

Path and View Planning for Underwater Surface Inspection

Geoffrey A. Hollinger, Gaurav S. Sukhatme
Computer Science Dept.
University of Southern California
Los Angeles, CA 90089 USA
Email: {gahollin,gaurav}@usc.edu

Brendan Englot, Franz Hover
Dept. of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139 USA
Email: {benglot,hover}@mit.edu

Urbashi Mitra
Electrical Engineering Dept.
University of Southern California
Los Angeles, CA 90089 USA
Email: ubli@usc.edu

Abstract—We discuss the problem of active inspection of a marine structure by an autonomous underwater vehicle (AUV). The AUV must plan a path of sensing views that estimates the difference between two surfaces: one reconstructed from prior data, and one reconstructed online during inspection. We propose modeling the structure’s surface using non-parametric Bayesian regression, which allows for a number of novel cost functions for planning the path of the AUV to minimize some metric of change detection between reconstructions. We analyze the resulting cost functions using the theoretical tools of adaptivity gaps and (adaptive) submodularity, and we draw connections to the more general active classification problem. In addition, we discuss experimental trials that utilize profiling sonar data from a ship hull inspection. This work moves towards formal analysis of active sensing for inspection tasks.

I. INTRODUCTION

There is growing interest in the use of autonomous underwater vehicles (AUVs) to inspect marine structures, such as docked ships and submarines. 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 is an *active sensing* problem, where the path and sensor views must be planned to maximize information gained about the surface [2, 5].

We examine the problem of coordinating an AUV to detect changes in underwater structures with a profiling sonar sensor. We propose modeling the surface using non-parametric Bayesian regression, which provides principled methods for modeling sensor noise and uncertainty in the surface reconstruction. If prior surface data is provided, the active sensing problem becomes one of reconstructing a second surface and determining if relevant changes exist between the surfaces.

To our knowledge, utilizing probabilistic regression to detect changes in surface structure has not been examined in prior work. Similar problems, such as active 3-D model acquisition [1], have been examined, primarily using geometric techniques. Gaussian processes have been used to define implicit surfaces for grasping tasks [3]. However, this prior work does not actively sense for changes in different surfaces.

Surface change detection can be seen as an instance of informative path planning [9], where a robot must gain the maximal amount of information relative to some performance




Fig. 1. A mesh reconstructed using profiling sonar data from an autonomous underwater vehicle. We propose modeling the mesh surface using non-parametric Bayesian regression, which provides principled objective functions for change detection on the structure.

metric. Many relative metrics of informativeness satisfy the theoretical property of *submodularity*. Submodularity is a rigorous characterization of the intuitive notion of diminishing returns that arises in many active planning applications [7]. 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 Golovin and Krause [4], which provides performance guarantees in many domains that require adaptive decision making.

Our framework allows for formal analysis of change detection objective functions and active sensing algorithms in the context of these rigorous theoretical properties. Such analysis allows us to (1) better understand the underlying structure of active sensing problems and (2) design principled planning algorithms with performance guarantees.

II. PROBLEM FORMULATION

The goal is to determine if a surface has changed relative to a model built from prior inspection. We define a number of possible surface classes $\mathcal{H} = \{h_0, \dots, h_N\}$. Let H be a random variable equal to the true class of the surface. For the binary classification case, h_0 represents an unchanged surface, and h_1 represents a changed surface (e.g., a damaged ship hull or a ship hull with a foreign device placed on it). In the multi-class problem, surface h_1 may represent one type of ship hull, and h_2, \dots, h_N may represent additional types of ship hulls.

We control the sensing path of an AUV, which can observe the surface from a set of possible locations $\mathcal{L} = \{L_1, \dots, L_M\}$. There is a cost of moving from location L_i to location L_j , which we denote as d_{ij} . This cost is determined by the kinematics of the vehicle and the dynamics of both the vehicle and environment. The locations allow us to view some features of the surface, which distinguish between the possible classes. This formulation fits into the more general active classification framework from our prior work [5].

III. METHODS AND ANALYSIS

We propose to model change detection on surfaces using non-parametric Bayesian regression. Specifically, we apply Gaussian process (GP) regression [8], though any form of regression that generates a mean and variance could be used in our framework. A GP models a noisy process with an underlying noise model as $z_i = f(x_i) + \varepsilon$.

We are given some data of the form $D = [(x_1, z_1), (x_2, z_2), \dots, (x_n, z_n)]$ where $x_i \in \mathbb{R}^d$ and $z_i \in \mathbb{R}$. In the case of surface reconstruction, 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 matrix of x_i values as X and the vector of z_i values as Z . The next step in defining a GP is to choose a covariance function to relate points in X . For surface reconstruction, the choice of the kernel is determined by the characteristics of the surface. The commonly used squared exponential is a smooth kernel, which is applicable if changes in the surface produce smooth deformations (e.g., in the case of hull damage or warping). If the goal is to locate an object placed on the surface, alternative kernels should be used. We plan to investigate the use of alternative kernels for various applications.

Once the kernel has been defined, combining the covariance values for all points into a matrix K and adding a Gaussian observation noise hyperparameter σ_n^2 yields $\text{cov}(Z) = K + \sigma_n^2 I$. We now wish to predict the function value (surface height) $f(x_*)$ at a selected point x_* given the measured data:

$$p(f(x_*)|x_*, X, Z) = \mathcal{N}(f(x_*); \mu_{x_*}, \sigma_{x_*}^2), \quad (1)$$

where

$$\mu(x_*) = k_*^T (K + \sigma_n^2)^{-1} Z, \quad (2)$$

$$\sigma(x_*)^2 = k(x_*, x_*) - k_*^T (K + \sigma_n^2)^{-1} k_*, \quad (3)$$

and k_* is the covariance vector between the selected point x_* and the training inputs X .

When estimating the likelihood of a measurement z_* at a selected point x_* , we must also add the observation noise $p(z_*|x_*, X, Z) = \mathcal{N}(z_*; \mu_{x_*}, \sigma_{x_*}^2 + \sigma_n^2)$. This model provides a mean and variance at all points on the surface in \mathbb{R}^2 .

In the context of change detection, we define one GP \mathcal{G}_0 that was derived from prior data (e.g., a lengthy inspection when the ship was docked at home port). We now wish to construct a new GP \mathcal{G}_1 and determine if relevant changes have occurred on the surface. For a given point x , let $\mu_1(x), \sigma_1(x)^2$ and $\mu_2(x), \sigma_2(x)^2$ be the mean and variance of the two GPs

respectively. We define a new GP \mathcal{G}_c as the difference between the surfaces with $\mu_c(x) = |\mu_1(x) - \mu_2(x)|$ and $\sigma_c(x)^2 = \sigma_1(x)^2 + \sigma_2(x)^2$.

To guide active sensing, we now define two metrics for evaluating the quality of \mathcal{G}_c :

$$J(\mathcal{G}_c) = \int_{\mathcal{X}} \sigma_c(x)^2 dx \quad (4)$$

$$J(\mathcal{G}_c) = \max_{\mathcal{X}} \sigma_c(x)^2 dx \quad (5)$$

If we choose sensing locations to minimize Equation 4, we are minimizing the average variance of the GP. Since we only have control over \mathcal{G}_1 (\mathcal{G}_0 is set from prior data), this is equivalent to minimizing the average variance of the GP. We note that this is a submodular optimization problem [7]. Thus, planning and view selection approximation algorithms designed for submodular objectives functions should work well in this setting. In contrast, the minimizing the maximum variance (Equation 5) is not submodular and may require alternative algorithms to optimize efficiently.

We will also analyze the benefit of adaptivity for these objective functions and others. We conjecture that these objectives will not benefit from adaptive selection (i.e., it does not help to consider new measurements as they become available). In contrast, an objective that seeks to minimize the probability of missing a large discontinuity would likely benefit from adaptive planning.

IV. EXPERIMENTS AND DISCUSSION

We will validate our proposed method both in simulation and using ship hull inspection data from a Bluefin AUV with a profiling sonar [6]. We will examine the contrasting cases of explosive device detection (locating sharp discontinuities) and deformed surface detection (locating smooth discontinuities). In addition, we will apply our techniques to bathymetry change detection using a YSI Ecomapper platform.

By developing a common framework for active classification and reconstruction, this research moves towards formal analysis of active sensing. Such analysis allows us to gain better understanding of the problem domain and to design principled algorithms for coordinating the sensing actions of autonomous underwater vehicles.

REFERENCES

- [1] S. Y. Chen and Y. F. Li. Vision sensor planning for 3-d model acquisition. *IEEE Trans. Systems, Man, and Cybernetics*, 35(5):894–904, 2005.
- [2] C. Connolly. The determination of next best views. In *Proc. IEEE Conf. Robotics and Automation*, 1985.
- [3] S. Dragiev, M. Toussaint, and M. Gienger. Gaussian process implicit surfaces for shape estimation and grasping. In *IEEE Int. Conf. Robotics and Automation*, 2011.
- [4] D. Golovin and A. Krause. Adaptive submodularity: A new approach to active learning with stochastic optimization. In *Proc. Conf. Learning Theory*, 2010.

- [5] G. Hollinger, U. Mitra, and G. S. Sukhatme. Active classification: Theory and application to underwater inspection. In *under review*, 2011.
- [6] M. Kaess, H. Johannsson, B. Englot, F. Hover, and J. Leonard. Towards autonomous ship hull inspection using the bluefin HAUV. In *Proc. Int. Symp. Technology and The Mine Problem*, 2010.
- [7] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *J. Machine Learning Research*, 2008.
- [8] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [9] A. Singh, A. Krause, C. Guestrin, and W. Kaiser. Efficient informative sensing using multiple robots. *J. Artificial Intelligence Research*, 34:707–755, 2009.