forcing results in highly variable environmental conditions.

Spatial distribution of proxies: We address this issue by performing adaptive sampling of unobserved features in two stages. We observe that the data gathered from AUV *in situ* measurements and lab-analyzed water samples are two distinct sets. The former is collected over a short duration (a few hours) during a pilot AUV deployment. The latter dataset is a result of lab analysis of water samples collected by the AUV over multiple missions spanning years. The dataset obtained from a pilot AUV survey provides a snapshot of ocean environmental conditions within a desired time window, wheres the water samples provide the means to predict biology from the newly estimated environmental conditions from the pilot survey, essentially capturing the connection between observed biology and the environmental context in the form of in-situ measurements. We learn two models from the two datasets as illustrated in Fig. 2, one that maps geographic locations to environmental parameters, and the other that maps environmental parameters to organism abundance.

This work has two goals. First, we demonstrate computation of probabilistic maps of organism abundance from two separate models, *in situ* and offline. Second, we use the probabilistic maps of organism distribution to investigate sampling policies for collection of new water samples to improve model accuracy, thus closing the loop.

The structure of this paper is as follows. In section II we discuss related work in adaptive sampling. In section III, we briefly describe Bayesian regression and our choice of models for the two stages. Next, we discuss our approach for propagating uncertainty through the models. In section IV, we present simulation results and compare sampling strategies. In section VI we present a case study on historical AUV and water samples analysis data. Finally, we conclude and present future directions.

II. Related Work

Earlier work in adaptive sampling has explored reconstruction of scalar fields [8]–[10] adaptively, but only in cases where the target field is directly observable. Reactive policies for adaptive water sample acquisition have been investigated in a multi-utility criterion based approach [11]. Bayesian regression with uncertainty in query points has been discussed by Girard et al. [12] which presents an approximate analytical approach for Gaussian process regression with input uncertainty. Ko et al. [13] discuss unscented Kalman filter (UKF) with GP observation and process models. To the best of our knowledge, our work presents the first analysis of modeling of processes consisting of two distinct relationships. Additionally, we extend prior work in adaptive sampling to hierarchical regression. Further, the novelty of this work is in its inter-disciplinary nature and the novel use of Gaussian Processes for studying ecology.

III. Technical Approach

Supervised learning from labeled training data to predict real valued estimates for query points falls under the area of regression analysis. Assume we have observations drawn from a noisy process given by,

$$y = f(x) + \epsilon \tag{1}$$

where ϵ is a Gaussian noise term. Our goal is to learn the function f to predict the output at unobserved location x^*. Traditionally, regression has been treated in a deterministic setting. For example, a commonly used technique called ordinary least-squares (OLS) regression learns a deterministic estimator that predicts the output at unobserved locations for noiseless query input points. The estimator has an accuracy dependent on the observation noise ϵ and quality of training data. In contrast, Bayesian regression allows estimation of predictive distributions for query points instead of mere point estimates. When used for spatial statistics, Bayesian regression offers an uncertainty map in addition to a prediction map useful in two ways. First, prediction uncertainties along with the predictions provide a measure of model accuracy. Second, having a probabilistic uncertainty map allows us to use information-theoretic adaptive sampling or active learning [2] techniques to acquire new training data.

A. Gaussian process regression

Our goal is to learn two probabilistic regression models from two separately acquired datasets. We use Gaussian Process regression (GP) [14], a non-parametric Bayesian regression technique commonly used in geostatistics under the name 'Kriging'. It assumes that the samples from the function to be estimated are normally distributed with the covariance between samples given by a 'kernel' or covariance function. We consider the case where the observations are unbiased, that is, the mean of the joint Gaussian distribution is zero. This can be simply satisfied by 'deameaning' the observed data. As a result of its formulation, given sufficient

[2] Active learning is a branch of machine learning that deals with choosing new training data to learn a model from. In our study, this is especially important because collecting and analyzing water samples is expensive and time consuming, i.e. the cost of labeling is high. Searching the input space prior to sample acquisition and labeling allows one to detect the best candidate for further analysis.

data points, GP regression automatically achieves model regularization from data only, without having to choose model complexity parameters *a priori*. Additionally, unbiased GP regression is defined completely by a kernel function that controls how quickly the input space gets decorrelated. This enforces smoothness constraints in the trained function, ideal for spatial models where usually the observed values for nearby input samples are more correlated than the ones farther apart. The first function f is learned from the data from a pilot AUV deployment and models the spatial distribution of environmental parameters. The second function to be learned, g, is not spatial, and maps an arbitrary number of environmental parameters to a real valued observation of organism abundance. It is trained from a dataset obtained after lab analysis of water samples collected from Monterey Bay over a few years, from various geographic locations at different times. However, GP is also an ideal choice because it is non-parametric and purely data-driven. The caveats are that with very few data points, GP might be susceptible to outliers. This is something we plan to investigate further in future work. However, even with data points being sparse for the second model, and being obtained slowly, we make the underlying assumption that the relationship between the environmental conditions and the biological abundance of organisms does not change as rapidly as the distribution of spatial environmental parameters in the ocean learned from the first model. Hence, the biological abundance model can be learned *a priori*, and updated only when new lab-analyzed data are available.

GP regression is performed as follows. Assume we have training data given by $D = < x_1, y_1 >, < x_2, y_2 >, ..., < x_n, y_n >$, drawn from the noisy process,

$$y_i = f(x_i) + \epsilon \tag{2}$$

where ϵ is a Gaussian noise term.

With this prior, we can compute the posterior mean and covariance for the test data points with the following equations,

$$GP_\mu = k_*(K + \sigma_n^2 I)^{-1} y \tag{3}$$

$$GP_\Sigma = k(x_*, x_*) * + k_*(K + \sigma_n^2 I)^{-1} k_* \tag{4}$$

The kernel function k is usually chosen to be a squared-exponential function given by,

$$k(x_p, x_q) = e^{-\frac{1}{2\lambda^2}|x_p - x_q|^2} \tag{5}$$

where λ is the decorrelation length scale. The hyperparameters for the kernel function can be learned using iterative methods such as conjugate gradient descent.

During the pilot experiment, the AUV measures environmental condition given by the set $Z = \{z_1, z_2, ..., z_n\}$, observed at geographic locations $X = \{x_1, x_2, ..., x_n\}$. The training dataset is given by the n-tuple $T_{env} = \{< x_1, z_1 > , < x_2, z_2 >, ..., < x_n, z_n >\}$. This can be used to learn an environmental model f that infers the predictive distribution over environmental condition at unobserved locations X^*, given by $p(Z^*|X^*, X, Z)$.

Similarly, we train a biological model g from lab-analyzed gulper samples. This dataset consists of the organism abundance (measured as optical density or OD at A450 nm) $\hat{B} = \{b_1, b2, ..., b_d\}$ observed for environmental parameters $\hat{Z} = \hat{z_1}, \hat{z_2}, ..., \hat{z_d}$. The training dataset is given by $T_{bio} = \{< \hat{z_1}, \hat{b_1} >, < \hat{z_2}, \hat{b_2} >, ..., < \hat{z_d}, \hat{b_d} >\}$. The model g learned from this dataset infers predictive distribution of organism abundance at unobserved environmental conditions Z^*, given by $p(B^*|Z^*, \hat{Z}, \hat{B})$.

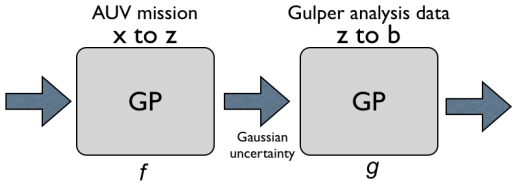

Fig. 3: Illustration of hierarchical regression where the prediction covariance of the first model is propagated through the second model to estimate the combined mean and the covariance for the final predictions.

Hierarchical probabilistic regression: Our goal is to perform our regression using two probabilistic regression models learned separately. Given a set of unobserved locations, we can use the predictive distribution of environmental parameters from the first model, as query-points to the second model. This is illustrated in Fig. 3.

The predicted environmental parameters from the first model are given by $p(z^*|x^*, X, Z)$ where x^* is an unobserved location, and X Z are from the training dataset for model f. Since x^*, X, and Z are deterministic, the distribution is simply $p(z^*)$. For this predicted distribution for environmental feature z^*, the second model g gives us the predictive distribution for organism abundance b^* as $p(b^*|z^*)$. Note that we omit the deterministic training dataset for model g for clarity of notation. The predictive distribution for organism abundance at the unobserved location x^* is obtained by marginalizing over the predictive distribution of environmental parameters for that location given by,

$$p(b^*) = \int p(b^*|z^*)p(z^*)\mathrm{d}z^*. \tag{6}$$

We cannot compute a closed-form solution for this integral with the non-linear function g. Additionally, we do not expect this distribution to be Gaussian, specifically when g is nonlinear. We can compute this integral using Monte Carlo methods [12]. However, with high-dimensional input (six, in case of model g), the method suffers from the *curse of dimensionality*[3], requiring thousands of samples drawn from $p(z^*)$, for each unobserved location x^* for which we wish to infer $p(b^*)$. Hence, we use another technique called *unscented transform* commonly used in optimal state estimation.

[3]In machine learning, as the number of dimensions increases, the required number of data for converging to the right parameters often increases exponentially because of increased number of parameters to train for the model. This is commonly called the curse of dimensionality.

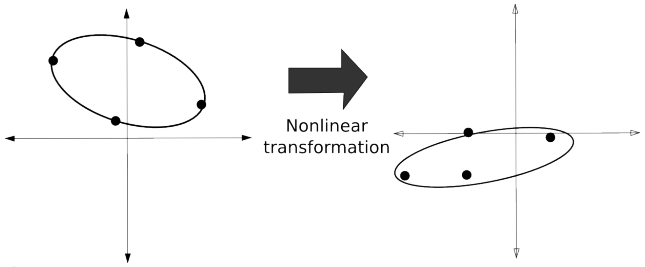

Fig. 4: The unscented transform propagates 2n+1 sigma points corresponding to each query point. The propagated points are used to compute the predicted mean and variance.

B. Propagation of uncertainty: The Unscented Transform

The unscented transform (UT) [15] approximates the result of applying a nonlinear transformation to a given mean and covariance estimate. The UT involves computing a set of points called sigma points, which have the same first and second order moments as the given mean and covariance estimate. After applying the nonlinear transformation to the set, an unscented estimate of the resulting mean and covariance is obtained by computing the statistics of the transformed set of sigma points, as illustrated in Fig. 4. Although UT appears similar to Monte-Carlo (MC) integration, the sample points are drawn randomly according to the input probability distribution in MC approaches, whereas they are computed deterministically in the unscented transform.

We use $2n+1$ sigma points given by the vector S, where n is the dimensionality of the input data. The sigma points are given by,

$$S_0 = \mu \tag{7}$$

$$S_i = \mu + \sqrt{n+\lambda} \tag{8}$$

$$S_{i+n} = \mu - \sqrt{n+\lambda} \tag{9}$$

where i=1...n, and $\lambda \in \mathbb{R}$ is chosen to scale the higher order terms of the set of sigma points. The sigma points are now the input to function g resulting in $\bar{X} = g(S)$. The new mean is computed using a linear combination of the transformed sigma points $\bar{X}$, given by

$$\bar{Y} = \sum_{i=0}^{2n} W_i \bar{X}_i \tag{10}$$

where the weights W are given by,

$$W_0 = \frac{\lambda}{n+\lambda} \tag{11}$$

$$W_i = \frac{1}{2(n+\lambda)} \tag{12}$$

$$W_{i+n} = \frac{1}{2(n+\lambda)} \tag{13}$$

The new covariance by

$$P = \sum_{i=0}^{2n} W_i (\bar{X}_i - \bar{Y})(\bar{X}_i - \bar{Y})^T \tag{14}$$

For our study, the unscented transform is used to compute the mean and covariance of the resulting predictive distribution of organism abundance from function g, given input with Gaussian uncertainty, predicted by model f. In the following section, we validate through simulation the use of the unscented transform to fuse the two models f and g.

IV. Simulation Studies

To evaluate our approach, we consider two one-dimensional processes representing the environmental and biological models. We picked two non-linear functions, a sigmoid for the first model and a sine for the second model, though other functions can also be used for this analysis. The functions are given by,

$$f = \frac{1}{1+e^{-10x-5}} \tag{15}$$

$$g = \frac{sin(2\pi x)}{2} + 0.5 \tag{16}$$

Our goal is to predict values of the second function for unobserved locations in the first function. This is given by $g(f(x))$. Additionally, we want to use a Bayesian setting allowing estimation of the corresponding predictive variance. The uncertainty can then be used to improve model accuracy adaptively with new samples.

A. Model learning and prediction

First, we discretize the input space X for function f in which candidate sampling locations have to be selected. Both models, f and g are analyzed within the closed interval $[0, 1]$. We start by drawing five samples from both model f and g with Gaussian noise of variance $= 0.0001$. This gives us the training datasets for the two functions. The hyperparameters of two GP models f and g are now learned using the GPML toolbox [16]. Next, we choose all discrete points for the input space of function f, and remove from these the five data points already used to train the model. This provides us with a set of unobserved locations for model f. The trained GP model f is now used to predict the mean and variance for the unobserved locations. This is the input to the second function g. To propagate the mean and variance, we use the unscented transform described in Section III-B. This gives us the predicted mean and variances for the combined model (result in Fig. 5).

B. Sampling policy

Once we have a combined model allowing us to estimate the predictive distribution of b from x directly, we evaluate three sampling strategies. First, we tried random sampling as the baseline policy. Here, we sequentially drew k samples from a uniform random distribution over the candidate unobserved locations for sensors, where k is the number of new samples we wanted to adaptively obtain. Next, we tried highest variance sampling by greedily choosing k points with highest predictive variance of model f only. This strategy doesn't consider the propagation of uncertainty through model g. The third strategy considers the predictive variance of the output of function g having considered the propagated uncertainty from model f. Fig. 6 shows the result of 100 trials comparing the three strategies. At the end of the training phase, the average normalzied root mean-squared error (NRMSE) for 100 trials was 0.2 With random sampling, this reduced to 0.076 after choosing 10 new samples. With highest variance sampling, this reduced to 0.036 for the case where only model f variance was considered. With our

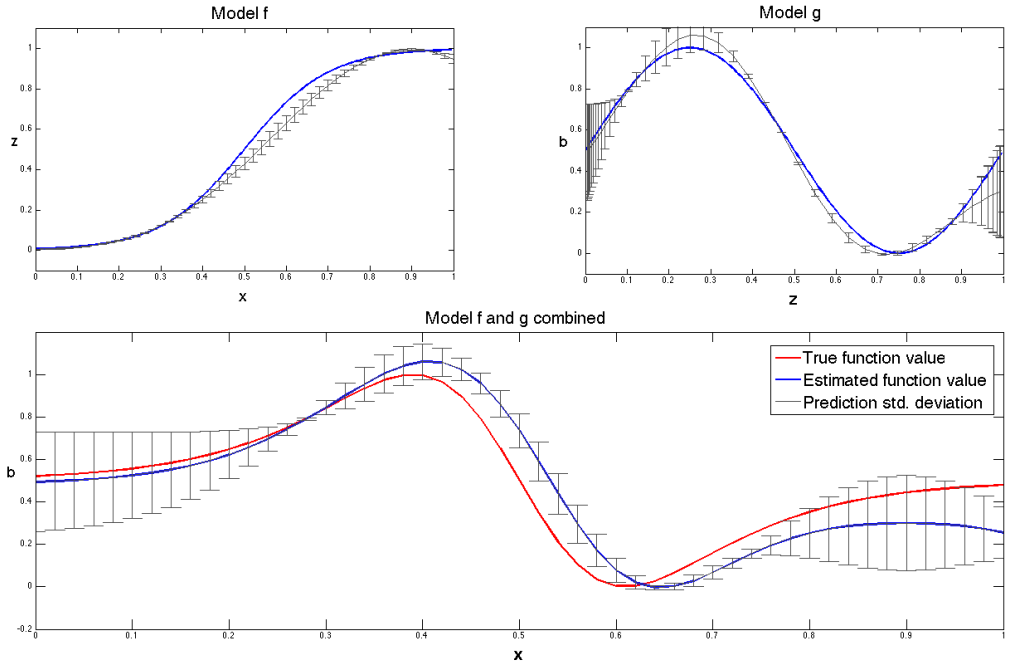

Fig. 5: The two one-dimensional models, and the combined model after learning from training data.

proposed approach considering combined variance for model f and g, the NRMSE was 0.015. Random sampling resulted in the highest NRMSE for any number of new samples drawn. The error was lowest for highest variance sampling with combined f and g variances, demonstrating that our approach reduces the model error the most.

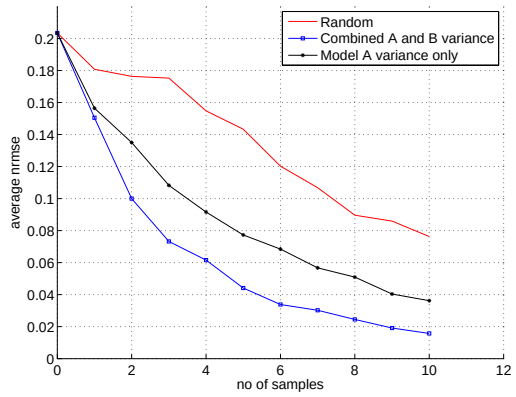

Fig. 6: Comparison of the average normalized root-mean-square error (NRMSE) for 100 trials with random sampling, sequential sampling using highest variance for spatial model only, and sequential sampling using variance of the combined spatial and biological models.

V. Organism Abundance Model Validation

We used two water sample datasets to train and test the organism abundance models discussed in this work. The first consists of data from 62 lab-analyzed water samples from Spring 2009. The second dataset consists of 70 lab-analyzed water samples from Fall 2010. Since the number of data points for each dataset is limited, we perform s-fold cross-validation, a technique for model validation ideal for small datasets. We do this individually on the Spring and Fall datasets, followed by training on the Spring dataset and testing on the Fall dataset, and vice versa. In this section, we present the results of the four validation cases.

A. s-fold cross validation analysis

Since we have a limited number of data for training and testing our models, s-fold cross-validation provides an efficient way to test the accuracy of our model. We do this by splitting the data into s folds (s = 5, in our analysis). Then, we use s-1 folds to train our data and one fold to test it. This is repeated s times and the average prediction error is computed. As in our earlier analysis, we used normalized root-mean-square error (NRMSE) as our error metric (result in Fig. 7).

B. Prediction across seasons

In additional to the s-fold cross validation analysis described above, we trained and tested our models across seasons. There are two cases for this analysis – a) train model on the Spring 2009 dataset (62 samples), and test it on the Fall 2010 dataset (71 samples), and b) train the model on the Fall 2010 dataset and test it on the Spring 2009 dataset. The plots of the normalized root-mean-square error for the two cases and the result of this analysis are shown in Fig. 8. Organism abundance can vary among seasons or due to a variety of other factors.. For example, we notice prediction of calanoida (copepods) in Fall is extremely inaccurate (NRMSE > 1) when the model is trained on the Spring dataset. Similarly, podoplea predictions show high inaccuracy when trained on Fall 2010 dataset and tested on Spring 2009 dataset.

C. Summary

The analysis of the organism abundance models with the various datasets demonstrates the following. Seasonal char-

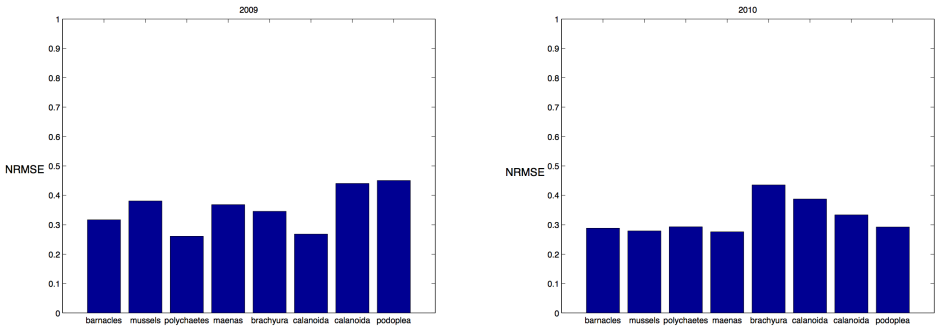

(a) Spring 2009 dataset (b) Fall 2010 dataset.

Fig. 7: Results of s-fold cross-validation on Spring 2009 and Fall 2010 datasets.

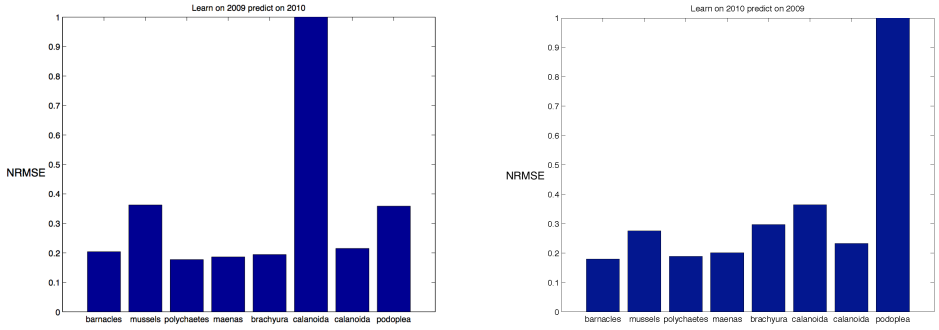

(a) Train on the Spring 2009 dataset and test on the Fall 2010 dataset.

(b) Train on the Fall 2009 dataset and test on the Spring 2010 dataset.

Fig. 8: Validation of models across seasons.

acteristics are pronounced for certain organisms, including calanoida and podoplea.

Overall, the s-fold analysis on the individual seasonal datasets show that the organism abundance models have an average error of $\sim 30\%$. Some organisms show consistently better predictability (e.g., barnacles and polychaetes). We are now working with marine biologists at MBARI to better understand the effect of seasons on the predictability of the organisms considered in this study.

VI. Case Study

Having validated our sampling approach in simulation, and having analyzed the accuracy of the organism abundance models, in this section we present the estimation of organism abundance from *in situ* measurements of environmental parameters by an AUV. The environmental parameters were measured by the *Gulper* AUV during a pilot experiment in Fall 2010 ($\sim$ 2600 data points, shown in Fig. 9). We learned organism abundance models for eight species using results of lab analysis on 96 water samples collected between 2009 and 2010. This dataset consists of the recorded organism abundance for various species along with the environmental parameters measured by the AUV when they were collected. Six environmental parameters were used for this analysis, viz. temperature, salinity, nitrate concentration, dissolved oxygen concentration, chlorophyll fluorescence, and optical backscatter. Using the data from the pilot AUV survey, these parameters within a regular grid with a resolution of 0.001^o

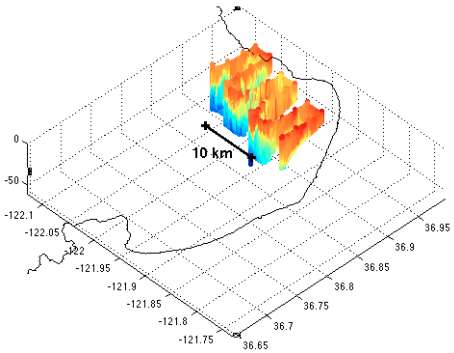

Fig. 9: Temperature measurements by the *Gulper* AUV after a lawnmower survey in Monterey Bay in October 2010 (looking north west). Five other environmental parameters (salinity, dissolved oxygen concentration, nitrate concentration, chlorophyll fluorescence, and optical backscatter) were also measured. This data was used to learn probabilistic spatial models for environmental parameters.

in the latitude and longitude axes, and $1m$ in the dept axis are shown in Fig. 10.

Eight different organism abundances were obtained by applying molecular genetic techniques to organisms collected from each water sample via filtration. Fig. 11 shows the abundance and prediction uncertainty map for calanoida in using the hierarchical probabilistic regression method discussed described above.

We applied the method presented in section IV to the field datasets. Fig. 11 shows the prediction for calanoida

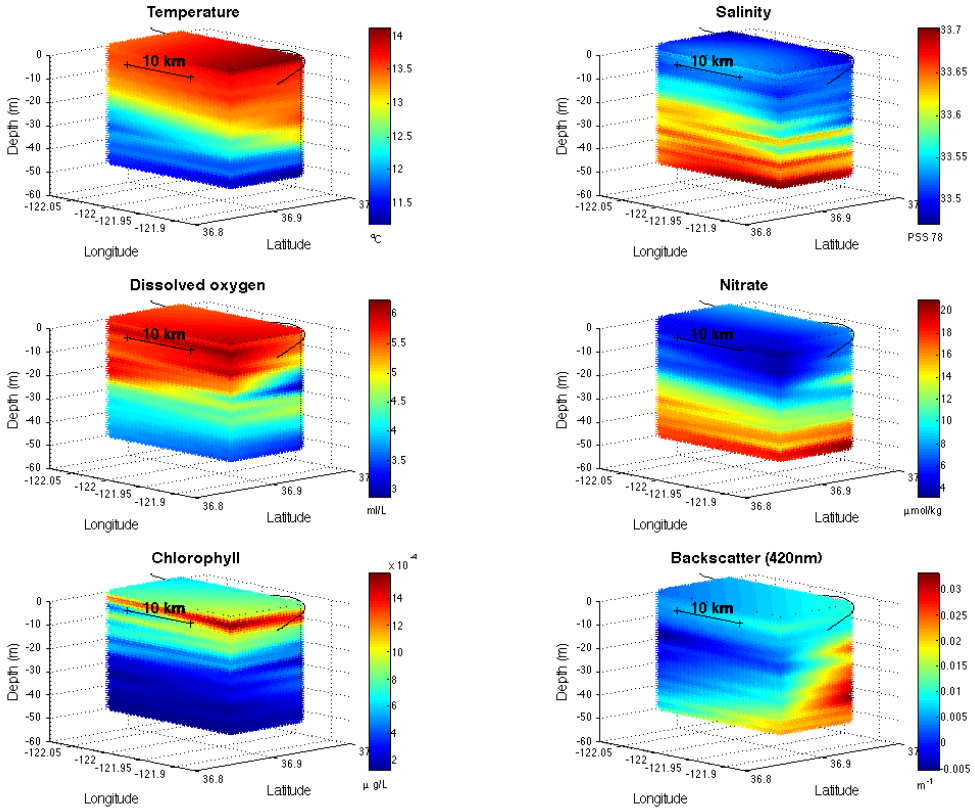

Fig. 10: Prediction of the six environmental parameters within a regular grid in Monterey Bay (looking north west). Depth axis is exaggerated for clarity.

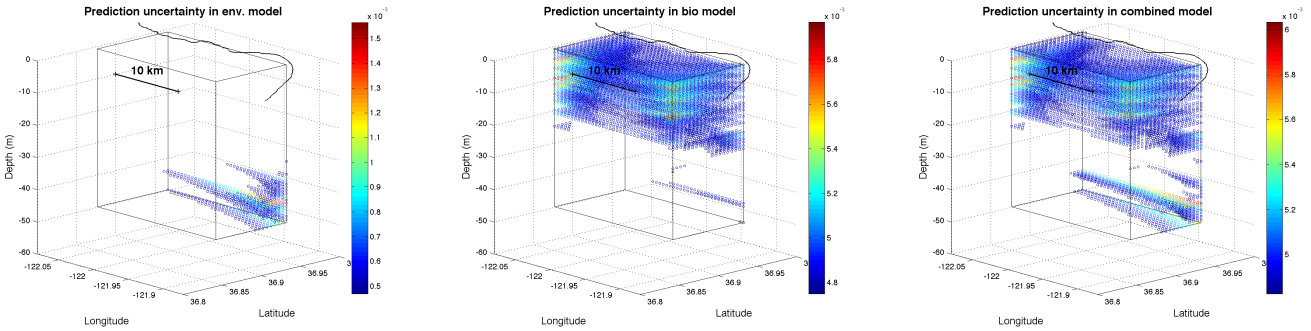

(a) Prediction uncertainty for the spatial environmental model only. The bottom-right region with high uncertainty was not covered by the AUV survey.

(b) Prediction uncertainty for the organism abundance model only.

(c) Combined prediction uncertainty using the unscented transform.

Fig. 12: Uncertainty maps for organism abundance estimates in Monterey Bay (looking north west). For clarity, depth axis is exaggerated and only points with high relative uncertainty are shown.

abundance and the associated prediction variance using the combined model. Fig. 12 shows uncertainty associated

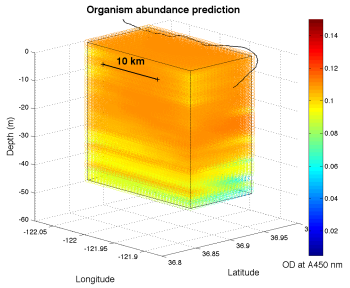

Fig. 11: Predicted spatial distribution of calanoida in Monterey Bay in a regular grid (looking north west). Depth axis is exaggerated for clarity.

with the organism abundance prediction in Monterey Bay. The first map (Fig. 12a shows only the uncertainty in the prediction of the environmental parameters from the AUV survey data. We restrict plot to only points with relatively high uncertainty. The region of the regular grid at the bottom without any data from the AUV survey has high uncertainty. The uncertainty map for only the biological model prediction is shown in Fig. 12b. The uncertainty for this model is high for the top region of the grid. If a sampling policy uses either of the two models for acquisition of new samples, samples will miss out one of the two regions of high uncertainty. Fig. 12c shows the uncertainty from both environmental and organism abundance models fused using the unscented transform.

VII. Conclusions and Future Work

In this work, we presented an approach to predict organism abundance in the coastal ocean from *in situ* measurements of environmental parameters by an AUV. Our approach uses a combination of two models, the first capturing the distribution of environmental parameters, and the second capturing the relationship between environmental parameters and biological organism abundance. To estimate the organism abundance at geographical locations, we use hierarchical probabilistic regression. The fused model ensures that predictions of organism abundance from environmental parameters also consider the uncertainty in the environmental parameter estimates. The unscented transform was used to propagate errors in the output of the the environmental through the organism abundance model. The goal of this work is two fold. We bring to the attention the sources of errors when building hierarchical models of phenomena of interest. This is particularly relevant in the environmental sciences where often the property of interest is predicted using environmental parameters. Following from this, we propose the use of uncertainty that considers all models, and not any one individually when developing adaptive sampling strategies. To demonstrate this, we carry out simulation studies to compared the performance of greedy variance based sampling using different forms of uncertainties. Historical field data was used to present a case study where we estimated spatial distribution of various species of plankton. In the computed prediction uncertainty maps, we highlight the role of errors from each model. We are now implementing our sampling strategy onboard an AUV for targeted sampling of plankton.

VIII. Acknowledgments

This work was supported in part by NOAA under award NA11NOS4780052 and NSF under award CCF-0120778. MBARI authors are funded by a block grant from the David and Lucile Packard Foundation.

References

[1] M. Fox, D. Long, F. Py, K. Rajan, and J. Ryan, "In Situ Analysis for Intelligent Control," in *Proceedings of IEEE/OES OCEANS Conference*. IEEE, 2007.

[2] J. P. Ryan, S. Johnson, A. Sherman, K. Rajan, F. Py, H. Thomas, J. Harvey, L. Bird, J. Paduan, and R. Vrijenhoek, "Mobile autonomous process sampling within coastal ocean observing systems," *Limnology & Oceanograhy: Methods*, vol. 8, pp. 394–402, 2010.

[3] Y. Zhang, R. McEwen, J. Ryan, and J. Bellingham, "Design and Tests of an Adaptive Triggering Method for Capturing Peak Samples in a Thin Phytoplankton Layer by an Autonomous Underwater Vehicle," *IEEE Journal of Oceanic Engineering*, vol. 35, no. 4, pp. 785–796, Oct 2010.

[4] J. P. Ryan, Y. Zhang, H. Thomas, E. Rienecker, R. Nelson, and S. Cummings, "A High-Resolution Survey of a Deep Hydrocarbon Plume in the Gulf of Mexico During the 2010 Macondo Blowout," *Geophysical Monograph Series*, vol. 195, pp. 63–75, 2011.

[5] C. Scholin, G. Doucette, S. Jensen, B. Roman, D. Pargett, R. M. III, C. Preston, W. Jones, J. Feldman, C. Everlove, A. Harris, N. Alvarado, E. Massion, J. Birch, D. Greenfield, R. Vrijenhoek, C. Mikulski, and K. Jones, "Remote detection of marine microbes, small invertebrates, harmful algae, and biotoxins using the Environmental Sample Processor (ESP)," *Oceanography*, vol. 22, pp. 158–167, 2009.

[6] C. R. Anderson, D. A. Siegel, R. M. Kudela, and M. A. Brzezinskic, "Empirical models of toxigenic pseudo-nitzschia blooms: Potential use as a remote detection tool in the santa barbara channel," *Harmful Algae*, vol. 8, pp. 478âĂŞ–492, 2009.

[7] S. J. Phillips, M. Dudík, and R. E. Schapire, "A Maximum Entropy Approach to Species Distribution Modeling," in *Proceedings of the twenty-first international conference on Machine learning*, ser. ICML '04, 2004, pp. 83–.

[8] J. Binney, A. Krause, and G. S. Sukhatme, "Informative Path Planning for an Autonomous Underwater Vehicle," in *ICRA*. IEEE, 2010, pp. 4791–4796.

[9] A. Krause and C. Guestrin, "Nonmyopic active learning of Gaussian Processes: an exploration-exploitation approach," in *Intnl. Conf. on Machine Learning*. New York, NY, USA: ACM, 2007, pp. 449–456.

[10] K. H. Low, "Multi-Robot Adaptive Exploration and Mapping for Environmental Sensing Applications," Ph.D. dissertation, Carnegie Melon University, 2009.

[11] A. Olaya, F. Py, J. Das, and K. Rajan, "An On-line Utility based Approach for Sampling Dynamic Ocean Fields," *IEEE Journal of Ocean Engineering*, 2011.

[12] A. Girard, C. E. Rasmussen, J. Quinonero-Candela, J. Q. N. Candela, M. Modelling, and R. Murray-smith, "Gaussian Process Priors with Uncertain Inputs - Application to Multiple-Step Ahead Time Series Forecasting," in *Advances in Neural Information Processing Systems*. MIT Press, 2003, pp. 529–536.

[13] J. Ko, D. J. Klein, D. Fox, and D. Haehnel, "GP-UKF: Unscented Kalman Filters with Gaussian Process Prediction and Observation Models," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, October 29 - November 2, 2007, Sheraton Hotel and Marina, San Diego, California, USA*. IEEE, 2007, pp. 1901–1907.

[14] C. E. Rasmussen and C. Williams, *Gaussian processes for machine learning*. Cambridge: MIT Press, 2006.

[15] S. Julier and J. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," in *Proc. of AeroSense: The 11th Int. Symp. on Aerospace/Defence Sensing, Simulation and Controls*, 1997.

[16] C. E. Rasmussen and H. Nickisch, "Gaussian Processes for Machine Learning (GPML) Toolbox," *J. Mach. Learn. Res.*, vol. 11, pp. 3011–3015, Dec. 2010. [Online]. Available: http://dl.acm.org/citation.cfm?id=1756006.1953029