MIT Open Access Articles

Uncertainty-driven view planning for underwater inspection

Uncertainty-Driven View Planning for Underwater Inspection

Geoffrey A. Hollinger, Brendan Englot, Franz Hover, Urbashi Mitra, and Gaurav S. Sukhatme

Abstract— We discuss the problem of inspecting an underwater structure, such as a submerged ship hull, with an autonomous underwater vehicle (AUV). In such scenarios, the goal is to construct an accurate 3D model of the structure and to detect any anomalies (e.g., foreign objects or deformations). We propose a method for constructing 3D meshes from sonar-derived point clouds that provides watertight surfaces, and we introduce uncertainty modeling through non-parametric Bayesian regression. Uncertainty modeling provides novel cost functions for planning the path of the AUV to minimize a metric of inspection performance. We draw connections between the resulting cost functions and submodular optimization, which provides insight into the formal properties of active perception problems. In addition, we present experimental trials that utilize profiling sonar data from ship hull inspection.

I. INTRODUCTION

The increased capabilities of autonomous underwater vehicles (AUVs) have led to their use in inspecting underwater structures and environments, such as docked ships, submarines, and the ocean floor. Since these tasks are often time critical, and deployment time of AUVs is expensive, there is significant motivation to improve the efficiency of these autonomous inspection tasks. Coordinating the AUV in such scenarios is an *active perception* problem, where the path and sensor views must be planned to maximize information gained about the surface [1], [2]. We will examine a subclass of active perception problems, which we refer to as *active inspection*, where an autonomous vehicle is inspecting the surface of a 3D structure represented by a closed mesh.

An important component of active inspection is the development of a measure of *uncertainty* on the surface of the mesh. To this end, we discuss methods for generating closed surfaces that are robust to data sparsity and noise, and we propose modeling uncertainty using non-parametric Bayesian regression. Our method uses Gaussian Process Implicit Surfaces [3] with an augmented input vector [4]. The input vector is augmented with the estimated surface normals found during mesh construction to yield higher uncertainty in areas of higher variability. We also employ sparse approximation methods [5] to achieve scalability to large data sets.

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

B. Englot and F. Hover are with the Center for Ocean Engineering, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, {benglot,hover}@mit.edu

U. Mitra is with the Department of Electrical Engineering, Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089 USA, ubli@usc.edu
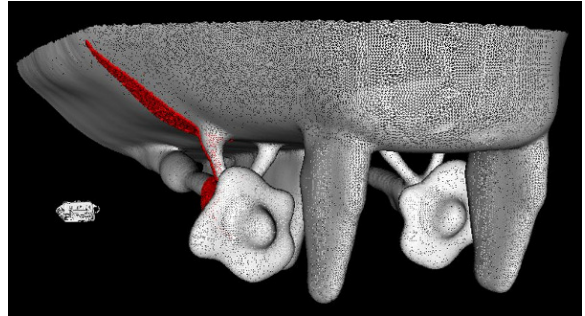
Fig. 1. Visualization of an autonomous underwater vehicle inspecting a ship hull using profiling sonar. We propose constructing a 3D mesh from the sonar range data and modeling the mesh surface using non-parametric Bayesian regression, which provides a measure of uncertainty for planning informative inspection paths.

The resulting uncertainty measures provide principled objective functions for planning paths of an AUV for further inspection. We connect the related objective functions to problems in submodular optimization, which provides insight into the underlying structure of active perception problems. The key novelties of this paper are (1) the development of mesh construction methods for sparse, noisy acoustic range data, (2) the use of Gaussian Process Implicit Surfaces with augmented input vectors to model uncertainty on mesh surfaces, and (3) the development of a probabilistic planner that maximizes uncertainty reduction while still providing coverage of the mesh surface.

The remainder of this paper is organized as follows. We first discuss related work in active perception, mesh construction, and informative path planning (Section II). We then formulate the underwater inspection problem (Section III). Next, we propose methods for constructing closed 3D meshes from acoustic range data (Section IV), and we develop techniques to represent uncertainty on meshes by extending Gaussian Process (GP) modeling (Section V). We then discuss theoretical properties, and we formulate a probabilistic path planner for minimizing uncertainty (Section VI). To test our approach, we utilize ship hull inspection data from two data sets, and we show that the proposed path planning method successfully reduces uncertainty on the mesh surface (Section VII). Finally, we conclude and discuss avenues for future work (Section VIII).

II. RELATED WORK

The active perception problem, where the goal is to plan views that best examine the environment, has seen extensive treatment throughout the robotics and computer vision communities (see Roy et al. [6] for a survey). Early

work in active vision [7] and the next-best-view problem [1] were primarily concerned with geometric approaches for determining informative views. More recent research has employed probabilistic tools, such as information gain [8]. Fewer active perception approaches have been applied to acoustic data, with a notable exception being work in medical ultrasound images [9]. While these prior works provide a basis to understand the active inspection problem, they typically do not consider uncertainty modeling or mobility restrictions in a manner appropriate to underwater inspection with an AUV.

In our prior work, we showed that an alternative formulation is to view active inspection as an instance of informative path planning [2], where a robot must gain the maximal amount of information relative to some performance metric. Informative path planning has seen rigorous theoretical analysis, utilizing the diminishing return property of submodularity [10], [11], and performance guarantees have been shown for efficient algorithms [12]. Recent advances in active learning have extended the property of submodularity to cases where the plan can be changed as new information is incorporated. The property of adaptive submodularity was introduced by [13], which provides performance guarantees in many domains that require in situ decision making. Analyzing the active inspection problem using these tools provides formal hardness results, as well as performance bounds for efficient algorithms. Thus, we gain additional insight into the nature of the active inspection problem.

Acoustic range sensing, essential in the inspection of turbid-water environments, is used to produce 2D images and 3D point clouds of underwater structures in harbor areas [14]. Laser-based range sensing, ubiquitous in ground, air, and space applications, can yield higher-resolution 3D point clouds (typically of sub-millimeter rather than sub-decimeter resolution), and specialized algorithms have been designed to generate watertight 3D mesh models from these high-resolution point clouds [15], [16]. Recently, an increasing number of tools are being developed for processing laser-based point clouds containing gaps, noise, and outliers [17], [18]. The key challenges in constructing 3D models from acoustic range data are dealing with noise and data sparsity. We extend tools previously used for laser-based modeling to generate meshes from noisy, low-resolution acoustic range data. Consequently, we focus here on the uncertainties in the mesh model rather than the navigation of the robot collecting the range data. The drift of both acoustic and inertial navigation sensors has been mitigated in inspection scenarios using localization based on sonar-frame registration [19].

To our knowledge, prior work has not considered the use of probabilistic regression to model uncertainty on the surface of a closed mesh. Similar problems, such as active 3D model acquisition [20], [21], have been examined, primarily using geometric techniques. Gaussian Processes have been used to define implicit surfaces for 2.5D surface estimation [5] and grasping tasks [3], [22], but these techniques do not actively plan to reduce uncertainty on the surface. A novelty of our approach is the use of surface normals as part of an augmented input vector to provide a better measure of surface uncertainty. The augmented input vector approach has been used to provide non-stationary kernels for Gaussian Processes [23], [4], though not in the context of surface modeling or with surface normal estimates.

III. PROBLEM FORMULATION

A number of active perception problems can be formulated as the reduction of uncertainty regarding a hypothesis class. If we define a (possibly infinite) space of classes $\mathcal{H} = \{h_0, h_1, \ldots\}$, the goal is to determine which class is being inspected. For a number of previously viewed measurements $\mathcal{Z}$, the *uncertainty* about the hypothesis space is defined by an objective function $J(\mathcal{Z})$. In the general case, this objective function relates to the entropy of the current distribution over the space of hypotheses. For the case of surface inspection, the hypothesis space is that of all possible surfaces, an infinite and high-dimensional space.

If the sensing path of the AUV is controlled, it can observe the surface from possible locations in $\mathbb{R}^d$, where d is the dimension of the space (typically 3D in underwater applications). There is a traversal cost of moving from location i to location j, which is determined by the kinematics of the vehicle and the dynamics of both the vehicle and environment. In addition, there may be an observation cost incurred when examining the surface at location i (e.g., the time taken to perform a sonar scan). The goal of active inspection is to reduce the uncertainty over the hypothesis space while also minimizing the total cost of the inspection. We note that this problem formulation fits into the general active classification framework from our prior work, which was previously used for object recognition [2].

IV. BUILDING 3D MODELS

We first address the problem of building 3D models from acoustic range data, which will later be used to derive a model of uncertainty for active inspection. We utilize data gathered using the Bluefin-MIT Hovering Autonomous Underwater Vehicle (HAUV), which was designed for autonomous surveillance and inspection of in-water ships [24]. Complex 3D structures are frequently encountered during hull inspections, particularly at the stern of the ship, where shafts, propellers, and rudders protrude from the hull. The HAUV uses a dual-frequency identification sonar (DIDSON) [25] with a concentrator lens to sample acoustic range scans for 3D modeling of these complex ship structures. The vehicle is shown in Fig. 2 along with its navigation and sensing components.

The DIDSON can perform low-resolution, long-range sensing for detection of surrounding structures as well as high-resolution, short-range sensing for identification of mines and other anomalies on the surfaces of structures. When a prior model of a structure to be inspected is unavailable, a safe-distance detection survey is conducted to build a preliminary model. Using this model, a close-range identification survey is planned to obtain coverage of the structure at higher resolution. This survey will aid in the
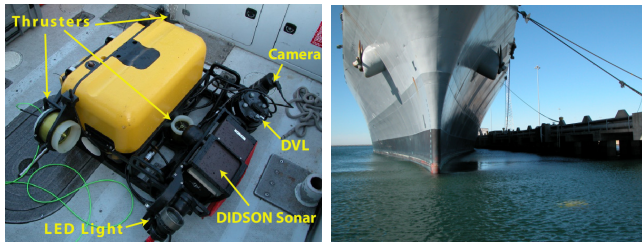
Fig. 2. Components of the Bluefin MIT-HAUV (left) and example deployment for inspection (right). The vehicle utilizes a dual-frequency identification sonar (DIDSON) to gather range scans of the ship hull. Localization is provided using a doppler velocity log (DVL) assisted by an IMU. The thrusters provide fully-actuated navigation and station-keeping capabilities.

mission-specific identification task and will also improve the accuracy of the model. Expressing uncertainty over a mesh surface will prove beneficial whenever a high-resolution identification survey must be designed using a low-resolution mesh model from a detection survey, allowing the planner to prioritize high-uncertainty areas of the model.

To construct a 3D model from the detection survey, we utilize several point cloud processing and surface construction tools from the field of laser-based modeling. All of the tools used to transform a fully dense point cloud into a 3D reconstruction can be accessed within Meshlab [26]. A fully dense point cloud of a ship hull is first obtained by applying a simple outlier filter to the individual sonar frames collected over the course of an inspection mission. All pixels of intensity greater than a specified threshold are introduced into the point cloud, and referenced using the HAUV's seafloor-relative navigation. Areas containing obvious noise and second returns are cropped out of the point cloud.

The fully dense point cloud is then sub-sampled (to about 10% of the original quantity of points) and partitioned into separate component point clouds. The partitions are selected based on the likelihood that they will yield individually well-formed surface reconstructions. Objects such as rudders, shafts, and propellers are thin objects that may not be captured in the final model without separate processing from the hull. Normal vectors are computed over the component point clouds, and some flat surfaces, for which only one of two sides was captured in the data, are duplicated. Both sub-sampling and estimation of normals are key steps in the processing sequence, found in practice to significantly impact the accuracy of the mesh [17]. Sub-sampling generates a low-density, evenly-distributed set of points, and normals aid in defining the curvature of the surface.

The Poisson surface reconstruction algorithm [27] is next applied to the oriented point clouds. Octree depth is selected to capture the detail of the ship structures without including excess roughness or curvature due to noise in the data. The component surfaces are merged back together, and a final Poisson surface reconstruction is computed over the components. If the mesh is used as a basis for high-resolution inspection planning, then it may be further subdivided to
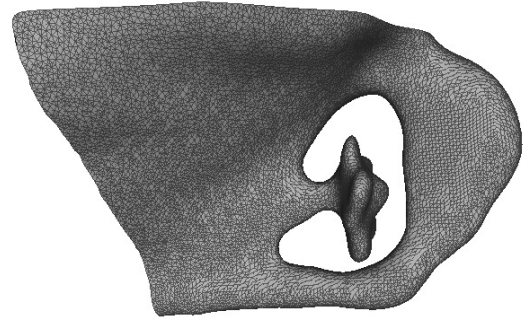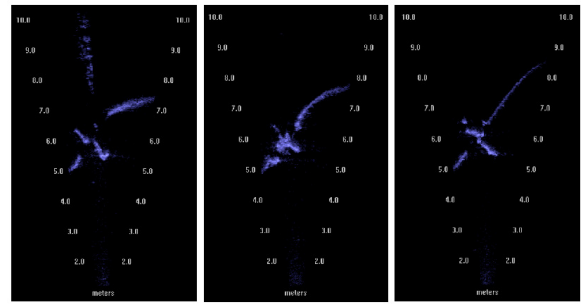

Fig. 3. An example of the raw data, and subsequent result, of the surface construction procedure. Top: three range scans featuring the hull, propeller, and rudder of the Nantucket Lightship, a historic ship 45 meters in length. These scans were collected in sequence as the HAUV dove with the running gear in view. Bottom: a 3D triangle mesh model from the complete survey of the Nantucket Lightship.

ensure the triangulation suits the granularity of the inspection task.

Fig. 3 depicts several representative range scans of a ship propeller, and the final 3D model of the ship's stern produced from the same survey. Evident in the sonar frames is the noise which makes this modeling task difficult in comparison to laser-based modeling, requiring human-in-the-loop processing to remove noise and false returns from the data.

V. REPRESENTING UNCERTAINTY

Given a mesh constructed from prior data, we propose modeling uncertainty on the surface of the mesh using non-parametric Bayesian regression. Specifically, we apply Gaussian process (GP) regression [28], though any form of regression that generates a mean and variance could be used in our framework. 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 $D = [(\mathbf{x}_1, z_1), (\mathbf{x}_2, z_2), \ldots, (\mathbf{x}_n, z_n)]$. We will first formulate the case of 2.5D surface reconstruction (i.e., the surface does not loop in on itself) and then relax this assumption in the following section. In the 2.5D case, $\mathbf{x}_i$ is a point in the 2D plane ($d = 2$), and z_i represents the height of the surface at that point. We refer to the $n \times d$ matrix of $\mathbf{x}_i$ vectors as $\mathbf{X}$ and the vector of z_i values as $\mathbf{z}$.

The next step in defining a GP is to choose a covariance

function to relate points in $\mathbf{X}$. For surface reconstruction, the choice of the kernel is determined by the characteristics of the surface. 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^2 represents the process noise, and each hyperparameter w_k represent 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^2 yields $cov(\mathbf{z}) = \mathbf{K} + \sigma_n^2\mathbf{I}$. We now wish to predict the mean function value (surface height) $\bar{f}_*$ and variance $\mathbb{V}[f_*]$ at a selected point $\mathbf{x}_*$ given the measured data:

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

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

where $\mathbf{k}_*$ is the covariance vector between the selected point $\mathbf{x}_*$ and the training inputs $\mathbf{X}$. This model provides a mean and variance at all points on the surface in $\mathbb{R}^2$. In this model, the variance gives a measure of uncertainty based on the sparsity of the data and the hyperparameters.

1) Gaussian process implicit surfaces: The GP formulation described above is limited to 2.5D surfaces (i.e., a 2D input vector with the third dimension as the output vector). Thus, it is not possible to represent closed surfaces, such as a ship hull. To apply this type of uncertainty modeling to closed surfaces, we utilize Gaussian Process Implicit Surfaces [3], [22]. The key idea is to represent the surface using a function that specifies whether a point in space is on the surface, outside the surface, or inside the surface. The implicit surface is defined as the 0-level set of the real-valued function f, where

$$f : \mathbb{R}^d \to \mathbb{R}; \ f(\mathbf{x}) \begin{cases} = 0, & \mathbf{x} \text{ on surface} \\ > 0, & \mathbf{x} \text{ outside surface} \\ < 0, & \mathbf{x} \text{ inside surface} \end{cases} \quad (4)$$

In this framework, a measurement in the point cloud at a location $\mathbf{x} = (x, y, z)$ has a value of $\mathbf{z} = 0$. We can add additional points outside the surface with $\mathbf{z} > 0$ and points inside the surface with $\mathbf{z} < 0$. The hyperparameter σ_f determines the tendency of the function to return to its mean value, which is set to a number greater than zero (i.e., outside the surface). This framework allows the representation of any surface, regardless of its 3D complexity. In general, hyperparameters can be learned automatically from the data [28]. However, learning hyperparameters for implicit surfaces requires the generation points inside and outside the surface to provide a rich enough representation of the output space. We leave the generation of these points and improved hyperparameter estimation as an avenue for future work.

2) Augmented input vectors: A limitation of the above framework is that it uses a stationary kernel, which determines correlation between points solely based on their proximity in space. Thus, areas with dense data will be have low uncertainty, and areas with sparse data will have high uncertainty. While data density is an important consideration when determining uncertainty, the amount of variability in an area should also be taken into account. For instance, the propeller of a ship hull has very complex geometry and will require a dense point cloud to reconstruct accurately.

To model surface variability, we utilize the augmented input vector approach to achieve a non-stationary kernel [23], [4]. The idea is to modify the input vector by adding additional terms that affect the correlations between the data. One measure of variability is the change in surface normal between points. From the mesh constructed in Section IV, we have an estimate of the surface normals at each point on the surface. Using this information in the augmented input vector framework, we modify the input vector to be $\mathbf{x}' = (x, y, z, \bar{n}_x, \bar{n}_y, \bar{n}_z)$, where $\bar{n}_x$, $\bar{n}_y$, and $\bar{n}_z$ are the x, y, and z components of the surface normal of the mesh at that point. We scale the surface normal to its unit vector, which reduces the possible values in the augmented input space to those between zero and one.

The weighting hyperparameters, denoted above as w_k, can be adjusted to modify the effect of spatial correlation and surface normal correlation. By modifying these hyperparameters, the user can specify how much uncertainty is applied for variable surface normals versus sparse data. If the spatial correlations are weighted more, areas with sparse data will become more uncertain. Conversely, if the surface normal correlations are weighted more, areas with high surface normal variability will become more uncertain.

3) Local approximation: The final issue to address is scalability. The full Gaussian Process models requires $O(n^3)$ computation to calculate the function mean and variance, which becomes intractable for point clouds larger than a few thousand points. To address this problem, we utilize a local approximation method using KD-trees [5]. The idea is to store all points in a KD-tree and then apply the GP locally to a fixed number of points near the test point. This method allows for scalability to large data sets.

The drawbacks of this approach are that a separate kernel matrix must be computed for each input point, and the entire point cloud is not used for mean and variance estimation. However, depending on the length scale, far away points often have low correlation and do not contribute significantly to the final output. We note that in prior work, this approach was not utilized in conjunction with implicit surfaces or augmented input vectors.

VI. INSPECTION PLANNING

A. Cost Functions

We now examine the problem of reducing the variance of the resulting Gaussian Process Implicit Surface, which will provide a performance metric for inspection planning. We denote the GP derived from the point cloud as $\mathcal{G}$. We

(a) SS Curtis without augmented input vectors

(b) SS Curtis with augmented input vectors

(c) Nantucket without augmented input vectors

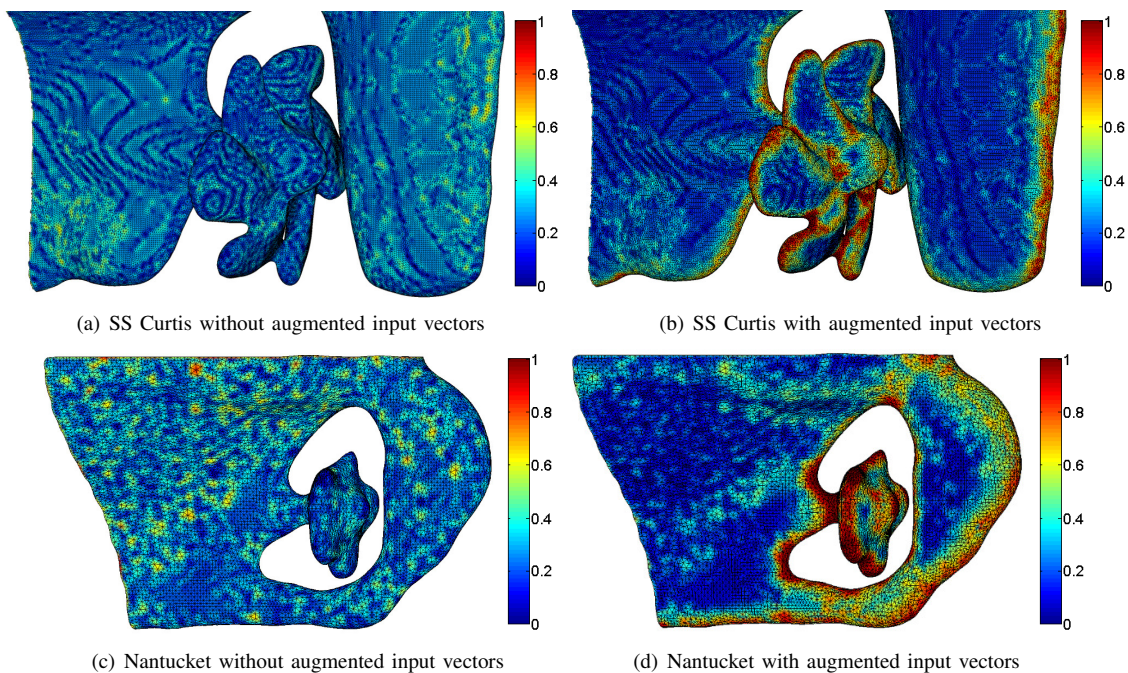(d) Nantucket with augmented input vectors

Fig. 4. Comparison of normalized uncertainty on the surface of the SS Curtiss (top) and Nantucket Lightship (bottom) meshes using a squared exponential kernel without augmented input vectors (left) and with augmented input vectors (right). Red areas are high uncertainty, and blue areas are low uncertainty. The areas of high variability have higher uncertainty when the augmented input vector is used. This figure is best viewed in color.

wish to observe additional points on the mesh surface that effectively reduce the variance. Incorporating these points will yield an updated GP $\mathcal{G}_f$. We define the following two metrics for evaluating the expected quality of $\mathcal{G}_f$:

$$J_{avg}(\mathcal{G}_f) = \int_{\mathcal{X}} \mathbb{V}[f_x] \ dx, \qquad (5)$$

$$J_{max}(\mathcal{G}_f) = \max_{\mathcal{X}} \mathbb{V}[f_x]. \qquad (6)$$

If we choose sensing locations to minimize Equation 5, we are minimizing the average variance of the GP. We note that variance reduction in GPs is a submodular optimization problem in most cases (see Das and Kempe [29]). Thus, performance guarantees from submodular function optimization can be directly applied [10]. In contrast, minimizing the maximum variance (Equation 6) is not submodular and may require alternative algorithms to optimize efficiently [11]. In our experiments, we examine the average variance reduction objective.

B. Path Planning

To plan the path of the vehicle, we utilize the sampling-based redundant roadmap method proposed in our prior work [30]. We sample a number of configurations from which the vehicle can view the surface of the mesh, and we ensure that these viewing locations are *redundant* (i.e., each point on the mesh is viewed some minimum number of times). After sampling is complete, viewing locations from the roadmap are iteratively selected as waypoints for the inspection; they are chosen using one of three approaches described in the following section. Waypoints are added until

a specified number of views has been reached or a threshold on expected uncertainty reduction is acquired.

Once the set of waypoints has been selected, the Traveling Salesperson Problem (TSP) is approximated to find a low-cost tour which visits all of the designated viewing locations. Initially, we assume that all point-to-point paths are represented by the Euclidean distances between viewing locations. The Christofides heuristic [31] is implemented to obtain an initial solution which, for the metric TSP, falls within three halves of the optimal solution cost. The chained Lin-Kernighan heuristic [32] is then applied to this solution for a designated period of time to reduce the length of the tour. All edges selected for the tour are collision-checked using the Rapidly-Exploring Random Tree (RRT) algorithm, Euclidean distances are replaced with the costs of feasible point-to-point paths, and the TSP is iteratively recomputed until the solution stabilizes (see [30] for more detail).

We assume that the vehicle does not collect data while moving between viewing locations; this is intended to accommodate the servoing precision of the HAUV. The vehicle can stabilize at a waypoint with high precision, but the precise execution of point-to-point paths in the presence of ocean disturbances is harder to guarantee. We also note that the HAUV is stable when it hovers in place, and the speed of the vehicle is constant along the path, which further motivates this abstraction.

VII. SIMULATIONS AND EXPERIMENTS

We test our proposed methods using ship hull inspection data from the Bluefin-MIT HAUV (see Section IV). This section examines two data sets: the Nantucket Lightship and

the SS Curtiss, which first appeared in our prior conference paper [30]. The Nantucket data set is composed of 21,246 points and was used to create a 42,088 triangle mesh with a bounding box of $6.3\ m \times 6.9\ m \times 4.4\ m$. The larger Curtiss data set contains 107,712 points and was used to create a 214,419 triangle mesh with a bounding box of $7.9\ m \times 15.3\ m \times 8.7\ m$.

We first examine the effect of the augmented input vectors on the uncertainty estimates. The GP hyperparameters were set to $\sigma_f = 1.0$ and $\sigma_n = 0.1$, based on the sensor noise model and the mesh shape prior. The weighting hyperparameters were set to $w_k = 1$ for all k, which provides equal weighting of data sparsity and surface variability components in the kernel. Through the use of the KD-tree local approximation with 100 adjacent points, computation was completed in approximately 1 minute for the Nantucket data set and 3 minutes for the Curtiss data set on a 3.2 GHz Intel i7 processor with 9 GB of RAM.

Fig. 4 shows a comparison of uncertainty on the surface of the two meshes with and without augmented input vectors. Utilizing the surface normals as part of an augmented input vector incorporates surface variability into the uncertainty prediction. As a result, areas with high variability require denser data to provide low uncertainty. For instance, parts of the propeller were incorrectly fused in the mesh reconstruction. The augmented input vector method shows high uncertainty at those locations, demonstrating the need for further inspection.

We now quantify the benefit of utilizing uncertainty when planning inspection paths for the HAUV. For comparison, we assume the following sensor model, which is based on the inspection strategy for high-resolution surveying with the DIDSON. At each viewing location, which is defined by the position (x, y, z) and the heading angle θ, the HAUV is assumed to sweep the DIDSON sonar through a full 180 degrees in pitch. An individual DIDSON frame spans thirty degrees in vehicle-relative heading and has a tunable minimum and maximum range. The total range of a DIDSON scan determines the resolution of the scan, and high-resolution surveying is possible when short ranges are sampled. Here we assume the DIDSON scan spans a range of $1 - 4\ m$ from the robot.

Viewing locations are generated by one of three methods: (1) coverage-biased: views are selected that maximize the number of new points observed, (2) uncertainty-biased: views are chosen to maximize the expected variance reduction on the surface, and (3) random selection: each point is selected uniformly. In all cases, viewing locations that have already been selected are excluded from consideration. The coverage-biased method represents the state of the art in inspection planning as described in our prior work [30]. For the uncertainty reduction method, it is too computationally costly to calculate the exact variance reduction for all possible viewing locations. Instead, we approximate this calculation during planning using an exponential drop off of $\exp(-v/\alpha)$, where v is the number of times a location has already been viewed in the plan, and α is a length scale parameter. The length scale parameter was estimated based on test runs where the exact reduction was calculated.

The sampled points are then connected using the TSP/RRT method described in Section VI. This method generates a complete inspection path that can be executed by the vehicle. The total mission time is calculated assuming that the AUV requires 5 seconds to perform a scan from a viewing location and moves between locations at a speed of $0.25\ m/s$. These numbers are taken directly from the experimental trials with the Bluefin-MIT HAUV. To evaluate the quality of the inspection paths, simulated measurements are generated by creating additional data points on the viewed parts of the mesh, and the surface normals of the new points are estimated from the original mesh. An estimate of expected uncertainty is then calculated by re-running the GP and determining the reduction in variance on the original points. The expected uncertainty reduction is a unitless quantity, which incorporates both data sparsity and surface variability.

Fig. 5 shows a quantitative comparison of the three view selection methods. We see that view selection based on uncertainty reduction leads to greater reduction in variance for a given mission time. The coverage-biased selection method provides competitive uncertainty reduction for shorter paths (i.e., good coverage implicitly leads to good uncertainty reduction), but it does not allow for continued planning after full coverage is achieved. The performance of the random selection method improves (versus the uncertainty-biased method) as mission time increases, due to the amount of possible uncertainty reduction being finite. In the larger mesh, random view selection does not perform as well, even with long mission times. This improvement is due to redundancy in viewing high uncertainty areas, which leads to additional benefit. Fig. 6 shows example paths on each data set using the uncertainty-biased method.

It is expected that the method that takes into account uncertainty would lead to greater variance reduction; however, we may be sacrificing some surface coverage to achieve this additional uncertainty reduction. Fig. 7 shows a comparison of the mission time vs. percent coverage for the three view selection methods. We see that the uncertainty reduction method converges quickly to greater than 99% coverage, and reaches 100% coverage long before random sampling.[1] Uncertainty-biased view selection gets to 99% coverage nearly as quickly as coverage-biased view selection, requiring only 41 seconds longer on the Nantucket mesh and 31 seconds longer on the Curtiss mesh, and it does so with a greater reduction in uncertainty. Thus, the uncertainty reduction method provides both high levels of coverage and improved uncertainty reduction for a given mission time.

VIII. CONCLUSIONS AND FUTURE WORK

We have shown that it is possible to construct closed 3D meshes from noisy, low-resolution acoustic ranging data, and we have proposed methods for modeling uncertainty

---

[1]We note that since variance reduction is monotonic, uncertainty-biased view selection is guaranteed to achieve 100% coverage in the limit.
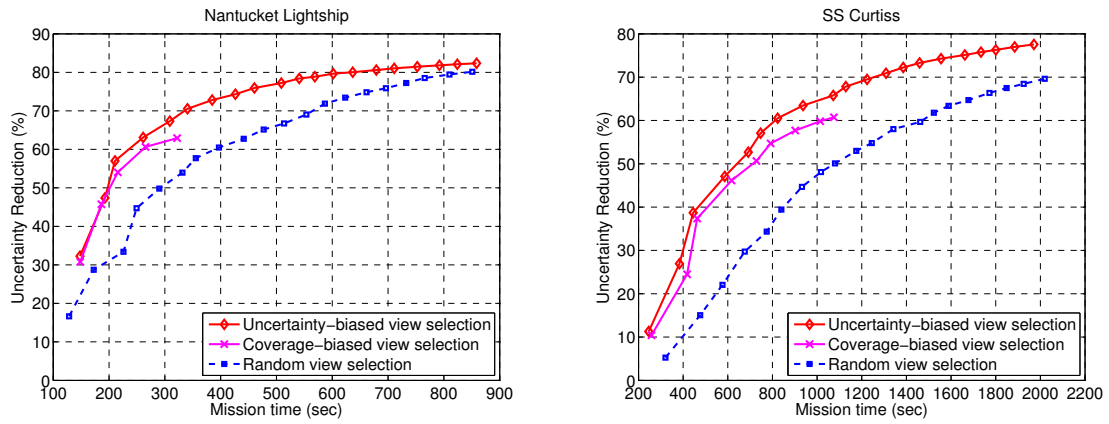
Fig. 5. Mission time versus expected uncertainty reduction for inspection simulations. Uncertainty-biased view selection provides improved uncertainty reduction for a given mission time. Coverage-biased view selection does not allow for planning after full coverage is achieved, while the other methods continue to reduce uncertainty. Uncertainty reduction is a unitless quantity that takes into account data sparsity and surface variability; it is displayed as a percent reduction from its initial value.
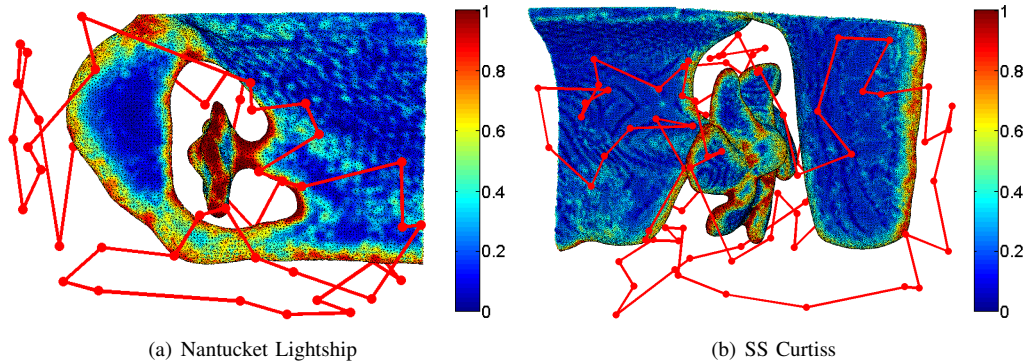

(a) Nantucket Lightship

(b) SS Curtiss

Fig. 6. Planned inspection paths on two ship hull mesh data sets. Uncertainty-biased view selection provides an inspection path that views areas with sparse data and high variability. This figure is best viewed in color.
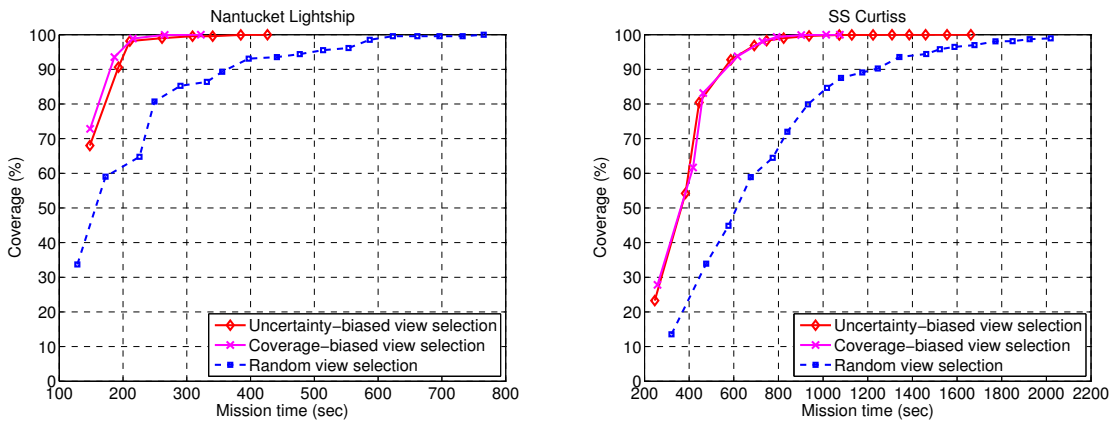

Fig. 7. Mission time versus percent coverage for inspection simulations. Plots are truncated after full coverage is achieved. Coverage-biased view selection and uncertainty-biased view selection both converge to full coverage. Random view selection does not achieve full coverage on the SS Curtiss mesh.

on the surface of the mesh using an extension to Gaussian Processes. Our techniques utilize surface normals from the 3D mesh to develop estimates of uncertainty that account for both data sparsity and high surface variability, and we achieve scalability to large environments through the use of local approximation methods. We have also shown that

probabilistic path planning can be used to generate views that effectively reduce the uncertainty on the surface in an efficient manner. This research moves towards formal analysis of active sensing, through connections to submodular optimization. Such analysis allows us to gain better understanding of the problem domain and to design principled algorithms for

coordinating the actions of autonomous underwater vehicles.

An interesting avenue for future work is to examine alternative kernels for representing correlations between points on the mesh surface. The neural network kernel has been employed in prior work and was shown to provide improved performance for modeling discontinuous surfaces [5]. However, this kernel does not have a simple geometric interpretation, and its direct application to the implicit surface model does not produce reasonable uncertainty predictions. The automatic determination of hyperparameters for Gaussian Process Implicit Surfaces is another area for future work. Utilizing the 3D point cloud on its own does not provide sufficient information to learn the observation noise and process noise hyperparameters. To learn these automatically, a method would need to be developed to generate points both inside and outside the surface.

Additional open problems include further theoretical analysis of performance guarantees, particularly in the case of path constraints. The algorithm utilized in this paper selects views based on uncertainty reduction and then connects them using a TSP/RRT planner. A planner that utilizes the path costs when planning the view selection could potentially perform better, particularly in the case of a small number of views. Finally, the analysis in this paper has applications beyond underwater inspection. Tasks such as ecological monitoring, reconnaissance, and surveillance are just a few domains that would benefit from active planning for the most informed views. Through better control of the information we receive, we can improve the understanding of the world that we gain from robotic perception.

IX. ACKNOWLEDGMENTS

REFERENCES

[1] C. Connolly, "The determination of next best views," in *Proc. IEEE Conf. Robotics and Automation*, 1985.

[2] G. Hollinger, U. Mitra, and G. S. Sukhatme, "Active classification: Theory and application to underwater inspection," in *Int. Symp. Robotics Research*, 2011.

[3] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," in *Proc. Gaussian Processes in Practice Workshop*, 2007.

[4] D. Thompson and D. Wettergreen, "Intelligent maps for autonomous kilometer-scale science survey," in *Proc. Int. Symp. Artificial Intelligence, Robotics and Automation in Space*, 2008.

[5] 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.

[6] S. D. Roy, S. Chaudhury, and S. Banarjee, "Active recognition through next view planning: A survey," *J. Pattern Recognition*, vol. 37, no. 3, pp. 429–446, 2004.

[7] Y. Aloimonos, I. Weiss, and A. Bandopadhyay, "Active vision," *International Journal of Computer Vision*, vol. 1, no. 4, pp. 333–356, 1987.

[8] J. Denzler and C. Brown, "Information theoretic sensor data selection for active object recognition and state estimation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 2, pp. 145–157, 2002.

[9] X. S. Zhou, D. Comaniciu, and A. K. Krishnan, "Conditional feature sensitivity: A unifying view on active recognition and feature selection," in *Proc. Int. Conf. Computer Vision*, 2003.

[10] A. Krause, A. Singh, and C. Guestrin, "Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies," *J. Machine Learning Research*, vol. 9, pp. 235–284, 2008.

[11] A. Krause, B. McMahan, C. Guestrin, and A. Gupta, "Robust submodular observation selection," *J. Machine Learning Research*, vol. 9, pp. 2761–2801, 2008.

[12] A. Singh, A. Krause, C. Guestrin, and W. Kaiser, "Efficient informative sensing using multiple robots," *J. Artificial Intelligence Research*, vol. 34, pp. 707–755, 2009.

[13] D. Golovin and A. Krause, "Adaptive submodularity: A new approach to active learning with stochastic optimization," in *Proc. Conf. Learning Theory*, 2010.

[14] D. M. Kocak, F. R. Dalgleish, F. M. Caimi, and Y. Y. Schechner, "A focus on recent developments and trends in underwater imaging," *Marine Technology Society J.*, vol. 42, no. 1, pp. 52–67, 2008.

[15] B. Curless and M. Levoy, "A volumetric method for building complex models from range images," in *Proc. Conf. Computer Graphics and Interactive Techniques*, 1996, pp. 303–312.

[16] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald, and W. Stuetzle, "Surface reconstruction from unorganized points," in *Proc. Conf. Computer Graphics and Interactive Techniques*, 1992, pp. 71–78.

[17] H. Huang, D. Li, H. Zhang, U. Ascher, and D. Cohen-Or, "Consolidation of unorganized point clouds for surface reconstruction," *ACM Trans. Graphics*, vol. 28, no. 5, 2009.

[18] T. Weyrich, M. Pauly, R. Keiser, S. Heinzle, S. Scandella, and M. Gross, "Post-processing of scanned 3d surface data," in *Proc. IEEE Eurographics Symp. Point-Based Graphics*, 2004, pp. 85–94.

[19] H. Johannsson, M. Kaess, B. Englot, F. Hover, and J. Leonard, "Imaging sonar-aided navigation for autonomous underwater harbor surveillance," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2010, pp. 4396–4403.

[20] S. Y. Chen and Y. F. Li, "Vision sensor planning for 3-d model acquisition," *IEEE Trans. Systems, Man, and Cybernetics*, vol. 35, no. 5, pp. 894–904, 2005.

[21] E. Dunn, J. van den Berg, and J.-M. Frahm, "Developing visual sensing strategies through next best view planning," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots and Systems*, 2009, pp. 4001–4008.

[22] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *IEEE Int. Conf. Robotics and Automation*, 2011.

[23] T. Pfingsten, M. Kuss, and C. E. Rasmussen, "Nonstationary gaussian process regression using a latent extension of the input space," in *Proc. ISBA World Mtg. Bayesian Statistics*, 2006.

[24] F. Hover, J. Vaganay, M. Elkins, S. Wilcox, V. Polidoro, J. Morash, R. Damus, and S. Desset, "A vehicle system for autonomous relative survey of in-water ships," *Marine Technology Society J.*, vol. 41, no. 2, pp. 44–55, 2007.

[25] E. Belcher, W. Hanot, and J. Burch, "Dual-frequency identification sonar (DIDSON)," in *Proc. Int. Symp. Underwater Technology*, 2002, pp. 187–192.

[26] P. Cignoni, M. Corsini, and G. Ranzuglia, "Meshlab: An open-source 3D mesh processing system," *ERCIM News*, vol. 73, pp. 45–46, 2008.

[27] M. Kazhdan, M. Bolitho, and H. Hoppe, "Poisson surface reconstruction," in *Proc. Eurographics Symp. Geometry Processing*, 2006, pp. 61–70.

[28] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[29] A. Das and D. Kempe, "Algorithms for subset selection in linear regression," in *Proc. ACM Symp. Theory of Computing*, 2008, pp. 45–54.

[30] B. Englot and F. Hover, "Planning complex inspection tasks using redundant roadmaps," in *Proc. Int. Symp. Robotics Research*, 2011.

[31] N. Christofides, "Worst-case analysis of a new heuristic for the traveling salesman problem," Computer Science Dept., Carnegie Mellon Univ., Tech. Rep. CS-93-13, 1976.

[32] D. Applegate, W. Cook, and A. Rohe, "Chained lin-kernighan for large traveling salesman problems," *INFORMS J. Computing*, vol. 15, no. 1, pp. 82–92, 2003.