

# Trajectory Learning for Human-robot Scientific Data Collection

Geoffrey A. Hollinger and Gaurav S. Sukhatme

**Abstract**—We propose an integrated learning and planning framework that leverages knowledge from a human user along with prior information about the environment to generate trajectories for scientific data collection. The proposed framework combines principles from probabilistic planning with uncertainty modeling through nonparametric Bayesian methods to refine trajectories for execution by autonomous vehicles. The resulting techniques allow for trajectories specified by a user to be modified for reduced risk of collision and increased reliability. We test our approach in the underwater ocean monitoring domain, and we show that the proposed framework reduces the risk of collision with ship traffic by as much as 51% for an autonomous underwater vehicle operating in ocean currents. This work provides insight into the tools necessary for combining human-robot interaction with autonomous navigation.

## I. INTRODUCTION

We envision the future of scientific data collection as a collaborative endeavor between human scientists and autonomous robotic systems. High-impact examples include autonomous underwater vehicles assisting oceanographers to track biological phenomenon [1], aerial vehicles providing imagery of changing ecosystems [2], and ground vehicles monitoring volcanic activity [3]. In these example domains, a key challenge lies in combining the expert knowledge of the scientist with the optimization capabilities of the autonomous system. The scientist brings specialized knowledge and experience to the table, while the autonomous system is capable of processing and evaluating large quantities of data. Leveraging these complementary strengths requires the development of collaborative systems capable of guiding scientific data collection.

In this paper, we propose a learning and planning framework that integrates input from humans with prior data about the environment to reduce the risk of collision for an autonomous vehicle. The input from the human is used as an initial candidate trajectory of waypoints, which is then optimized based on knowledge of disturbances in the environment to avoid areas that have high risk of collision with land and passing ship traffic. The resulting framework combines probabilistic planning with nonparametric learning to improve the safety and reliability of operation. We test

G. Hollinger is with the School of Mechanical, Industrial & Manufacturing Engineering, College of Engineering, Oregon State University, Corvallis, OR 97331 USA, [geoff.hollinger@oregonstate.edu](mailto:geoff.hollinger@oregonstate.edu).

G. Sukhatme is with the Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, [gaurav@usc.edu](mailto: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 CNS-0831728, NSF CCR-0120778, NSF CNS-1035866, and NSF IIS-1317815.

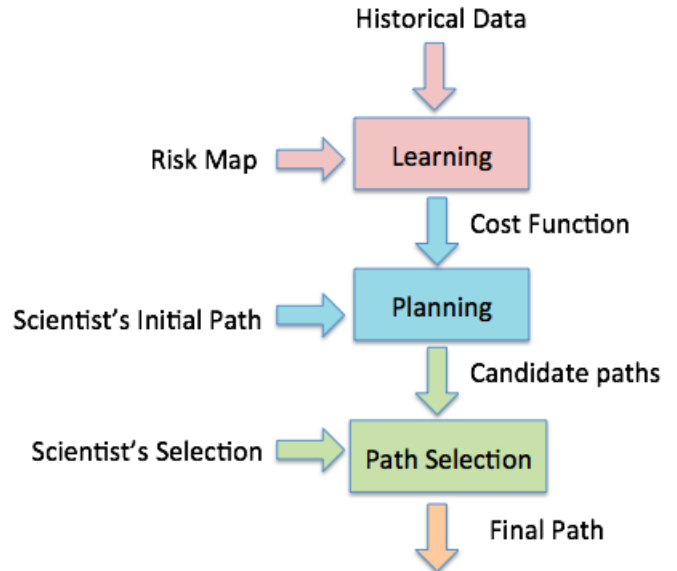


Fig. 1: Proposed framework for human-robot collaboration to generate safe and informative paths for autonomous vehicles to gather scientific data.

our approach in the domain of autonomous underwater ocean monitoring using historical data of ocean currents and shipping activity. This domain is particularly well suited for the testing of human-robot collaboration due to the limited communication available underwater and the necessary supervised autonomy capabilities.

The main novelty of this paper is a unified planning and learning framework that generates reliable and safe trajectories for autonomous vehicles based on human input. To our knowledge, this is the first work to combine Bayesian learning, probabilistic path planning, and user-generated trajectories in a unified approach. The remainder of this paper is organized as follows. We first discuss related work in motion planning, learning, and human-robot interaction (Section II). We then describe the proposed human-robot planning architecture (Section III) and trajectory refinement algorithms (Section IV). A number of data-driven simulations are then presented to show the benefit of the proposed approach (Section V). Finally, we draw conclusions and discuss avenues for future work (Section VI).

## II. RELATED WORK

This paper draws on a large body of prior literature in robotic motion planning, learning, and human-robot interaction. We will now discuss related work in these three

subareas and highlight the need for a unified architecture.

Motion and path planning are fundamental problems in robotics and have been studied extensively in the past two decades [4], [5]. Increasing the degrees of freedom of the robot or the dimensionality of the environment typically causes an exponential increase in the computation required to solve the planning problem optimally. Thus, motion and path planning problems are generally computationally hard (NP-hard or PSPACE-hard) [6]. Modern planning methods have focused on the generation of approximate plans with limited computation (e.g., RRT\* algorithms [7]). Our work extends these ideas to domains where human-robot collaboration is beneficial for the generation of high-quality plans.

A key component of our work is the generation of confidence measures for the reliability of the robot’s trajectory. Prior work has examined the development of such measures for various environmental processes using straightforward statistical machine learning tools [8] as well as more sophisticated Bayesian models [9]. We propose utilizing Gaussian Process (GP) regression [10] augmented with an alternative measure of uncertainty based on the interpolation variance [11] to provide confidence measures on the prediction of uncertainty in scientific data collection domains. Such approaches have been used successfully to improve the accuracy of such tasks as underwater navigation [12], [13] and aerial vehicle surveillance [14]. However, these prior approaches have not integrated input from humans into the learning and planning frameworks.

Our work builds on prior work in imitation learning that builds up cost maps from human operator example paths [15], [16]. These prior works utilize maximum margin planning techniques to learn cost maps from user input. Prior work has focused on surveillance problems, where the goal is to generate a safe path through a hazardous environment. Our proposed framework expands on these ideas by integrating the human scientist into the planning loop.

The problems discussed here are closely related to the problem of adaptive sampling, which has been studied extensively in the robotics and machine learning communities [17], [18]. In adaptive sampling, the goal is to choose observation locations that minimize prediction uncertainty or maximizing some measure of information gain. Unlike typical adaptive sampling applications, our work also considers probabilistic risk maps and reliability of operation. We also integrate the preferences of the scientist at a fundamental level to develop a unified planning architecture for human-robot collaboration.

There is a recent thrust in environmental monitoring towards the development of Decision Support Systems that allow the human operator to seamlessly track the progress of autonomous vehicles and to issue commands on the fly [19], [20], [21]. Such systems are capable of monitoring the progress of autonomous vehicles operating in the ocean and in other unstructured environments by providing data to scientists in real time. Our work complements these systems by generating suggestions for alternative paths to the humans in addition to useful passive data.

### III. HUMAN-ROBOT ARCHITECTURE

The workflow of the proposed architecture is as follows: (1) a human operator specifies a series of waypoints for a vehicle to gather scientific data, (2) the waypoints are refined by the system to suggest alternative trajectories that have lower risk of collision, and (3) the human operator chooses the desired path. Figure 1 gives a visualization of the proposed architecture.

The problems that we will examine consist of the following components: a trajectory of waypoints provided by the scientist that indirectly specifies the quality of information gathered, a “risk” map that provides the expected safety of operating in given area, and a model of the environment that determines how reliably the autonomous vehicle can move between points in the environment. Figure 2 gives an example of the necessary maps in an oceanographic monitoring domain where an autonomous underwater vehicle is monitoring a number of ecological hotspots (e.g., harmful algal blooms [22]). The waypoints specified by the scientist provide areas of high algal bloom density, the risk map provides the probability of colliding with a ship or running into land, and the reliability of operation is determined by the ocean currents. We note that all of these maps are uncertain in the sense that the quantities are not known exactly at the time of planning. For reliable operation, it is necessary to predict the values of the information, risk, and reliability maps, as well as the uncertainty in those values.

Several additional examples show the broader applicability of the delineation between information quality, risk, and reliability maps. For instance, in the aerial surveillance domain, the information quality represents the usefulness of the viewing an area, the risk is the probability of collision (or detection in a security application), and the reliability is derived from the wind and vehicle dynamics. Similar examples in ground vehicle reconnaissance are also apparent [15].

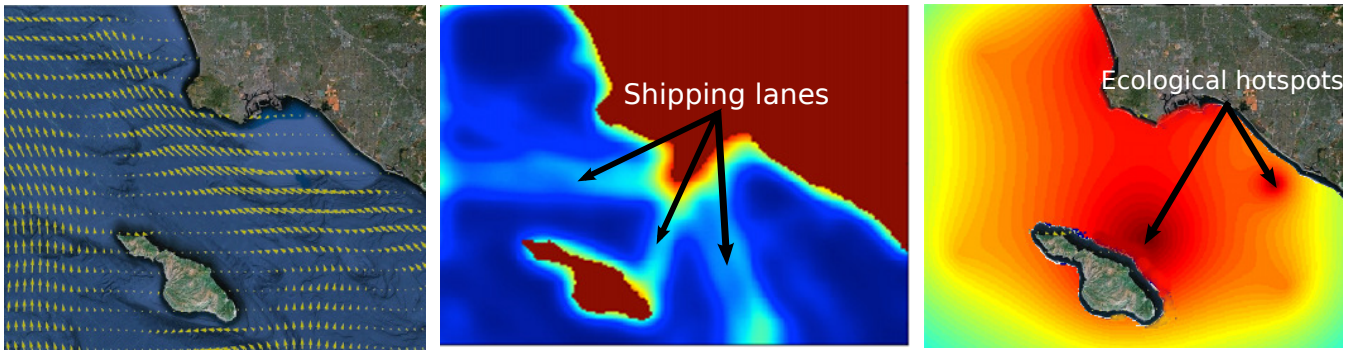
### IV. TRAJECTORY REFINEMENT ALGORITHMS

We first discuss our proposed approach for incorporating input from a human scientist into an optimization framework. We will build on ideas originally presented in prior work to learn cost maps that guide autonomous ground reconnaissance vehicles [15]. Unlike prior work, our methods will incorporate a measure of risk into the planning and learning framework. We will also incorporate measures of uncertainty into our predictions.

The formal problem is to plan a path  $\xi$  that is a solution to the following optimization problem:

$$\xi^* = \operatorname{argmin}_{\xi \in \Psi} D(\xi, \xi_0) + \alpha R(\xi), \quad (1)$$

where  $D(\xi, \xi_0)$  is the deviation from the scientist’s initial trajectory of waypoints  $\xi_0$ ,  $R(\xi)$  is an expected risk of executing  $\xi$ ,  $\Psi$  the space of all possible paths, and  $\alpha$  is a weighting parameter. We assume that we are given an example trajectory of waypoints  $\xi_0$  from the human operator and that an explicit reward function is not provided as part of the trajectory. After approximate solutions to the above



(a) Strength of ocean currents provides reliability of operation (b) Collision probability with shipping lanes and land determines risk of operation (c) Ecological hotspots are specified with waypoints provided by the scientist

Fig. 2: Three maps that must be combined to perform efficient information gathering in an ocean monitoring scenario where an autonomous underwater vehicle is tracking a harmful algal bloom. The ocean currents affect the planned path of the vehicle, the risk map determines the safety of operation, and the pre-specified waypoints given by the scientist provide the usefulness of the information gathered. Our proposed system integrates these maps to improve scientific data collection.

optimization problem are found for various values of  $\alpha$ , the resulting trajectories are suggested to the operator for final selection. The operator also has the option to adjust the values of  $\alpha$  to make the trajectories deviate more or less from the initial plan.

There are several properties of the above problem that make it difficult to solve optimally. If the risk function is non-convex, optimizing it will typically be NP-hard for any rich space of paths [5]. In addition to the inherent complexity of the path planning problem, the functions  $D$  and  $R$  may be computationally intensive to compute. Furthermore, the deviation and risk functions may not be known exactly in advance (e.g., risk is only known with some certainty), and it may become necessary to estimate their expected values based on a distribution of possibilities. Similarly, for a given path, it may not be certain that the vehicle can execute the path exactly, which adds an additional level of uncertainty to the optimization.

Successfully solving these challenges and optimizing the vehicle's paths requires the development of both uncertainty modeling and planning solutions. We will now describe how the proposed architecture addresses each of these subproblems.

1) *Modeling Uncertainty*: A key component of our proposed work is to provide a principled estimate of uncertainty for predictions of the vehicle's actions. These uncertainty estimates will be incorporated through probabilistic planning to provide the final suggested paths for the vehicle. We propose using non-parametric Bayesian Regression in the form of Gaussian Processes (GPs) to provide measures of uncertainty [10]. We will now discuss background in GPs and also show how we can use similar ideas to develop novel representations of uncertainty. This formulation follows closely with our prior work in uncertainty modeling for ocean currents [12].

The disturbances at a given latitude  $lat$ , longitude  $lon$ , and time  $t$  can be written as a tuple  $\mathbf{c}(lat, lon, t) = (u, v)$ , where  $u$  and  $v$  are scalar values representing components of

the disturbance vector along the cardinal axes. At a given time  $T$ , we assume we have access to some historical data for times  $t = \{T-1, T-2, \dots\}$  of the disturbances. Given this data, we want to provide predictions for the points of time in the future as well as confidence bounds for these predictions.

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  $\varepsilon$  is Gaussian noise. Since the standard GP models a one-dimensional value  $z_i$ , we can model the full 2D or 3D space using separate GPs or as a coupled process (e.g., using the techniques in [23]).

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 a component of the disturbance vector at that point and time. 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 kernel function  $k(\mathbf{x}_i, \mathbf{x}_j)$  that relates the points in  $\mathbf{X}$  to each other. As in our prior work, we utilize a space/time squared exponential kernel to model correlations between the data [12]. 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  $\sigma_n^2 \mathbf{I}$  yields  $\mathbf{K}_z = \mathbf{K} + \sigma_n^2 \mathbf{I}$ . The following equation predicts the mean function value (e.g., a disturbance value along the predicted trajectory)  $\mu(\mathbf{x}_*)$  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.

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.

While the GP variance provides some useful insight into the uncertainty in predictions, it has been shown in prior work that it fails to correlate with the error in complex disturbances (e.g., ocean currents). To provide a more informed uncertainty measure, we utilize a method based on the interpolation variance [14], [24]. Once a GP has been learned, the interpolation variance can be estimated as:

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

This measure of variance provides a richer representation that accounts for both data sparsity and data variability and has been shown to provide improved prediction for the trajectories of autonomous underwater vehicles [12].

2) *Probabilistic Planning*: The learned uncertainty predictions described above can be incorporated into probabilistic path planners to refine trajectories provided by humans. We propose utilizing Monte Carlo Sampling methods to estimate the risk of a trajectory given an estimate of the disturbances and the uncertainty in those estimates. The planner assumes that the stochasticity in the predictions uses the spatio-temporal variance estimates from the Gaussian Process (either the GP variance or the interpolation variance). These variances are used to generate a distribution of surfacing locations from a set of prior trials (or simulations). This distribution of surfacing locations is used to generate a transition function  $T(s'|s, a)$ , which describes the probability of moving to state  $s'$  given we choose to take action  $a$  from state  $s$ .

The proposed planner uses the transition function described above to evaluate a number of candidate plans. The costs of the plans are given by a weighting of the risk obtained from the risk map and the deviation from the operator's initial trajectory of waypoints  $\xi_0$  as described in Equation 1. The proposed planner sequentially examines each waypoint provided by the operator and checks all possible alternative waypoints. The cost values  $C$  are calculated for each possible action  $a$  that could be taken from initial waypoint  $s$  using the following rule:

$$C(s, a) \leftarrow \sum_{s'} T(s'|s, a) (\Delta D(s', \xi_0) + \alpha \Delta R(s')), \quad (5)$$

where  $\Delta D(s', \xi_0)$  is the deviation caused by adding state  $s'$  to the trajectory and  $\Delta R(s')$  is the risk caused by adding state  $s'$  to the trajectory. In the domains of interest, the actions in the above equation represent target waypoints (note that the actual waypoint reached will be different from the target waypoint due to the modeled disturbances). We discretize the possible target waypoints in the environment and then select the action with the lowest expected cost value. This process is then repeated for the next waypoint until the entire trajectory has been modified.

Given the appropriate uncertainty measures and the planning methods described above, we now have a framework to modify waypoints provided by the user. We would expect the transition models and risk maps to provide improvements

in the reliability and safety of the resulting plan. The data-driven simulations in the following section will confirm this trend.

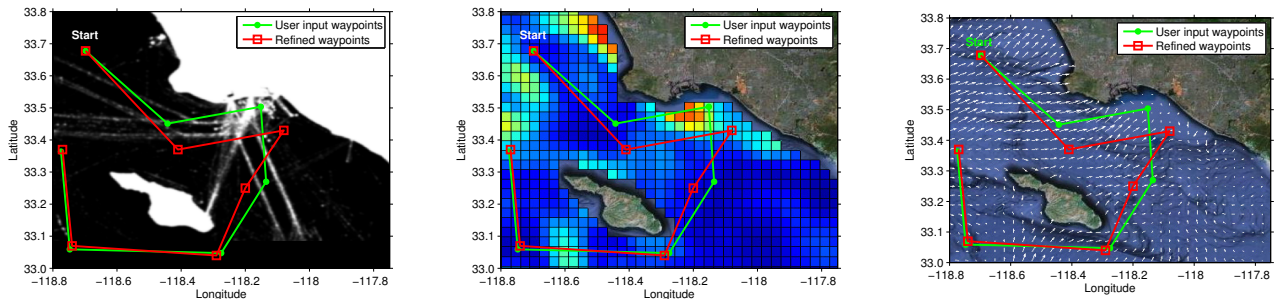
## V. DATA-DRIVEN SIMULATIONS

We now present a validation of the proposed framework in the underwater monitoring domain in the Southern California Bight region where an autonomous underwater vehicle (AUV) is monitoring an oceanographic phenomenon with the help of a scientist. The simulations model a Slocum Glider [25], which is a buoyancy-controlled AUV that moves at a speed of approximately 0.3 m/s. The scientist provides the glider with a series of waypoints, and the vehicle dives between the waypoints while using dead reckoning to determine when to surface. Due to its slow speed, the glider is highly susceptible to ocean currents, and it is often difficult to predict where exactly the glider will surface relative to the specified waypoint. The glider is in danger of running aground if it comes too close to land, and in addition, if the glider surfaces within a shipping lane, it becomes susceptible to collision with passing boat traffic.

The goal of these simulations is to determine the extent to which we can improve the safety of operation using the proposed learning and planning methods. In this domain, safety is measured by the probability of the underwater vehicle successfully completing its mission without coming too close to land or encountering a passing ship. The scientist provides the initial series of waypoints, and the proposed framework is then used to modify these waypoints to decrease the probability of collision while also minimizing the necessary deviation from the initial plan.

The simulations were performed on a single desktop PC with a 3.2 GHz Intel i7 processor and 9 GB of RAM. The simulations incorporate data from ocean currents provided by the JPL Regional Ocean Modeling Systems (ROMs) [26]. The JPL ROMs system provides estimates of the ocean currents but not uncertainty in those estimates. The uncertainty in the ocean currents was determined using the interpolation variance as discussed in Section IV. The uncertainty learning portion of the proposed method took approximately 5 minutes to complete, and the planning portion completed in less than a second using a  $40 \times 40$  discretized grid of possible waypoint locations. We note that the learning portion only needs to be run once per day, and many trajectories can then be refined using those uncertainty estimates.

Risk maps were generated for the simulation using historical Automatic Identification Systems (AIS) shipping data. AIS is a tracking system that mandates a large number of vessels in the United States (and other countries) to broadcast their locations information via VHF transceivers (see [25] for more details). We used historical AIS data collected over a period of five months (between January and May, 2010) in the region 33.25 degrees N to 34.13 degrees N and 117.7 degrees W and 118.8 degrees W. Using this data, we calculated an aggregate risk value at all possible discretized



(a) Initial and refined waypoints overlaid on risk map. Lighter areas have higher risk of collision with land or passing ships. (b) Initial and refined waypoints overlaid on ocean current uncertainty map. Redder areas denote higher areas of normalized uncertainty. (c) Initial and final trajectory overlaid on ocean current predictions. Vectors denote direction and magnitude of ocean currents.

Fig. 3: Example of improving the initial trajectory of waypoints provided by a scientist using the proposed learning and planning framework for data from April 24, 2013. The refined path avoids the riskier (lighter) areas and also remains in areas where the uncertainty of the ocean currents is low. The resulting path is safer and more reliable without deviating much from the initial waypoints given by the scientist.

waypoints in the region, which correlates with the chance of hitting a passing ship.

Figure 3 shows an example of how an initial trajectory was modified to reduce risk in the ocean monitoring domain, and Figure 4 gives a visualization of how changing the weighting value ( $\alpha$ ) affects the resulting trajectories. In this example, the human operator chooses to move the AUV into a risky harbor region to gather data. The algorithm then modifies this trajectory to preserve the intent of gathering data in the harbor while also avoiding the high-risk shipping lanes and the most dangerous areas of the harbor. We also see that the estimates of the ocean currents are more certain in the areas traversed by the modified trajectories. Effectively combining these two types of information would require an expert operator capable of processing data in real-time. With the proposed system, the operator can choose the desired trajectory without any underlying knowledge of the risk and then have the system refine it for increased safety and reliability.

Finally, we provide quantitative evaluations of the deviation and risk tradeoff between a high ( $\alpha = 1000$ ) and low ( $\alpha = 100$ ) value of the weighting parameter. In Figure 5, we see that the proposed methods provides trajectories that range from closely tracking the initial trajectory with high risk to loosely following the initial trajectory with lower risk. In some cases we are able to achieve up to a 51% reduction in risk while deviating from the path by less than 0.8 km.

## VI. CONCLUSIONS AND FUTURE WORK

The results in this paper have shown that it is possible to combine waypoints provided by a human operator with historical data to improve the operation of autonomous vehicles in scientific monitoring scenarios. We have proposed Bayesian learning techniques that allow for uncertainty in predictions to be incorporated into the final trajectory, and we have integrated these uncertainty estimates into a probabilistic planning framework. The resulting framework allows for reduced risk of collision for an autonomous

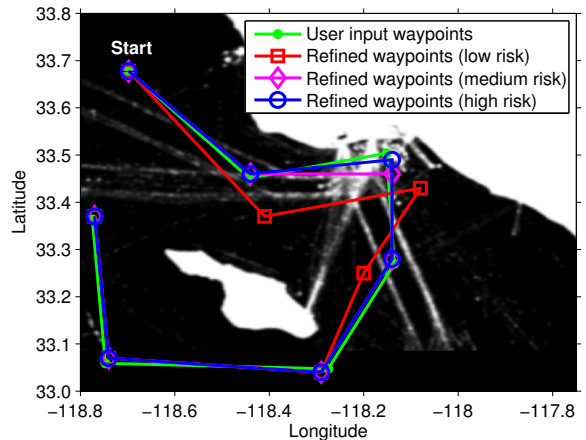


Fig. 4: Initial trajectory and set of suggested trajectories using different weightings between deviation and risk. Higher weighting of risk leads to safer paths at the cost of deviating from the initially specified trajectory.

glider performing an ocean monitoring task with input from a human operator. By integrating feedback from the user into an algorithmic planning framework, we have effectively improved the safety and reliability of autonomous vehicle operation.

This paper opens up a number of avenues for future work at the intersection of human-robot interaction and autonomous path planning. One promising avenue is the development of lifelong learning approaches that allow trajectories specified by the human operator to be stored to improve trajectory generation for future plans. In addition, it may be beneficial to change the ordering of the waypoints to improve the safety of the trajectory. The incorporation of re-ordering into our framework is fairly straightforward; however, it would require additional metrics to determine the deviation penalty from the scientist's original path. Ulti-

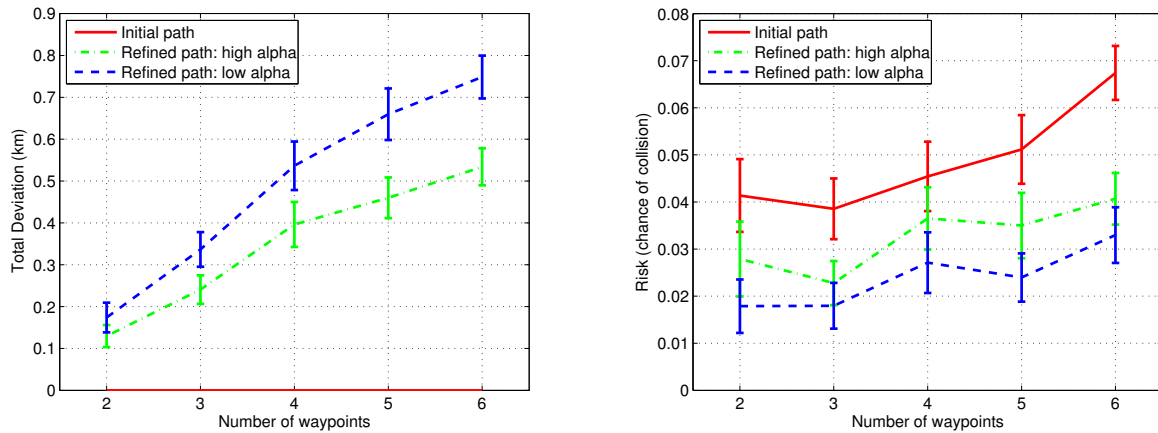


Fig. 5: Deviation and risk for varying weighting parameters for trials using data from April 24, 2013. The proposed method allows the scientist to trade off between the initially selected waypoints and safer paths that deviate from them. Each data point is averaged over 20 user-input trajectories, and error bars are one SEM.

mately, we believe that techniques like the one proposed here will improve the efficiency of scientific data collection and allow us to gather data about phenomena that were previously outside the reach of scientific investigation.

#### REFERENCES

- [1] R. N. Smith, J. Das, H. K. Heidarsson, A. A. Pereira, F. Arrichiello, I. Cetinic, L. Darjany, M.-E. Garneau, M.-D. Howard, C. Oberg, M. Ragan, E. Seubert, E. C. Smith, B. Stauffer, A. Schnetzer, G. Toro-Farmer, D. A. Caron, B. H. Jones, and G. S. Sukhatme, "USC CINAPS builds bridges: Observing and monitoring the Southern California Bight," *IEEE Robotics and Automation Mag.*, vol. 17, no. 1, pp. 20–30, Mar 2010.
- [2] M. Bryson, A. Reid, F. Ramos, and S. Sukkarieh, "Airborne vision-based mapping and classification of large farmland environments," *J. Field Robotics*, vol. 27, no. 5, pp. 632–655, 2010.
- [3] S. L. Nooner and W. W. Chadwick Jr., "Volcanic inflation measured in the caldera of axial seamount: Implications for magma supply and future eruptions," *Geochemistry, Geophysics, Geosystems*, vol. 10, no. 2, 2009.
- [4] J. C. Latombe, *Robot Motion Planning*. Boston, MA: Kluwer Academic Publishers, 1991.
- [5] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006, also available at <http://planning.cs.uiuc.edu/>.
- [6] J. H. Reif, "Complexity of the mover's problem and generalizations," in *Proc. IEEE Symp. Foundations of Computer Science*, 1979, pp. 421–427.
- [7] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *Int. J. Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [8] 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," *J. Geophysical Research*, vol. 90, no. C12, pp. 8995–9005, 1985.
- [9] P. F. Lermusiaux, "Uncertainty estimation and prediction for interdisciplinary ocean dynamics," *J. Computational Physics*, vol. 217, pp. 176–199, 2006.
- [10] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [11] J. Yamamoto, "An alternative measure of the reliability of ordinary kriging estimates," *Mathematical Geology*, vol. 32, no. 4, pp. 489–509, 2000.
- [12] G. Hollinger, A. Pereira, and G. S. Sukhatme, "Learning uncertainty models for reliable operation of autonomous underwater vehicles," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2013, pp. 5573–5579.
- [13] G. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *Int. J. Robotics Research*, vol. 32, no. 1, pp. 3–18, 2013.
- [14] 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*, 2012.
- [15] N. Ratliff, D. Silver, and J. A. Bagnell, "Learning to search: Functional gradient techniques for imitation learning," *Autonomous Robots*, vol. 27, no. 1, pp. 25–53, July 2009.
- [16] D. Silver, J. A. Bagnell, and A. Stentz, "Learning from demonstration for autonomous navigation in complex unstructured terrain," *Int. J. Robotics Research*, vol. 29, no. 1, pp. 1565–1592, 2010.
- [17] K. H. Low, J. M. Dolan, and P. Khosla, "Information-theoretic approach to efficient adaptive path planning for mobile robotic environmental sensing," in *Proc. Int. Conf. Automated Planning and Scheduling*, 2009, pp. 233–240.
- [18] D. R. Thompson, S. Chien, Y. Chao, P. P. Li, B. Cahill, J. Levin, O. Schofield, A. Balasuriya, S. Petillo, M. Arrott, and M. Meisinger, "Spatiotemporal path planning in strong, dynamic, uncertain currents," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2010.
- [19] J. Das, T. Maughan, M. McCann, M. Godin, T. O'Reilly, M. Messie, F. Bahr, K. Gomes, F. Py, J. Bellingham, G. Sukhatme, and K. Rajan, "Towards mixed-initiative, multi-robot field experiments: Design, deployment, and lessons learned," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2011.
- [20] P. Li, Y. Chao, Q. Vu, Z. Li, J. Farrara, H. Zhang, and X. Wang, "OurOcean - an integrated solution to ocean monitoring and forecasting," in *Proc. Oceans MTS/IEEE Conf.*, Boston, MA, 2006.
- [21] J. Sattar and G. Dudek, "Towards quantitative modeling of task confirmations in human-robot dialog," in *Proc. IEEE Int. Conf. Robotics and Automation*, May 2011, pp. 1957–1963.
- [22] J. Das, K. Rajan, S. Frolov, J. Ryan, F. Py, D. A. Caron, and G. S. Sukhatme, "Towards marine bloom trajectory prediction for auv mission planning," in *Proc. IEEE Int. Conf. Robotics and Automation*, 2010.
- [23] M. A. Álvarez and N. D. Lawrence, "Computationally efficient convolved multiple output gaussian processes," *J. Machine Learning Research*, vol. 12, pp. 1425–1466, 2011.
- [24] 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.
- [25] A. Pereira, J. Binney, G. Hollinger, and G. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *J. Field Robotics*, vol. 30, no. 5, pp. 741–762, Sept. 2013.
- [26] 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.