

Learning Uncertainty Models for Reliable Operation of Autonomous Underwater Vehicles

Geoffrey A. Hollinger, Arvind A. Pereira, and Gaurav S. Sukhatme

Abstract—We discuss the problem of learning uncertainty models of ocean processes to assist in the operation of Autonomous Underwater Vehicles (AUVs) in the ocean. We focus on the prediction of ocean currents, which have significant effect on the navigation of AUVs. Available models provide accurate prediction of ocean currents, but they typically do not provide confidence estimates of these predictions. We propose augmenting existing prediction methods with variance measures based on Gaussian Process (GP) regression. We show that commonly used measures of variance in GPs do not accurately reflect errors in ocean current prediction, and we propose an alternative uncertainty measure based on interpolation variance. We integrate these measures of uncertainty into a probabilistic planner running on an AUV during a field deployment in the Southern California Bight. Our experiments demonstrate that the proposed uncertainty measures improve the safety and reliability of AUVs operating in the coastal ocean.

I. INTRODUCTION

Sensing and monitoring of ocean processes is becoming increasingly important for science and exploration. Autonomous Underwater Vehicles (AUVs) are now available that are capable of long-term monitoring tasks over several weeks or months [1], [2]. As the autonomy capabilities of these vehicles improves and the length of their missions increases, there is growing concern for the safety and reliability of operation. For instance, large ships provide potential dangers for AUVs during surfacing, and strong ocean currents increase the risk of running into shallow water and aborting the mission. Improved predictions of uncertainty in the modeling of ocean currents has the potential to reduce these risks and increase the functionality of AUVs operating in the ocean.

Predictions of ocean currents are publicly available and are currently in widespread use by ocean scientists. Such predictions utilize forward simulations of prior data, and they incorporate large-scale measurements from satellite and HF-radar. We focus on the Regional Ocean Monitoring System (ROMS) predictions of ocean currents [3] (shown in Figure 1 for the Southern California Bight region). These predictions are provided daily and give forecasts up to two days in advance.

Available models typically do not provide estimates of the accuracy of the prediction, which have the potential to




Fig. 1. Ocean currents in the Southern California Bight region predicted by the Regional Ocean Modeling System (ROMS) [3]. We propose data-driven methods to develop confidence estimates on these predictions, and we integrate the resulting estimates into a probabilistic planner to improve the safety of autonomous underwater vehicle operation.

improve AUV operation. We focus on developing such confidence estimates based on novel variance measures related to Gaussian Process (GP) regression. We propose the use of the *interpolation variance* as an alternative to the standard GP variance. Interpolation variance was originally introduced in the geostatistic community [4], [5] and was recently applied to robotic sampling tasks [6]. We integrate this confidence measure into the action model of a probabilistic planner, and we test the proposed approach both in simulation and in experimental trials in the Southern California Bight.

The key novelties of this paper are (1) the application of interpolation variance to the prediction of ocean currents, (2) the integration of interpolation variance into probabilistic planning, and (3) the field implementation of these techniques on an AUV. We first survey prior work (Section II). We next discuss methods and algorithms for improving the prediction of ocean currents, and we integrate these predictions into AUV navigation (Section III). We then present both simulated and field experiments demonstrating the performance of our approach (Section IV). We conclude by discussing future work and future experiments (Section V).

II. RELATED WORK

Predictions of ocean currents have previously been used by several research groups to improve the navigation capabilities of autonomous vehicles and improve the safety of their operation [7], [8]. Lower bounds on navigation error have

G. Hollinger, A. Pereira, and G. Sukhatme are with the Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, {gahollin, menezesp, gaurav}@usc.edu.

This research has been funded in part by the following grants: ONR N00014-09-1-0700, ONR N00014-07-1-00738, ONR N00014-06-10043, NSF 0831728, NSF CCR-0120778 and NSF CNS-1035866.

also been derived to estimate the path following performance of underwater vehicles [9]. While these prior works provide a basis for navigation of AUVs in the ocean, they are missing a principled analysis of the uncertainty of predictions. We develop the necessary tools to utilize these uncertainty predictions and improve path planning methods.

Prior work has examined the development of confidence measures for various ocean processes using straightforward statistical tools [10] as well as more sophisticated Bayesian models [11]. In [12], a rigorous computational method for the quantification, prediction, and estimation of uncertainties called Error Subspace Statistical Estimation (ESSE) is described. ESSE has components of time-varying basis functions, multi-scale initializations, and stochastic ensemble predictions. ESSE was later combined with principled filtering techniques to derive the GMM-DO filter, a general data-assimilation technique [13]. Both ESSE and GMM-DO require knowledge of the approximate governing equations for the time-evolving error covariance bases, which is not the case for our approach.

We propose utilizing Gaussian Process (GP) regression to provide confidence measures on the prediction of ocean processes. GPs have been used successfully to improve the accuracy of such tasks as large-scale terrain modeling [14], underwater habitat mapping [15], scientific planetary surveying [16], and navigation under uncertainty [17]. We demonstrate that standard GP variance predictions do not capture the uncertainty inherent in prediction of ocean processes. Interpolation variance provides an alternative measure of uncertainty that also incorporates variability of the prediction into the variance. Interpolation variance was originally introduced for the estimation of geostatistical data [4], [5] and was recently applied as a method for guiding adaptive sampling by mobile robots [6]. To our knowledge, we are the first to apply interpolation variance to the prediction of ocean currents and the first to integrate it into the action model of a probabilistic planner. The preliminary modeling algorithms in this paper were presented in our prior workshop paper [18].

III. METHODS AND ALGORITHMS

The goal is to provide a principled estimate of uncertainty for predictions of large-scale ocean processes. We propose augmenting available predictions through the use of Gaussian Process learning methods. We focus on the prediction of ocean surface currents due to the availability of accurate HF-Radar measurements for these values.

The ocean currents at a given latitude lat , longitude lon , and time t are represented by a vector $c(lat, lon, t) = (u, v)$, where u and v are the components of the currents along the latitude and longitude axes respectively. The positions are discretized based on the latitude and longitude, and time is discretized into hours. At a given time T , we are given historical data for times $t = \{T-1, T-2, \dots\}$ back several months or years. We are also given predictions from the ROMS data for $t = \{T+1, T+2, \dots, T+48\}$ (two days in the future). Given this data, we want to provide better

predictions for the points of time in the future as well as confidence bounds for these predictions.

A. Gaussian Process regression

We propose modeling ocean currents using non-parametric Bayesian regression in the form of Gaussian Processes (GPs) [19]. A GP models a noisy process $z_i = f(\mathbf{x}_i) + \varepsilon$, where $z_i \in \mathbb{R}$, $\mathbf{x}_i \in \mathbb{R}^d$, and ε is Gaussian noise.

We are given some data of the form $D = [(\mathbf{x}_1, z_1), (\mathbf{x}_2, z_2), \dots, (\mathbf{x}_n, z_n)]$, where \mathbf{x}_i represents a vector of latitude, longitude, and time values for a data point i , and z_i represents either the u or v component of the currents at that point and time. We note that this formulation decouples the prediction of u and v , an assumption that could be relaxed in future work (e.g., using the techniques in [20]). We refer to the $d \times n$ matrix of \mathbf{x}_i vectors as \mathbf{X} and the vector of z_i values as \mathbf{z} .

To fully define a GP, we must choose a covariance function that relates the points in \mathbf{X} to each other. We expect that data points that are close to each other in time and space will have high correlations, which provides a smoothing effect on the data. In addition, there is a periodic effect that arises due to tidal processes, which creates a correlation between points that are separated by 12 hours, 24 hours, etc. To capture both the spatial correlations and the periodic correlations, we apply the following kernel based on the squared exponential:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp[-w_{lat}(\mathbf{x}_i^{lat} - \mathbf{x}_j^{lat})^2 - w_{lon}(\mathbf{x}_i^{lon} - \mathbf{x}_j^{lon})^2 - w_t(\mathbf{x}_i^t - \mathbf{x}_j^t)^2 - w_p \sin^2(\pi(\mathbf{x}_i^t - \mathbf{x}_j^t)/12)]. \quad (1)$$

The hyperparameter σ_f^2 represents the process noise, and the hyperparameters w_{lat} , w_{lon} , w_t , and w_p represent weighting for the latitude, longitude, time, and periodic correlations. Having defined the kernel, combining the covariance values for all points into an $n \times n$ matrix \mathbf{K} and adding a Gaussian observation noise hyperparameter σ_n^2 yields $\mathbf{K}_z = \mathbf{K} + \sigma_n^2 \mathbf{I}$.

The following equation predicts the mean function value (u or v value) $\mu(\mathbf{x}_*)$ and variance $\mathbb{V}_{gp}(\mathbf{x}_*)$ at a selected point \mathbf{x}_* given the historical and prediction training data:

$$\mu(\mathbf{x}_*) = \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{z}, \quad (2)$$

$$\mathbb{V}_{gp}(\mathbf{x}_*) = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (3)$$

where \mathbf{k}_* is the covariance vector between the selected point \mathbf{x}_* and the training inputs \mathbf{X} . This model gives a mean and variance for a particular latitude, longitude, and future time. An important aspect of this model is the derivation of a variance, which provides a measure of confidence of each prediction based on the correlations of the points used to predict it.

The Gaussian Process variance described above gives an estimate of the uncertainty of a prediction based on the data sparsity around that point and the estimated hyperparameters. However, this estimate does not take into account variability

of the ocean currents around that point, which is an important predictor of measurement confidence. To improve on the GP variance uncertainty predictions, we propose using a method based on the interpolation variance [4], [5], [6]. Given that a GP has already been learned, the interpolation variance can be estimated as:

$$\mathbb{V}_{iv}(\mathbf{x}_*) = \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} (\mathbf{z} - \mu(\mathbf{x}_*))^T (\mathbf{z} - \mu(\mathbf{x}_*)). \quad (4)$$

The interpolation variance incorporates correlations between the data learned using the GP framework as well as variability of correlated data to determine the predicted variance. This estimate of variance provides a richer representation that accounts for both data sparsity and data variability.

B. Local approximation

The modeling technique proposed above can be used for small datasets. However, the ROMS datasets correspond to vast periods of time and a large portion of space. The computation time required by the GP scales $O(n^3)$ in the number of data points, which makes it infeasible for any large dataset more than approximately 2000 points. As an alternative, we propose estimating the predictions using a subset of the data corresponding to the points we expect to be most correlated. To choose these points, we store the data in a KD-tree, where the relative weighting of space and time are adjusted to fit the values in the kernel (e.g., the periodic correlations are preserved). Such techniques have been applied successfully for terrain modeling [14], though we believe their application to time-series data is novel.

The estimation of the kernel hyperparameters $\theta = (\sigma_f, \sigma_n, w_{lat}, w_{lon}, w_t, w_p)$ in large-scale Gaussian Processes is also a challenging task. The standard method is maximizing the likelihood of the measurements given the data and the hyperparameters [19]:

$$\log p(\mathbf{z}|\mathbf{X}, \theta) = -\frac{1}{2} \mathbf{z}^T \mathbf{K}_z^{-1} \mathbf{z} - \frac{1}{2} \log |\mathbf{K}_z| - \frac{n}{2} \log 2\pi, \quad (5)$$

where $\mathbf{K}_z = \mathbf{K} + \sigma_n^2 \mathbf{I}$. This likelihood can be maximized using conjugate gradient optimization for small and moderately sized datasets. For the large ROMS dataset, we learn separate hyperparameters for each day. This approach keeps down the size of the dataset and also models variations that occur between days.

C. Probabilistic Planning

The methods described above can be incorporated into path planners for AUVs to improve the safety and reliability of operation. Path planners that incorporate data from ocean currents clearly stand to benefit from improvements in the accuracy of ocean current predictions. In addition, confidence estimates are useful both for planners that reason probabilistically as well as for planners that reason about worst-case instances. In particular, we are interested in improving the operation of AUVs by avoiding areas with high probability of encountering a dangerous ship. Figure 2 shows a risk




Fig. 2. Risk map built for the Southern California coastal ocean that shows areas of high collision risk in red, high shipping activity in light blue, and low shipping activity in dark blue. The bright red areas are Catalina island and the California coast. Our goal is to plan risk-aware paths for AUVs that avoid the shipping lanes. To accomplish this goal, we utilize predictions of the ocean currents and the uncertainty of those predictions.

map built from historical data showing areas of the Southern California coast with high level of shipping activity. We now examine how the proposed GP prediction methods could be useful for improving the paths generated by such planners.

ROMS provides forecasts up to two days in advance, which allows us to develop action models for vehicles that utilize these ocean predictions. We evaluate two risk-averse path planners for AUVs, which utilize probabilistic action models based upon ocean predictions - (1) a Markov Decision Process (MDP) planner where we assume that the stochasticity in the ROMS predictions is uniform throughout the map and (2) a Markov Decision Process which uses the spatio-temporal interpolation variance estimates from the Gaussian Process (referred to as IV-MDP).

The transition model we use relies upon a discrete grid of states for the AUV in the ocean, with possible actions being the choice of moving from one state to any of its 8-connected neighbors. The grid is built by performing a number of simulations (> 30) for the AUV traversing each pair of grid locations under the influence of ocean currents with different starting times for each simulation (to generalize for temporal variations) and an additive Gaussian noise proportional to the predicted noise. For the naive MDP, this noise distribution is constant through the entire map, while for the IV-MDP the noise variance is based upon the interpolation variance measures. The distribution of surfacing locations for each set of trials gives us an estimate of the transition function for that action. This transition function $T(s''|s, a(s, s'))$ describes the probability of ending up in state s'' , given we choose to take action a (state s to state s').

When performing simulations for generating the transition model of the Fixed-noise MDP, we choose a representative noise for the currents with a standard-deviation between 0.1 m/s to 0.001 m/s . The same noise model was used for both the easting and northing current components. The following equations describe how the simulated currents are obtained, where simulated noise is drawn from Gaussian

distributions $N_u \sim \mathcal{N}(0, \sigma_u^2)$ and $N_v \sim \mathcal{N}(0, \sigma_v^2)$:

$$u_{sim}(x, y, t) = u_{pred}(x, y, t) + N_u \quad (6)$$

$$v_{sim}(x, y, t) = v_{pred}(x, y, t) + N_v \quad (7)$$

In the IV-MDP we use the maximum interpolation variance among the states s and s' of each transition edge, $e(s, s')$, over every time interval $t \in [t_1, t_2]$ to generate the transition model for the MDP. To compute the maximum variance for the generation of the transition model for a particular transition edge, we use Equations 8 and 9. Each transition model is then computed by conducting (> 30) simulations of transitions between each pair of states where the currents used for the simulation are drawn from a Gaussian $\mathcal{N}(0, \sigma^2)$, where the variance standard deviation σ is σ_u and σ_v for the easting and northing current components using Equations 6 and 7 respectively:

$$\sigma_u^2 = \max_{t \in [t_1, t_2]} \{ \sigma_u^2(s, t), \sigma_u^2(s', t) \}, \quad (8)$$

$$\sigma_v^2 = \max_{t \in [t_1, t_2]} \{ \sigma_v^2(s, t), \sigma_v^2(s', t) \}. \quad (9)$$

Both the MDP planners use their respective transition functions to perform Bellman updates in a planning graph, where the rewards $R(s)$ are given by negative risk of collision with land or ships (obtained from the risk map). The Bellman updates to compute the utility values U , if value iteration is used to find the optimal policy, are described by:

$$U_{i+1}(s) \leftarrow -R(s) + \underset{a}{\operatorname{argmax}} \sum_{s''} T(s''|s, a(s, s'')) U_i(s''). \quad (10)$$

With improved confidence measures describing the amount of noise in the ocean current predictions, the transition models used by the IV-MDP are more representative of the true errors than the models using a constant prediction noise value over the entire planning graph. Thus, we would expect the IV-MDP to provide improvements in the execution speed, reliability, and safety of the resulting plan. The experiments in the following section will confirm this trend.

IV. SIMULATIONS AND EXPERIMENTS

We now provide an experimental validation of our approach. The validation is two-fold: (1) We first examine the performance of the interpolation variance for prediction of errors in the ocean current estimates, and (2) we then validate the approach both in simulation and experimentally on an AUV operating in the Southern California Bight.

A. Comparison of uncertainty predictions

The data processing was performed on a single desktop with a 3.2 GHz Intel i7 processor with 9 GB of RAM. The kernel hyperparameters were learned independently for each day using conjugate gradient optimization. Using the KD-tree inference approximation, the GP took approximately 6 minutes to compute all predictions for a single day. These results use 100 points in the space/time KD-tree for the local

GP prediction at each point. There are a total of 2560 locations that are estimated over a 24 hours period for predictions up to two days ahead (total of 123,880 point predictions per day). The proposed method is highly parallelizable, since the prediction for each point can be performed independently once the hyperparameters are learned and the KD-tree is built.

Our methods provide confidence measures based on the underlying variance. In the formulation above, the Gaussian Process variance prediction provides an estimate of error due to data sparsity and measurement noise. Since the ROMS data are available at high resolution across time and space, this value provides a limited measure of confidence that is fairly homogeneous. The alternative method based on the interpolation variance incorporates variability in the surrounding currents as a component of the uncertainty. Figure 3 shows example variance maps for a day in August, 2012 as well as the true prediction error (compared to the next day's nowcast, which would not have been available during planning). The results demonstrate that the interpolation variance provides a much richer measure of uncertainty that qualitatively matches with the true prediction error.

We also compare the correlation coefficients (R-values) between the interpolation variance and the Gaussian Process variance correlated with the true prediction error. If the variance is an accurate representation of uncertainty, we would expect to see a positive correlation between the variance and the true error. Table I shows the R-values for three months in 2012. The interpolation variance shows a positive correlation with the true prediction error, while the Gaussian Process variance shows essentially no correlation.¹

TABLE I
CORRELATION COEFFICIENTS (R-VALUES) FOR GAUSSIAN PROCESS VARIANCE AND INTERPOLATION VARIANCE RELATIVE TO TRUE PREDICTION ERROR (SHOWN SEPARATELY FOR N/S COMPONENT AND E/W COMPONENTS OF THE OCEAN CURRENT VECTORS)

Month (2012)	June	July	August
GP variance R-value (E/W)	-0.0519	-0.0178	-0.0277
Interp. variance R-value (E/W)	0.1383	0.1260	0.1655
GP variance R-value (N/S)	-0.0652	-0.0245	-0.0271
Interp. variance R-value (N/S)	0.1745	0.1400	0.1323

B. Risk-aware AUV planning

We now report results using the two Minimum Risk planners which attempt to minimize the risk of collision with ships and land along the AUV's path. To facilitate testing, we constructed a virtual archipelago by adding additional islands and shipping lanes around St. Catalina island in Southern California (shown in Figure 4). The benefit of this approach is that we are able to test more complex behaviors in planners, such as allowing them to be more goal-directed (and consequently more risk-taking), without increasing the

¹The slight negative correlation for the GP variance is likely due to the high predicted variances in the bordering edges of the dataset, which do not correspond to increased errors in the prediction.

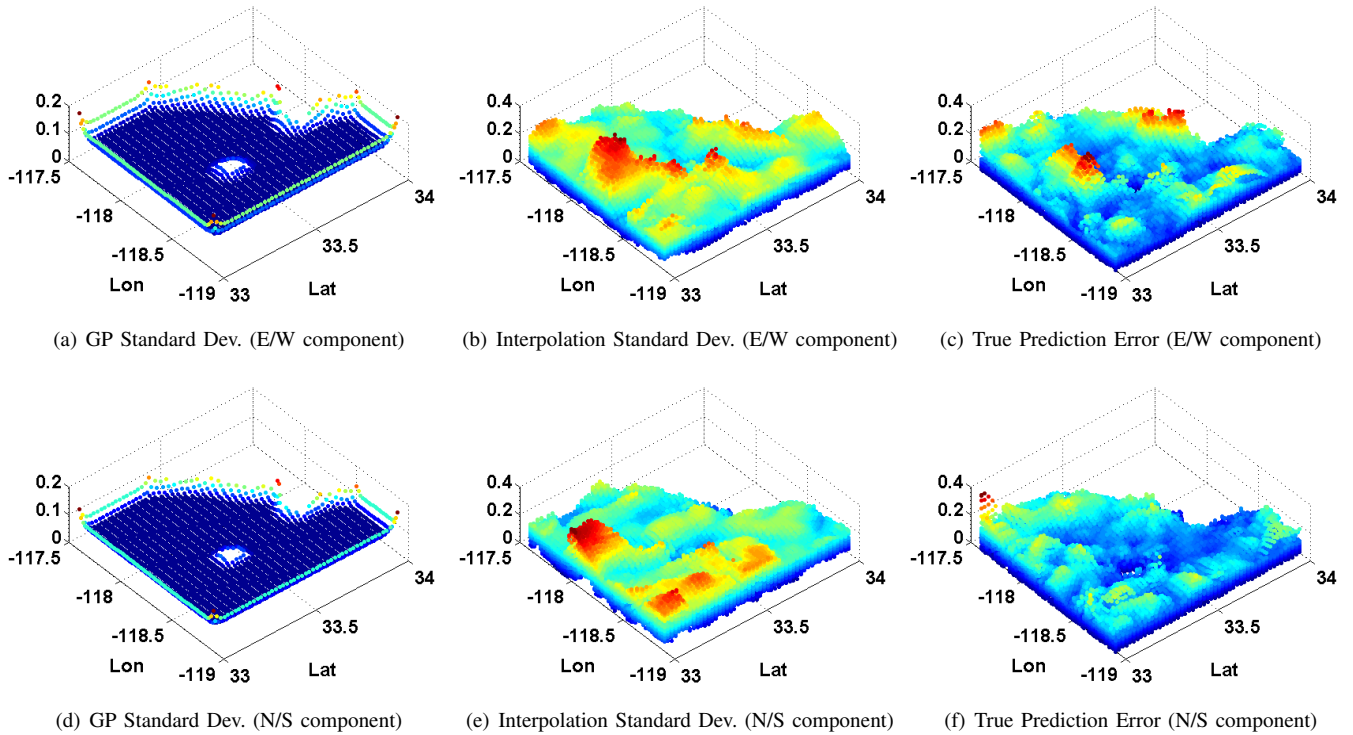


Fig. 3. Comparison of the Gaussian Process variance and the interpolation variance measures of uncertainty for the prediction of ocean currents in the Southern California Bight on August 7, 2012. The interpolation variance accounts for variability in the ocean currents when making its prediction of uncertainty, and it shows a stronger correlation with the true prediction error.

real risk of collision with land or ships. The drawback is that the virtual islands allow for strong currents that run through (and consequently into) islands, which, along with tight corridors, makes this a particularly challenging scenario for risk-aware AUV planning.

We now present results for risk-aware planners that utilize the predictions of ocean currents described in the previous section. We performed simulations on a regular grid (2.5 km spacing between adjacent nodes) overlaid on the risk map shown in Figure 4. We compared the performance of two planners that seek to minimize the risk of collision between the AUV and obstacles in the region: i) MDP - Markov Decision Process, ii) IV-MDP - Markov Decision Process using the interpolation variance for transition models (both ROMS-aware planners described in section III-C). The AUV is modeled as a Slocum glider (see next section for more detail), a slow-moving (approximately 0.3 m/s) buoyancy-controlled vehicle subject to changes in the ocean currents.

Stochastic planners (eg., MDPs) are especially useful in high variability situations where a reasonable probabilistic model can describe the expected behavior of the system. The transition model for the naive MDP was created using an empirically determined prediction noise of $\sigma = 0.01$ m/s, and the transition model for the IV-MDP uses the interpolation variances described in Section III. The ocean current predictions during the experimental trial period were generally flowing from southeast toward the northwest, usually at speeds comparable to those of the AUV.

We ran simulations between two locations at the eastern and western ends of the planning graph, and we present results from a set of trials both against and along with the current flow. The ROMS two-day forecasts were used to generate transition models, which are used by the two MDP planners, while we use the assimilated forecast for each day (which gives us the nearest current values to ground-truth for each day) for simulations. A simulated mission is aborted when the AUV moved too close to a virtual island for safe operation. Successful trials are those that reach the goal without violating this safety condition. Table II and Table III show the results from the simulations. The IV-MDP method outperforms the MDP in terms of successes, time, and average risk along the path. In the case of moving against the currents, the IV-MDP achieves more than a 40% reduction in average risk along the path. In the case of moving with the currents, there is not as much reduction in risk, but the success rate jumps from 50% to 80% by using the IV-MDP.

C. Experimental Trials

We performed several trials involving two Slocum gliders (see [21]) in the same planning region used for the simulated trials. Slocum gliders are slow-moving AUVs without propellers that use buoyancy control to adjust their depth. Horizontal static wings and pitch-control to an angle of 26 degrees while alternating between diving and climbing allows the glider to generate horizontal movement. This efficient use of energy gives gliders significant range and

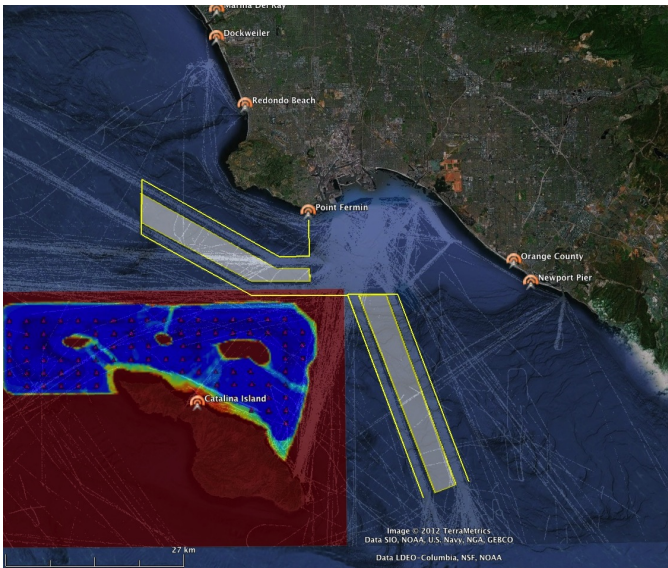


Fig. 4. Location of the planning graph used for the MDP planners relative to the real shipping lanes and Southern California coastline (to the north and east) and St. Catalina island to the south. Magenta triangles indicate nodes in the planning graph. Red denotes land which has the highest risk of collision, while dark blue regions denote low collision risk. Warmer colored regions such as yellow or cyan indicate locations with intermediate risk such as shipping lanes or high-traffic areas near shore.

TABLE II

RESULTS FROM SIMULATIONS OF A GLIDER AUV FLYING AGAINST CURRENTS, WITH DIFFERENT START TIMES (BETWEEN AUG 17 AND AUG 24, 2012)

Planner	Fixed-noise MDP	IV-MDP
Success	5/10	5/10
Avg. Time	132.9 hrs	81.72 hrs
Avg. Surfacings	49.5	28.4
Avg. Risk	15.546	8.737

longevity at sea at the expense of speed (maximum speed of operation is typically less than 1 km/hr or 0.3 m/s). The flip-side of having lower velocities means the vehicle has to contend with ocean currents, which are often of the same order of magnitude of the glider's nominal horizontal speed.

In the experiments, one glider ran the Fixed-noise MDP planner while the other glider ran the Interpolation-variance MDP. Both gliders were started from approximately the same start location. At each surfacing, we use transition models which have been pre-computed for the latest ROMS data, solve the MDP with the latest position measurement from the glider, and then transfer the resulting policy back to the glider in the form of a new mission file with a single next waypoint to go to. This process was repeated for each waypoint until the gliders either reached the goal or the mission was aborted due to safety concerns.

An example experimental trial is shown in Figure 5. We observe that the MDP planner was not goal-directed initially and kept the glider in a holding pattern in the relatively safe location near the start for nearly 20 hours. The IV-MDP was less conservative and kept making progress toward the

TABLE III

RESULTS FROM SIMULATIONS OF A GLIDER AUV FLYING WITH CURRENTS, WITH DIFFERENT START TIMES (BETWEEN AUG 17 AND AUG 24, 2012)

Planner	Fixed-noise MDP	IV-MDP
Success	5/10	8/10
Avg. Time	34.24 hrs	32.73 hrs
Avg. Surfacings	13.5	12.1
Avg. Risk	4.278	3.875

goal, which it reached 51 hours later. The Fixed-noise MDP reached the goal more than a day later due to the time it spent oscillating back and forth initially. These observations are consistent with our simulation results where the IV-MDP was able to act less conservatively than the Fixed-noise MDP due to its additional knowledge of the predicted variances.

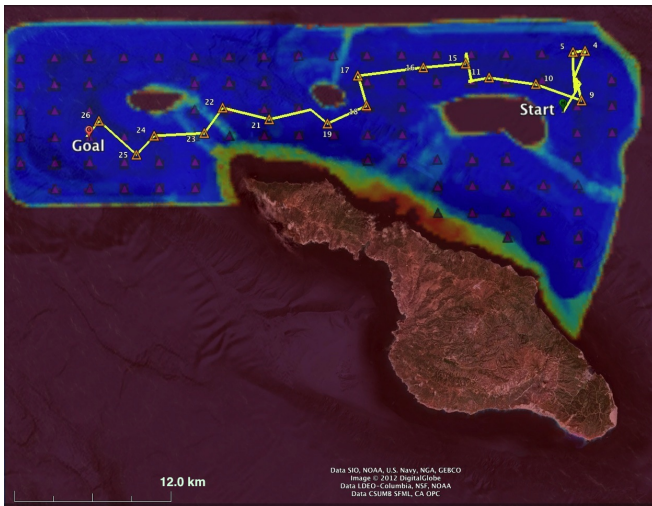
The simulations in the previous section showed that the additional knowledge of the errors in the ocean current estimates enable the IV-MDP to produce less conservative policies than those based on an MDP planner with a fixed prediction noise, generally helping it to produce more direct paths to the goal. This goal-directed behavior makes the IV-MDP have lower cumulative risk at surface, fewer hops to get to the goal, and lower average time to the goal. Our experimental trials with two gliders at sea (simultaneously executing each planner) confirm the results in simulation and further demonstrate that the benefit of uncertainty-aware MDP planners for generating feasible and low-risk paths.

V. CONCLUSIONS AND FUTURE WORK


This paper has shown that it is possible to generate error estimates for ocean current predictions using data-driven machine learning methods. We have proposed the application of time-series Gaussian Process regression combined with a measure of interpolation variance, and we have shown that this technique provides improvements over commonly used measures of uncertainty. In addition, we have integrated our estimates into risk-aware MDP planning methods for AUVs, and we have shown that we can improve the safety of autonomous operation in the coastal ocean by utilizing models of ocean currents.

We believe these results show significant promise for the use of probabilistic planning and statistical machine learning methods to generate confidence bounds and improve the reliability of operation for robotic vehicles. Planning for autonomous aerial vehicles operating in urban environments and wooded areas would also benefit from similar risk-aware planning with uncertainty. In these scenarios, additional uncertainty would arise from the vehicle's dynamics as well as from disturbances in the environment. However, many of the planning methods and uncertainty measures we have proposed will likely be applicable in these domains as well.

There are a number of improvements that are avenues for future work. In the short term, we plan to examine more sophisticated methods for automatically learning the hyper-parameters based on large amounts of historical data as well



(a) Fixed variance MDP



(b) Interpolation variance MDP

Fig. 5. A representative experimental trial of the planner operating with a fixed variance (left) and with the interpolation variance (right). The gliders started on Aug 17, 2012 at 00:20 am PST. Despite facing strong currents, the interpolation variance confidence measures produce a less conservative policy, which moves the glider more quickly toward the goal. The glider using the interpolation variance reached the goal at 3:21 am PST on Aug 19, well over a day earlier than the other glider running the fixed-variance MDP, which reached the goal on Aug 20, 2012 at 8:25am PST.

as alternative methods for generating confidence measures of the predictions. Longer-term future work includes applications of such techniques to other environmental processes, as well as exploring the life-long learning capabilities when presented with increasing amounts of historical data. The investigation of alternative uncertainty representations opens a rich area for future research that will allow for more informed planning and scheduling of autonomous vehicles.

REFERENCES

- [1] N. E. Leonard, D. A. Paley, R. E. Davis, D. M. Fratantoni, F. Lekien, and F. Zhang, "Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in Monterey Bay," *J. Field Robotics*, vol. 27, no. 6, pp. 718–740, 2010.
- [2] D. A. Paley, F. Zhang, and N. E. Leonard, "Cooperative control for ocean sampling: The glider coordinated control system," *IEEE Trans. Control Systems Technology*, vol. 6, no. 4, 2008.
- [3] A. F. Shchepetkin and J. C. McWilliams, "The Regional Oceanic Modeling System (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model," *Ocean Modelling*, vol. 9, no. 4, pp. 347–404, 2005.
- [4] J. Yamamoto, "An alternative measure of the reliability of ordinary kriging estimates," *Mathematical Geology*, vol. 32, no. 4, pp. 489–509, 2000.
- [5] J. Yamamoto and M. M. da Rocha, "Properties and applications of the interpolation variance associated with ordinary kriging estimates," in *Proc. Int. Symp. Spatial Accuracy Assessment in Natural Resources and Environmental Sciences*, 2008, pp. 70–77.
- [6] Y.-H. Kim, D. Shell, C. Ho, and S. Saripalli, "Spatial interpolation for robotic sampling: Uncertainty with two models of variance," in *Proc. Int. Symp. Experimental Robotics*, Quebec City, Canada, 2012.
- [7] C. Fernández-Perdomo, J. Cabrera-Gómez, D. Hernández-Sosa, J. Isern-González, A. C. Domínguez-Brito, A. Redondo, J. Coca, A. G. Ramos, E. A. Fanjul, and M. García, "Path planning for gliders using Regional Ocean Models: Application of pinzón path planner with the ESEOAT model and the RU27 trans-Atlantic flight data," in *Proc. IEEE OCEANS Conf.*, 2010, pp. 1–10.
- [8] A. M. Pereira, J. Binney, B. H. Jones, M. Ragan, and G. S. Sukhatme, "Toward risk aware mission planning for autonomous underwater vehicles," in *IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2011, pp. 3147–3153.
- [9] K. Szwaykowska and F. Zhang, "A lower bound on navigation error for marine robots guided by ocean circulation models," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2011, pp. 3584–3588.
- [10] C. J. Willmott, R. E. Davis, J. J. Feddema, K. M. Klink, D. R. Legates, C. M. Rowe, S. G. Ackleson, and J. O'Donnell, "Statistics for the evaluation and comparison of models," *Journal of Geophysical Research*, vol. 90, no. C12, pp. 8995–9005, 1985.
- [11] M. C. Kerman, W. Jiang, A. F. Blumberg, and S. E. Buttrely, "A comparison of robust metamodels for the uncertainty quantification (UQ) of New York Harbor oceanographic data," *Journal of Operational Oceanography*, vol. 1, no. 2, pp. 3–13, 2008.
- [12] P. F. Lermusiaux, "Uncertainty estimation and prediction for interdisciplinary ocean dynamics," *Journal of Computational Physics*, vol. 217, pp. 176–199, 2006.
- [13] T. Sondergaard, "Data assimilation with gaussian mixture models using the dynamically orthogonal field equations," Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, Nov. 2011.
- [14] S. Vasudevan, F. T. Ramos, E. W. Nettleton, and H. F. Durrant-Whyte, "Gaussian process modeling of large scale terrain," *J. Field Robotics*, vol. 26, no. 10, pp. 812–840, 2009.
- [15] P. Rigby, O. Pizarro, and S. B. Williams, "Toward adaptive benthic habitat mapping using gaussian process classification," *Journal of Field Robotics*, vol. 27, no. 6, pp. 741–758, 2010.
- [16] D. Thompson and D. Wettergreen, "Intelligent maps for autonomous kilometer-scale science survey," in *Proc. Int. Symp. Artificial Intelligence, Robotics and Automation in Space*, Los Angeles, CA, 2008.
- [17] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. Castellanos, "Active policy learning for robot planning and exploration under uncertainty," in *Proc. Robotics: Science and Systems*, Atlanta, GA, 2007.
- [18] G. Hollinger, A. Pereira, V. Ortenzi, and G. Sukhatme, "Towards improved prediction of ocean processes using statistical machine learning," in *Proc. Robotics: Science and Systems Workshop on Robotics for Environmental Monitoring*, Sydney, Australia, 2012.
- [19] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [20] M. A. Álvarez and N. D. Lawrence, "Computationally efficient convolved multiple output gaussian processes," *Journal of Machine Learning Research*, vol. 12, pp. 1425–1466, 2011.
- [21] C. Jones, E. Creed, S. Glenn, J. Kerfoot, J. Kohut, C. Mudgal, and O. Schofield, "Slocum gliders – a component of operational oceanography," in *Proc. Int. Symp. Unmanned Untethered Submersible Technology*, Durham, NH, 2005.