

Active and Adaptive Dive Planning for Dense Bathymetric Mapping

Geoffrey A. Hollinger, Urbashi Mitra, and Gaurav S. Sukhatme

Abstract We examine the problem of planning dives for an Autonomous Underwater Vehicle (AUV) to generate a dense bathymetric map using sidescan sonar. The three key challenges in this scenario are (1) proper modeling of the local uncertainty of the 3D reconstruction, (2) efficient dive planning to reduce this uncertainty, and (3) determination of when to re-plan adaptively based on new information. To address these challenges, we propose using non-parametric Bayesian regression to model the expected accuracy of the map, which provides principled cost functions for planning subsequent dives. In addition, we propose an efficient greedy method to reduce this uncertainty, and we show that it achieves theoretically bounded performance given assumptions on the sensor model and the form of the uncertainty function. We present experiments on the propeller-driven YSI EcoMapper AUV equipped with a sidescan sonar in an inland lake. The experiments demonstrate the benefit of efficient dive planning, with our results providing performance gains of up to 83% versus standard lawnmower patterns.

1 Introduction

The increasing capabilities of Autonomous Underwater Vehicles (AUVs) have enabled their deployment in oceans, coastal environments, and inland waterways. There are a number of potential applications for such vehicles, including ecological

Geoffrey A. Hollinger and Gaurav S. Sukhatme
Computer Science Department, Viterbi School of Engineering, University of Southern California,
Los Angeles, CA 90089, e-mail: {gahollin, gaurav}@usc.edu

Urbashi Mitra
Ming Hsieh Department of Electrical Engineering, Viterbi School of Engineering, University of
Southern California, Los Angeles, CA 90089, e-mail: ubli@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 CCR-0120778, and NSF CNS-1035866.



Fig. 1 YSI EcoMapper propeller-driven AUV used to perform bathymetric mapping experiments in Puddingstone Lake. The vehicle is equipped with an Imaginex Sportscan sidescan sonar for underwater imaging as well as a Doppler Velocity Log (DVL) for navigation.

monitoring [26], mine sweeping [20, 34], inspection of submerged structures [15], and underwater habitat mapping [24, 27]. In such scenarios, deployment time for the AUV represents a considerable resource cost, and there is significant motivation to improve the efficiency of the deployment relative to the desired mission goal. The problem of planning the AUV’s mission to improve the quality of the sensor data fits into the broad framework of *active perception* that dates back to early seminal research in robotics and computer vision [3]. The key idea behind active perception is that we can plan the path and sensor views of robotic vehicles to maximize information gained while minimizing time and/or energy consumption.

In addition to actively planning the vehicle’s dives, we also explore the benefit of *adaptive* dive planning for AUVs. When calculating efficient dive patterns, it may be beneficial to adaptively re-plan as new information is acquired. In our recent work, we have begun to quantify the potential benefit of adaptivity in underwater inspection scenarios [15, 16], which moves towards planning algorithms that selectively adapt to new information when it is most beneficial. In the current paper, we explore these ideas experimentally on an AUV performing underwater mapping tasks.

Our experiments validate the principles of active and adaptive perception in the domain of underwater bathymetric mapping using an AUV. The two hypotheses to be experimentally tested are (1) active dive planning improves the efficiency of bathymetric mapping and (2) selective use of adaptivity further improves this efficiency. These hypotheses are suggested by our earlier theoretical analysis [16] and by experiments in related underwater inspection domains [15]. The experiments presented here serve to advance our understanding of the principles behind these related active sensing problems.

The remainder of this paper is organized as follows. We first discuss related work in active perception and underwater mapping (Section 2). We then describe the active dive planning problem in more detail and propose solutions for measuring the uncertainty of the reconstruction and planning informative dives (Section 3). Next, we discuss experimental results from the AUV performing active and adap-

tive strategies for dense bathymetric mapping (Section 4). Finally, we conclude and discuss avenues for future work (Section 5).

2 Related Work

The study of general active perception problems dates back to early work in active vision [1, 3] and next-best view planning [9]. Early techniques were often concerned with the “next-best” action to maximize information with relatively less focus on long-term planning methods. While this early work made limited use of information theoretic methods, later work in this area has increasingly focused on incorporating information theory and probabilistic techniques into active perception systems [25]. Two recent surveys describe the breadth of work in active perception and its development over the past two decades [7, 22]. Our experiments are complementary to this prior work, in that we provide experimental validation of the benefit of active dive planning in the underwater mapping domain. In addition, our work examines the use of long-term planning and the benefit of adaptivity, which are often overlooked in prior work.

Our formulation connects the active perception problem with sequential hypothesis testing, another classical problem where an observer must select a sequence of noisy experiments to determine the nature of an unknown [30]. Hypothesis testing has applications in information theory, machine learning, and wireless communications. Of particular interest is the connection between sequential hypothesis testing and feature selection in Bayesian learning, which has been applied to sensor placement problems in the context of identifying objects in 2D images [6]. This early work in robotic perception provides formal justification for the formulation of active perception as a Bayesian learning problem.

Many active perception problems can also be seen as instances of submodular orienteering and informative path planning [23], where the goal is to optimize the path of a robot to gain the maximal amount of information. Many of these problems allow for efficient algorithms with performance guarantees due to the property of submodularity, a formal characterization of the property of diminishing returns. Recent advances in submodular optimization have extended informative sensor placement problems to domains that require adaptive re-planning [14]. In complementary work, the stochastic optimization community has begun examining the potential benefit of adaptive re-planning for such problems as stochastic covering and stochastic knapsack [11]. To our knowledge, these ideas have not previously been applied to underwater robotics applications.

Acoustic range sensing, essential in the inspection of turbid-water environments, has been used to produce 3D point clouds throughout various underwater domains. Early work in 3D underwater mapping dates back more than a decade to the use of diver-held sonar devices that utilized an inertial measurement unit to help align subsequent images [2]. Such techniques inevitably suffer from drift over time, which can be mitigated through Simultaneous Localization and Mapping (SLAM) tech-

niques [12]. It is also possible to generate underwater maps without motion estimation through sophisticated techniques for accurate registration [5]. In the large body of prior work in underwater mapping, the vehicle is remotely controlled by a human operator, and the potential for autonomous operation is not explored.

The underwater robotics community has recently begun examining problems of AUV dive planning for such applications as mine sweeping [33] and seabed classification [32]. These works focus on the dive planning problem, but they do not integrate methods for measuring 3D reconstruction accuracy. One early approach for measuring uncertainty in 3D reconstructions uses a linear approximation of the surface, where the resulting covariance provides a measure of accuracy [31]. Our approach using Gaussian Processes provides a more general uncertainty measure that utilizes both data sparsity and surface complexity when estimating uncertainty. Similar non-parametric regression approaches have been used to define reconstruction accuracy metrics for 2.5D surface estimation [29]; though these techniques do not actively plan to reduce uncertainty. In addition, our approach utilizes the augmented input vector method [28] to provide non-stationary kernels that improve the richness of the uncertainty representation.

3 Algorithm Design

The goal is to plan the dives of the autonomous underwater vehicle to provide an accurate reconstruction of the bottom of a body of water. The first step in solving this problem is to define a reconstruction accuracy metric that is computable across the area of interest. To do so, we propose using non-parametric Bayesian Regression in the form of Gaussian Processes (GPs) [21].¹ If we view the 3D reconstruction as a function approximation of surface height over a 2D bottom plane, the GP representation has the advantage of providing a principled measure of variance, which we use as a measure of uncertainty, as well as a mean value at each point. The resulting uncertainty measure can then be used in a greedy planning framework to plan subsequent dives that we expect to be highly informative for generating a dense and accurate mapping.

3.1 Gaussian Process Modeling

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 $D = [(\mathbf{x}_1, z_1), (\mathbf{x}_2, z_2), \dots, (\mathbf{x}_n, z_n)]$. For 2.5D surfaces that do not loop on themselves, as is typically the case for bathymetric

¹ We note that Gaussian Process regression is closely related to Kriging interpolation as described in the geostatistics literature [4].

mapping, \mathbf{x}_i is a point in the 2D plane ($d = 2$), and z_i represents the height at that point.²

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 define a GP, it is necessary to choose a covariance function that relates points in \mathbf{X} . We employ the commonly used squared exponential, which produces a smooth kernel that drops off with distance:

$$k(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp \left(- \sum_{k=1}^d w_k (\mathbf{x}_{ik} - \mathbf{x}_{jk})^2 \right). \quad (1)$$

The hyperparameter σ_f represents the process noise, and each hyperparameter w_k represents a weighting for the dimension k . Once the kernel has been defined, combining the covariance values for all points into an $n \times n$ matrix \mathbf{K} and adding a Gaussian observation noise hyperparameter σ_n yields $\mathbf{K}_z = \mathbf{K} + \sigma_n^2 \mathbf{I}$. We can now estimate the kernel hyperparameters $\theta = (\sigma_f, \sigma_n, w_{1:d})$ using the standard method of maximizing the likelihood of the measurements given the data and the hyperparameters [21]:

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

This likelihood is maximized using conjugate gradient optimization. We note that for larger data sets, downsampling or block learning approaches may be required to make hyperparameter learning tractable.

We now wish to predict the mean function value (height) \bar{f}_* and variance $\mathbb{V}[f_*]$ at a selected point \mathbf{x}_* given the measured data:

$$\bar{f}_* = \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{z}, \quad (3)$$

$$\mathbb{V}[f_*] = k(\mathbf{x}_*, \mathbf{x}_*) - \mathbf{k}_*^T (\mathbf{K} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{k}_*, \quad (4)$$

where \mathbf{k}_* is the covariance vector between the selected point \mathbf{x}_* and the training inputs \mathbf{X} . This model provides a mean height and variance at all points of interest in \mathbb{R}^2 . In this model, the variance gives a measure of uncertainty based on the sparsity of the data and the hyperparameters. It should be noted that this variance is not modeled as a random variable, and it is only dependent on data density. Thus, we can estimate the variance exactly before we have taken measurements assuming that we know exactly which points we will observe. We note that this assumption may not hold, and we provide a more detailed discussion of this issue later.

The formulation above determines correlation between points solely based on their proximity in space. Another important consideration when determining uncertainty is the amount of variability in an area. For instance, an underwater ridge would require a large number of points to reconstruct efficiently. To model uncertainty caused by surface variability, we utilize the augmented input vector approach [28]. The idea is to modify the input vector by adding additional terms that affect the cor-

² For more complex mapping tasks with arbitrary 3D geometry, the parametric Gaussian Process Implicit Surface extension can be used [15].

relations between the data. We propose using an initial estimate of surface height (found by running the GP with a standard kernel or another interpolation method) to modify the input vector to be $\mathbf{x}' = (x, y, \bar{z})$, where \bar{z} is the initial estimate of the surface height at that point. The weighting hyperparameters can be adjusted to modify the effect of spatial correlation versus surface height correlation.

To address the problem of scalability, we utilize a local approximation method using KD-trees [29]. Since correlated points will typically be near each other, we can pre-calculate a KD-tree and then retrieve the k closest points when calculating the estimate and variance at a given point. Since the retrieval from the KD-tree is on average $O(\log n)$, the resulting local approximation of the GP requires $O(k^3 \log n)$ computation for each point, or $O(nk^3 \log n)$ total. Comparing this to $O(n^3)$ computation of a typical GP, we see that if $k \ll n$, this approximation provides a significant reduction in computation. The value of k can be selected based on the necessary computing power, which allows the approximation to improve in accuracy with increasing computation.

3.2 Variance Reduction Algorithm

For a given GP representation, we define a total variance by integrating over the space of interest \mathcal{X} . The goal of the dive planning is to generate a dive pattern that maximally reduces this variance. More formally, we define a policy π that executes a given dive based on a given uncertainty representation. We can now define the following metric for dive planning:

$$J_{var}(\pi) = \int_{\mathcal{X}} \mathbb{V}_0(x) - \mathbb{V}_{\pi}(x) dx, \quad (5)$$

where $\mathbb{V}_0(x)$ is the initial variance at point x , and $\mathbb{V}_{\pi}(x)$ is the variance at point x after executing policy π .

For a policy π , we set the measure of information quality to $J_{var}(\pi)$, and we define $c(\pi)$ as the cost of executing the policy. In the application of interest, this cost will be determined by the number of planned dives and their length. We will choose dives that optimize the following maximization problem:

$$\pi^* = \operatorname{argmax}_{\pi} J_{var}(\pi) \text{ s.t. } c(\pi) \leq B, \quad (6)$$

where B is a budget constraint on time or energy.³ Setting the cost as a budget constraint is natural in underwater applications where the battery life or deployment time are limiting factors. For this paper, we define the budget as a pre-specified number of dives multiplied by a fixed dive length (i.e., we only consider dives of

³ We note that we can alternatively maximize the weighted sum $\alpha J(\pi) - \beta c(\pi)$ with appropriate weighting constants using a Lagrangian relaxation [13].

equal length). The extension of this algorithm to dive patterns with dives of unequal length is a topic for future work.

To provide informative dive planning, we propose greedily selecting dives that maximize variance reduction until the budget is reached and then running a gradient optimization that perturbs each dive and locks the pattern into a local optimum. The gradient optimization is not strictly necessary, though we expect it will provide some limited improvement over the greedy policy. We note that even with the gradient optimization, we do not expect the algorithm to yield the optimal dive pattern in all cases due to local maxima. A summary of the proposed algorithm is given in Algorithm 1.

The summary in Algorithm 1 does not specify how we calculate the expected variance reduction for a given dive d . Since we do not know which observations we will receive before making them, and re-running the GP is computationally expensive, the calculation of this quantity is not trivial. We provide an approximation to the variance reduction by first calculating the sum of variance in the area viewed by the dive. For overlapping dive patterns, we approximate the reduction in variance caused by each subsequent view using an exponential drop off. The resulting variance at a point x is found by $\mathbb{V}_n(x) = \mathbb{V}_0(x) \exp(-n/\alpha)$, where $\mathbb{V}_0(x)$ is the initial variance at the point, n is the number of times the point has been viewed in the dive plan, and α is a length scale parameter. The length scale parameter was set to $\alpha = 1$ based on fitting to test runs where the exact reduction was calculated.

A more computationally intensive way of calculating the variance reduction would be to simulate a dive by generating a point cloud of the expected dive result. The GP could then be re-run with the updated point cloud. This approach would require running the GP once for each candidate dive during each planning iteration. We found that the exponential drop off assumption provided sufficient accuracy while remaining computationally tractable.

3.3 Guarantees and Need for Adaptivity

It has been shown in prior work that greedy placement of static sensors to reduce variance in Gaussian Processes provides a constant-factor performance guarantee relative to optimal of $1 - 1/e \approx 63\%$ [18]. Thus, even in the worst-case, the greedy deployment will still achieve 63% of the variance reduction as the optimal policy and can be expected to perform significantly better in practice [17, 23]. When the dives are of equal length, a pre-planned series of greedily planned dives provides the same bounded approximation of the optimal dive pattern if the following assumptions hold:

1. The variance reduction objective function is monotone and submodular. Monotonicity implies that additional measurements always improve the objective. Submodularity implies that the objective follows the law of diminishing returns (i.e., the more measurements observed, the less incremental benefit of receiving a new

Algorithm 1 Active Dive Planning Algorithm

```

1: Input: Uncertainty map  $\mathbb{V}_0$  and dive length budget  $B$ 
2: Select a set of possible dives  $\mathcal{N}$  of equal length  $L$ 
3: Initialize remaining budget  $B_r \leftarrow B$ , uncertainty map  $\mathbb{V}_r \leftarrow \mathbb{V}_0$ , set of selected dives  $\mathcal{M} \leftarrow \emptyset$ 
4: while  $B_r \geq L$  do
5:   for each dive  $d \in \mathcal{N}$  do
6:     Calculate expected uncertainty reduction  $J(d) \leftarrow \mathbb{V}_r - \mathbb{V}_d$ 
7:   end for
8:   Select dive  $d^* = \operatorname{argmax}_d J(d)$ 
9:   Update selected dives  $\mathcal{M} \leftarrow \mathcal{M} \cup d^*$ , uncertainty map  $\mathbb{V}_r \leftarrow \mathbb{V}_d$ , budget  $B_r \leftarrow B_r - L$ 
10: end while
11: while not converged and planning time remains do
12:   for each dive  $d \in \mathcal{M}$  do
13:     for selected directions  $a$  do
14:       Perturb dive by  $\varepsilon$  in direction  $a$  to get  $d^p$ 
15:       Calculate new uncertainty  $J(d^p) \leftarrow \mathbb{V}_r - \mathbb{V}_{d^p}$ 
16:       if  $J(d^p) > 0$  then
17:         Update dive pattern  $\mathcal{M} \leftarrow d^p \cup \mathcal{M} \setminus d$  and uncertainty  $\mathbb{V}_r \leftarrow \mathbb{V}_{d^p}$ 
18:       end if
19:     end for
20:   end for
21: end while
22: Sort dives in  $\mathcal{M}$  to minimize execution time
23: Execute dive pattern  $\mathcal{M}$ 

```

measurement). This assumption holds in many cases for variance reduction in Gaussian Processes with fixed hyperparameters [10, 18].

2. The 2D locations where measurements will be received are known in advance (i.e., we can predict exactly which points in 2D we will receive data from before executing a dive). We note that for realistic sensor models, this assumption may not be valid.
3. The hyperparameters do not change during the dive planning. For the case of augmented input vectors (see above), this implies that the surface height does not change significantly (i.e., the surface heights act as implicit hyperparameters). This assumption will be violated whenever significant changes occur in the reconstruction.

Since the last two assumptions can be violated in real-world scenarios, there is motivation to re-plan subsequent dives based on new information. However, it is not obvious how much they will be violated, which depends both on the rate of change of the 3D reconstruction and the sensor model of the sidescan sonar. Thus, we examine the benefit of adaptivity in active dive planning through experimental trials.

4 Experiments and Results

We now explore the benefit of active and adaptive planning for 3D reconstruction by mapping a portion of the bottom of Puddingstone Lake in Southern California (Lat. 34.088854°, Lon. -117.810667°). The region of interest is approximately 10 m deep and covers an area of 100 m \times 50 m of lakebed. We utilize a YSI EcoMapper propeller-driven AUV (shown in Figure 1) capable of moving at speeds of 5 knots and diving down to 100 m. A downward-looking SportsScan Imaginex sidescan sonar is mounted on the vehicle. The vehicle is also equipped with a Doppler Velocity Log (DVL) and GPS unit, which provide navigation capabilities. A photograph of the vehicle is shown in Figure 1.

4.1 3D Reconstruction from Sidescan Sonar

The Imaginex sidescan sonar returns 2D intensity images that do not include 3D depth information. Example images are shown in Figure 2. It is important to note that the AUV has a significant “blind spot” directly under it, which necessitates overlap in images to provide a complete reconstruction. In order to generate a 3D reconstruction from a 2D sidescan image, we use standard shape-from-shading techniques that have been successfully applied in prior work [8]. The central idea is to make assumptions on the reflectance properties of the sonar and then use the geometry of the sidescan position to develop a fully 3D reconstruction from a collection of 2D images.

The sensor provides an intensity return from each point viewed on the 2D bottom plane. We use the traditional Lambertian model [19], which relates intensity to the angle of reflection. This model assumes that the scattering is diffuse and that the returned intensity is not dependent on the angle of observation or the frequency of the sonar pulse. For a viewed point p , the intensity can be computed as:

$$I(p) = K\Phi(p)R(p)\cos(\theta(p)), \quad (7)$$

where Φ denotes the intensity of the illuminated sonar pulse, R denotes the reflectivity of the seabed, θ is the angle of incidence, and K is a constant. Values for Φ , R , and K were determined empirically in this work based on accepted values from prior work [8, 19]. Improving the estimates of these values would serve to improve the 3D reconstruction, and our active planning techniques can easily be applied to such improved reconstructions.

From the above equation, we can determine an angle of incidence for each viewed point from known quantities. Given the angles of incidence, we reconstruct a 3D representation of the bottom through numerical integration. More detail on shape-from-shading approaches can be found in prior work [8]. Since the focus of this paper is on validating the benefit of active dive planning, we do not implement




Fig. 2 Example sidescan sonar images providing bathymetry information in Puddingstone Lake. Variations in the lakebed topography are apparent as bright returns in the images. A full 3D reconstruction can be found using shape-from-shading methods that make assumptions on the reflectance properties of the sonar.

post-processing of the reconstruction through constrained optimization, though we expect such techniques would improve the quality of the 3D reconstruction.

4.2 Active Dive Planning

We now test the benefit of active planning for informative dives to improve the 3D reconstruction. Planning and learning was performed on a desktop machine with a 3.2 GHz Intel i7 processor and 9 GB of RAM. A variance map could be calculated in approximately 30 seconds using a discrete grid approximation with 1 m resolution and a KD-tree local approximation of the GP using 100 points. Once a variance map was available, the dive patterns were calculated in less than a second using the exponential drop off approximation (see Section 3). The plan was then transmitted to the vehicle on the surface using a standard wireless network.

The experiment progresses as follows: (1) The AUV executes three evenly spaced dives to generate an initial 3D reconstruction of the area of interest, (2) the initial reconstruction is used to plan a subsequent three dive pattern using Algorithm 1, and (3) the vehicle executes the resulting dive pattern. The final uncertainty after six dives is then compared to an alternative dive pattern that maximizes overlap in the sidescan images without using the uncertainty representation. The difference between the uncertainties represents the benefit of active planning in this domain.

Figure 4 shows a comparison of 3D reconstructions for the two dive patterns described above: one maximizing coverage and one optimizing variance reduction as in Algorithm 1. The reconstruction resulting from the minimum variance dive pattern does not contain the clearly erroneous valleys that occur in the maximum coverage dive pattern. Figure 5 shows a quantitative comparison of the total uncertainty of the two dive patterns, as well as total uncertainty after the initial 3 dives and after an exhaustive 9 dive pattern. The variance dive pattern using 6 dives pro-



Fig. 3 Dive patterns optimized to maximize coverage overlap (left) and uncertainty reduction (right). The region of interest is the $100\text{ m} \times 50\text{ m}$ area (Lat. 34.088854° , Lon. -117.810667°) surrounded by the rectangular sensor footprint and highlighted in blue and red. Points of interest covered once by the sidescan sonar are labeled blue (darker), and areas of interest covered twice are labeled red (lighter).

vides nearly the same uncertainty as the 9 dive pattern. The 9 dive pattern represents a target uncertainty at which sensor noise and errors from the shape from shading make further reduction difficult.

Comparing the reduction in uncertainty of the variance dive pattern and the coverage dive pattern, we see that the variance dive pattern improves the uncertainty reduction by 83%, showing a significant benefit from using active planning. In addition, the 6 dive pattern that reduces variance achieves a smaller percent error for the reconstruction when the percent error is computed relative to the 9 dive pattern. However, we note that the reconstruction from the 9 dive pattern does not provide perfect ground truth.

4.3 Benefit of Adaptivity

The experiment described above utilizes adaptive behavior only when re-planning the last three dives (i.e., it learns the uncertainty model and chooses the resulting dives). Alternatively, the uncertainty model could be re-learned more often during the mission. We next examine the benefit of such re-planning by running additional AUV data collections. Figure 6 shows a long 20 dive run that took approximately two hours to execute. The 20 dives were pre-run, which allowed us to test different selection methods that choose a subset of the dives. We note that the gradient optimization step (see Algorithm 1) was not used in this experiment since there are a limited number of dives to select at pre-set locations.

Figure 6 shows the quantitative results with a different number of re-planning cycles. The first additional re-planning cycle provides approximately an 8% reduction in uncertainty, and the second re-planning cycle does not provide any improvement.




Fig. 4 3D bathymetric reconstructions found using sidescan sonar (top) and accompanying variance maps (bottom). Left: reconstruction resulting from dive patterns that maximize overlap in sensor coverage. Right: reconstruction from dive patterns that minimize the expected variance in a Gaussian Process representation of the reconstruction. The right image does not suffer from erroneous valleys that result from data sparsity, and the variance map matches this qualitative observation.




Fig. 5 Comparison of the resulting bathymetric reconstructions for different dive patterns. The initial dive pattern uses 3 evenly spaced dives. The coverage dive pattern adds an additional 3 dives (total of 6 dives) that maximize coverage overlap. The variance dive pattern instead uses the additional three dives (total of 6 dives) to maximize reduction of uncertainty in the map. The variance dive pattern provides improved uncertainty reduction, nearly that of an exhaustive dive pattern that uses 9 dives. The graph shows both total uncertainty remaining (dark blue) and average reconstruction error computed relative to the 9 dive pattern (light red).




Fig. 6 Crisscross pattern used to test the benefit of adaptive dive planning (left) and the uncertainty remaining in the $40\text{ m} \times 40\text{ m}$ area after 6 dives for different numbers of re-planning cycles (right). The vehicle may choose to execute any 6 of 20 dives (10 North/South and 10 East/West) to maximize the accuracy of the 3D bathymetric reconstruction. The dives are pre-run, which allowed post-processing the data to determine the benefit of different dive selection methods. Re-planning consists of re-calculating the uncertainty representation before determining the next dive. The first additional re-planning cycle provides approximately an 8% reduction in uncertainty versus a pre-planned dive pattern, but the second re-planning cycle does not provide any improvement. Note that there is less remaining uncertainty in these results than those in Figure 5 due to the smaller area being surveyed.

We note that the benefit from the first re-planning cycle is significant; however, the pre-planned path still provides an accurate 3D reconstruction. Thus, adaptivity is beneficial in this scenario, but the essential component of the algorithm is pre-planning to reduce the uncertainty of the reconstruction. Based on these results, it is clear that we can predict the returns that we will receive from the sidescan sonar to some extent, which motivates the use of long-term planning in this scenario.

5 Discussion and Conclusions

Our experiments have examined the benefit of active dive planning in the real-world application domain of 3D bathymetric mapping. We have shown that planning a dive pattern that greedily maximizes variance reduction provides a significant quantitative and qualitative improvement over standard lawnmower approaches. In addition, we have shown that adaptive planning that re-plans provides some additional improvement. We note that the re-planning occurs once the vehicle has surfaced, which allows for the necessary computation to be placed offboard and communicated to the vehicle. In addition, the planned paths can be checked by a human operator before execution to ensure safety. These experimental findings provide a baseline for future research in active mapping and inspection. In addition, our experiments provide insight into the validation of recent theoretical advances in submodular optimization and active sensing.

In underwater applications, perhaps more so than in other robotics domains, there is significant interest in simplifying the behaviors that occur underwater. Complex

behaviors, such as adaptive planning, can potentially lead to safety concerns and increase the necessary computation carried by the AUV. Our experiments have helped determine the level of adaptivity necessary to provide high performance bathymetric mapping as well as the granularity with which this adaptivity must be made. Our results demonstrate that there is limited benefit to adaptive planning, which avoids some of the necessity for onboard computation on the AUV. The ability to provide good 3D reconstructions without in situ adaptation significantly reduces the computational burden on the vehicle and also simplifies the testing requirements. As a result, our algorithm can easily be applied to any propeller-driven AUV capable of executing a series of waypoints.

Future work includes analysis of alternative uncertainty representations that utilize the specifics of the shape-from-shading method to provide better predictions of areas that need further inspection. For instance, alternative kernels in the GP may allow for more accurate uncertainty modeling. Additional experimental validation includes testing of the proposed method on 3D reconstructions generated from other sensors (e.g., monocular vision, stereo, and infrared imaging). The techniques proposed here and the experimental methods employed are applicable throughout underwater inspection domains as well as on ground vehicles and aerial platforms. Ultimately, the experimental validation of active and adaptive planning moves towards more efficient use of robotic perception for autonomous vehicles.

Acknowledgements The authors gratefully acknowledge Jonathan Binney, Jnaneshwar Das, Arvind Pereira, and Hordur Heidarsson at the University of Southern California for their insightful comments and assistance with experiments.

References

1. Y. Aloimonos, I. Weiss, and A. Bandopadhyay. Active vision. *International Journal of Computer Vision*, 1(4):333–356, Jan. 1988.
2. B. A. am Ende. 3D mapping of underwater caves. *IEEE Computer Graphics and Applications*, 21(2):14–20, Mar./Apr. 2001.
3. R. Bajcsy. Active perception. *Proc. IEEE, Special Issue on Computer Vision*, 76(8):966–1005, Aug. 1988.
4. S. Banerjee, B. P. Carlin, and A. E. Gelfand. *Hierarchical Modeling and Analysis for Spatial Data*. Chapman and Hall/CRC Press, 2004.
5. H. Buelow and A. Birk. Spectral registration of noisy sonar data for underwater 3D mapping. *Autonomous Robots*, 30(3):307–331, Apr. 2011.
6. A. Cameron and H. Durrant-Whyte. A Bayesian approach to optimal sensor placement. *Int. J. Robotics Research*, 9(5):70–88, Oct. 1990.
7. S. Chen, Y. Li, and N. M. Kwok. Active vision in robotic systems: A survey of recent developments. *Int. J. Robotics Research*, 30(11):1343–1377, Sept. 2011.
8. E. Coiras, Y. Petillot, and D. M. Lane. Multiresolution 3-D reconstruction from side-scan sonar images. *IEEE Trans. Image Processing*, 16(2):382–390, Feb. 2007.
9. C. Connolly. The determination of next best views. In *Proc. IEEE Conf. Robotics and Automation*, pages 432–435, Mar. 1985.
10. A. Das and D. Kempe. Algorithms for subset selection in linear regression. In *Proc. ACM Symp. Theory of Computing*, pages 45–54, May 2008.

11. B. Dean, M. Goemans, and J. Vondrak. Approximating the stochastic knapsack: the benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, Nov. 2008.
12. N. Fairfield, G. A. Kantor, and D. Wettergreen. Real-time SLAM with octree evidence grids for exploration in underwater tunnels. *J. Field Robotics*, 24(1–2):3–21, Jan./Feb. 2007.
13. M. L. Fisher. The lagrangian relaxation method for solving integer programming problems. *Management Science*, 27(1):1–18, Jan. 1981.
14. D. Golovin and A. Krause. Adaptive submodularity: Theory and applications in active learning and stochastic optimization. *J. Artificial Intelligence Research*, 42:427–486, Nov. 2011.
15. G. Hollinger, B. Englot, F. Hover, U. Mitra, and G. Sukhatme. Uncertainty-driven view planning for underwater inspection. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 4884–4891, May 2012.
16. G. Hollinger, U. Mitra, and G. S. Sukhatme. Active classification: Theory and application to underwater inspection. In *Int. Symp. Robotics Research*, Aug. 2011.
17. G. Hollinger, S. Singh, J. Djughash, and A. Kehagias. Efficient multi-robot search for a moving target. *Int. J. Robotics Research*, 28(2):201–219, Feb. 2009.
18. A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Machine Learning Research*, 9:235–284, Feb. 2008.
19. D. Langer and M. Hebert. Building qualitative elevation maps from side scan sonar data for autonomous underwater navigation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2478–2483, Apr. 1991.
20. V. Myers and D. Williams. A POMDP for multi-view target classification with an autonomous underwater vehicle. In *Proc. IEEE OCEANS Conf.*, pages 1–5, Sept. 2010.
21. C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
22. S. D. Roy, S. Chaudhury, and S. Banarjee. Active recognition through next view planning: A survey. *J. Pattern Recognition*, 37(3):429–446, Mar. 2004.
23. A. Singh, A. Krause, C. Guestrin, and W. Kaiser. Efficient informative sensing using multiple robots. *J. Artificial Intelligence Research*, 34:707–755, Apr. 2009.
24. H. Singh, R. Armstrong, F. Gilbes, R. Eustice, C. Roman, O. Pizarro, and J. Torres. Imaging coral I: Imaging coral habitats with the SeaBED AUV. *J. Subsurface Sensing Technologies and Applications*, 5(1):25–42, Jan. 2004.
25. M. A. Sipe and D. Casasent. Feature space trajectory methods for active computer vision. *IEEE Trans. Pattern Analysis and Machine Learning*, 24(12):1634–1643, Dec. 2002.
26. R. Smith, M. Schwager, S. Smith, B. Jones, D. Rus, and G. Sukhatme. Persistent ocean monitoring with underwater gliders: Adapting sampling resolution. *J. Field Robotics*, 28(5):714–741, Sept. 2011.
27. D. Steinberg, S. B. Williams, O. Pizarro, and M. Jakuba. Towards autonomous habitat classification using Gaussian mixture models. In *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, pages 4424–4431, Oct. 2010.
28. D. Thompson and D. Wettergreen. Intelligent maps for autonomous kilometer-scale science survey. In *Proc. Int. Symp. Artificial Intelligence, Robotics and Automation in Space*, Feb. 2008.
29. S. Vasudevan, F. T. Ramos, E. W. Nettleton, and H. F. Durrant-Whyte. Gaussian process modeling of large scale terrain. *J. Field Robotics*, 26(10):812–840, Oct. 2009.
30. A. Wald. Sequential tests of statistical hypotheses. *Ann. Mathematical Statistics*, 16(2):117–186, Jun. 1945.
31. P. Whaite and F. P. Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 19:193–205, Mar. 1997.
32. D. Williams. AUV-enabled adaptive underwater surveying for optimal data collection. *Intelligent Service Robotics*, 5(1):33–54, Jan. 2012.
33. D. P. Williams. Bayesian data fusion of multiview synthetic aperture sonar imagery for seabed classification. *IEEE Trans. Image Processing*, 18(6):1239–1254, Jun. 2009.
34. D. P. Williams. On optimal AUV track-spacing for underwater mine detection. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 4755–4762, May 2010.