Multi-Robot Informative Planning for Long Term Ocean Monitoring

Kai-Chieh Ma, Lantao Liu, Gaurav S. Sukhatme

***Abstract*— We present an informative path planning method for multiple autonomous underwater vehicles (AUVs) used for long-range and long-term ocean monitoring. We consider the spatio-temporal variations of ocean phenomena, and develop an information-driven approach that computes the most informative observation way-points for reducing the uncertainty for ocean modeling and prediction. The sampling paths of AUVs are then formulated and solved through a matching graph based routing method, which allows the vehicles to transit the obtained informative way-points in an efficient and interference-free way. We provide preliminary simulation results to validate the proposed method.**

I. Introduction

We propose an informative path planning method that guides a team of autonomous underwater vehicles (AUVs) to collect ocean data in an efficient way. By efficiency we mean the "informativeness" of collected data (i.e., reduction of phenomena modeling uncertainty) as well as the minimization of energy and time used to collect the data.

We employ a Gaussian Process to model an underlying phenomenon, and utilize the mutual information between visited locations and the remainder of the space to characterize the amount of information collected. Such a formulation allows us to obtain a set of "most informative" future observation/sampling points in the environment. However, these observation points do not form a path (or multiple paths), since no way-point traversal order/sequence information is provided. Therefore, we employ and modify an existing matching graph based routing method [2], which produces a set of interference-free paths leading all AUVs to mutual exclusive goal locations. At the lower level, we also consider the AUV's action uncertainty due to disturbances caused by the non-stationary ocean currents, and extend the Markov Decision Process in continuous space to control AUV's motion.

We validated the method through simulations with real ocean data. The preliminary results show that, for a single robot, the method not only maximizes information-gain but also saves time and energy while exploring the non-stationary ocean; for multiple robots, our method is flexible (i.e., with tunable parameters) in generating interference-free routing paths that are also time and energy efficient. Our ongoing work involves further assessment of the information-gain performance for different routing results tuned with different path parameters.

The authors are with the Department of Computer Science at the University of Southern California, Los Angeles, CA 90089, USA. {kaichiem, lantao.liu, gaurav}@usc.edu

II. Preliminaries

A. Gaussian Process based Uncertain Field

To model spatial phenomena, a common approach in spatial statistics is to use a rich class of Gaussian Processes [4, 6].

Formally, let W denote a set of sampling points describing the environmental phenomenon of interest. Each point $w \in W$ is a d-dimensional feature vector associating with either a realized measurement z_w if observed (sampled) or a random measurement Z_w if unobserved. Let set $\{Z_w\}_{w \in W}$ denote a GP, then for every finite subset of $\{Z_w\}_{w \in W}$, it has a multivariate Gaussian distribution. The GP can be fully specified by its mean $\mu_w \triangleq \mathbb{E}(Z_w)$ and covariance $\sigma_{ww'|\theta} \triangleq cov(Z_w, Z_{w'}|\theta)$ for all $w, w' \in W$, where θ parameterizes the covariance function which models the spatial phenomenon (parameterization details are presented later).

Assume we are given an observed data set $D = \{(w_i, Z_{w_i}), i = 1 : |D|\}$, where $D \subset W$. GP can be used to predict the mean and covariance of measurements for any unobserved subset of $U \subset W \setminus D$. Based on the property that every subset of $\{Z_w\}_{w \in W}$ is a multivariate Gaussian distribution, the joint distribution of Z_U and Z_D can therefore be expressed as:

$$\begin{pmatrix} Z_D \\ Z_U \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_D \\ \mu_U \end{pmatrix}, \begin{pmatrix} K_{DD} & K_{DU} \\ K_{DU}^T & K_{UU} \end{pmatrix} \right). \tag{1}$$

where

$$\begin{aligned} K_{DD} &= \left(\sigma_{ww'|\theta}\right)_{|D|\times|D|}, \quad \forall w, w' \in D, \\ K_{DU} &= \left(\sigma_{ww'|\theta}\right)_{|D|\times|U|}, \quad \forall w \in D, w' \in U, \\ K_{UU} &= \left(\sigma_{ww'|\theta}\right)_{|U|\times|U|}, \quad \forall w, w' \in U. \end{aligned} \tag{2}$$

We then obtain the Gaussian posterior mean and covariance,

$$\begin{aligned} \mu_{U|D,\theta} &= \mu(W_U) + K_{DU}^T K_{DD}^{-1}(Z_D - \mu(W_D)), \\ \Sigma_{U|D,\theta} &= K_{UU} - K_{DU}^T K_{DD}^{-1} K_{DU}, \end{aligned} \tag{3}$$

where W_U, W_D denote the set of sampling points in U, D respectively, and Z_D denotes the realized measurements of D. Note that the posterior covariance matrix $\Sigma_{U|D,\theta}$ is independent of the measurements and it can be used to assess the uncertainty with respect to the predicted measurements.

A GP's behavior is controlled via specifying its prior covariance (also known as *kernel*) $\sigma_{ww'|\theta}$, which describes the relation between sampling points w and w'. A widely adopted choice is the *squared exponential* kernel function:

$$\sigma_{ww'|\theta} = \sigma_s^2 \exp(-\frac{1}{2}(w-w')^T\Lambda^{-1}(w-w')) + \sigma_n^2\delta_{ww'} \tag{4}$$

where $\Lambda = diag(l_1^2, \ldots, l_d^2)$ and $\theta = \{\sigma_s^2, \sigma_n^2, l_1^2, \ldots, l_d^2\}$ is the set of hyper-parameters specifying the property of the w, w' pairwise relation. The parameters $l_1 \ldots l_d$ are the *length-scales* in each dimension of w and determine the level of correlation between points (each l_i models the degree of smoothness in the spatial variation of the measurements in the ith dimension of the feature vector w). σ_s^2 and σ_n^2 denote the variances of the signal and noise, respectively. $\delta_{ww'}$ is the Kronecker delta function which is 1 if $w = w'$ and zero otherwise.

B. Entropy and Mutual Information

To assess the level of measurement and prediction uncertainty, we adopt the concept of entropy and mutual information. Formally, given a vector of sampling points A of size k, the joint *differential entropy* of the corresponding vector Z_A of random measurements is

$$\begin{aligned} H(Z_A) &= -\int p(Z_A)\log p(Z_A) d(Z_A) \\ &= \frac{1}{2}\log\Big((2\pi e)^k |\Sigma_{AA}|\Big). \end{aligned} \tag{5}$$

For arbitrary two vectors of sampling points A, B, the mutual information between A and B can be expressed in terms of (conditional) entropy

$$I(Z_A; Z_B) = I(Z_B; Z_A) = H(Z_A) - H(Z_A|Z_B) \tag{6}$$

where the conditional entropy $H(Z_A|Z_B)$ can be calculated via

$$H(Z_A|Z_B) = \frac{1}{2}\log\Big((2\pi e)^k |\Sigma_{A|B}|\Big). \tag{7}$$

Because the field is modeled with a Gaussian Process, the conditional covariance matrix $\Sigma_{A|B}$ can essentially be calculated from the posterior covariance matrix described in Eq. (3).

III. Technical Approach

With some initial knowledge/condition, the planner computes a path for the AUV which gives us the most additional information (reducing predictive uncertainty). However, as time elapses, the uncertainty of visited/measured regions may increase again. The informative paths are thus repetitively generated based on the spatially and temporally varying ocean phenomenon.

A. Environment Representations & Methodology Framework

We discretize the ocean environment into a grid map with certain resolution. Each grid at a location stores a mean value that predicts the phenomenon interested as well as a variance that measures the uncertainty of such prediction. We assume noise-free observation and when a location is sampled/observed, the variance of this grid is reduced to a very small value.

Note that, the basic informativeness maximization procedure (details are presented later) only outputs a batch of observation points, but does not directly produce paths. Therefore, we post-process these observation points with our multi-robot routing planner to generate multiple paths and route each AUV from its initial location to the specified destination locations.

B. Generation of Informative Sampling Way-points

Given a desired total number of n way-points we want to generate. Let W denote the whole sampling set of a map. The objective is to find a subset of sampling points, $P \subset W$ with a size $|P| = n$, which gives us the most information. It's equivalent to finding the maximum of the mutual information between the selected subset and the rest (unobserved part) of the map.

Thus the optimal P^* with maximal mutual information is

$$P^* = \underset{P\in\mathcal{X}}{\arg\max}\, I(Z_P; Z_{W\setminus P}) \tag{8}$$

where $\mathcal{X}$ represents all possible combinatorial sets, each of which has a size n in the map. In order to find P^*, one way is to enumerate all possible combinations in $\mathcal{X}$. The time complexity of this naive approach would be exponential and hence intractable in practice. Fortunately, P^* can be computed by following a *dynamic programming* structure which dramatically reduces the computational time. Details are described below.

Let $w_i \in W$ denote an arbitrary sampling point at stage i and $w_{a:b}$ represent a sequence of sampling points from stage a to stage b. With Eq. (6), the mutual information between P and the unobserved part at the final stage n can then be written as $I(Z_{w_{1:n}}; Z_{W\setminus\{w_{1:n}\}})$. This mutual information can be expanded using the chain rule:

$$\begin{aligned} I(Z_{w_{1:n}}; Z_{W\setminus\{w_{1:n}\}}) &= I(Z_{w_1}; Z_{W\setminus\{w_{1:n}\}}) \\ &+ \sum_{i=2}^{n} I(Z_{w_i}; Z_{W\setminus\{w_{1:n}\}}|Z_{w_{1:i-1}}). \end{aligned} \tag{9}$$

One can utilize this form of mutual information to calculate w_i step by step, however, at every stage i before the final stage, the entire unobserved set $W \setminus \{w_{1:n}\}$ is not known in advance, therefore we make an approximation

$$\begin{aligned} I(Z_{w_{1:n}}; Z_{W\setminus\{w_{1:n}\}}) &\approx I(Z_{w_1}; Z_{W\setminus\{w_1\}}) \\ &+ \sum_{i=2}^{n} I(Z_{w_i}; Z_{W\setminus\{w_1,\ldots,w_i\}}|Z_{w_{1:i-1}}), \end{aligned} \tag{10}$$

which can be formulated with a recursive form, i.e., for stages $i = 2, \ldots, n$, the value $V_i(w_i)$ of w_i is:

$$\begin{aligned} V_i(w_i) = &\max_{w_i\in W\setminus\{w_1,\ldots,w_{i-1}\}} I(Z_{w_i}; Z_{W\setminus\{w_1,\ldots,w_i\}}|Z_{w_{1:i-1}}) \\ &+ V_{i-1}(w_{i-1}), \end{aligned} \tag{11}$$

and the base case for this recursion is:

$$V_1(w_1) = I(Z_{w_1}; Z_{W\setminus\{w_1\}}). \tag{12}$$

Note the optimal way-point in the last stage n is

$$w_n^* = \arg\max_{w_n \in W} V_n(w_n). \tag{13}$$

With the optimal solution in the last stage, w_n^*, we can backtrace all optimal sampling points till the first stage w_1^*, and get the whole set of observation points $w^* = \{w_1^*, w_2^*, \ldots, w_n^*\}$. Note, these observation points contain no path information, and paths will be computed with a routing strategy presented in the following subsection.

C. Way-point Routing for Multi-Robot Systems

With the observation points obtained from Sect. III-B, the interference-free routing paths for multiple vehicles to transit these way-points are planned. This is essentially a multi-source multi-goal (MSMG) path planning problem.

Let $G = (V, E)$ be the graph used for routing the AUVs, where vertex set V contains nodes represented by the obtained observation points, and edge set E contains edges with weights that measure the Euclidean distances (travel costs) to nearby nodes (possibly within some predefined radius). One efficient way to solve an MSMG is through transforming $G = (V, E)$ to a bipartite (matching) graph $\tilde{G} = (V, V', \tilde{E})$. Manipulation details are presented in [2]. The essence is recapitulated as follows.

Briefly, a bipartite graph $\tilde{G}$ has two sets of nodes V and V', where V' is simply a copy of V such that $|V| = |V'|$, and an edge $\tilde{e} = (v_i, v_j') \in \tilde{E}$ connects the vertices $v_i \in V$ and $v_j' \in V'$ if there is an edge $e = (v_i, v_j) \in E \in G$. Edge $\tilde{e} = (v_i, v_j')$ is weighted the same as the counterpart edge (v_i, v_j)

$$\tilde{w}(v_i, v_j') = \tilde{w}(v_i', v_j) = w(v_i, v_j) \tag{14}$$

Besides that, a set of edges (v_i, v_i') is also added with some preset feasible weight $\tilde{w}(v_i, v_i')$ for all states except the starts and goals (weight $\tilde{w}(v_i, v_i')$ can be scaled or tuned, see details below).

A bipartite graph of this form well represents the assignment problem and can be solved by the Hungarian Algorithm, which manipulates *augmenting paths* consisting of *matched* and *unmatched* edges in order to find a solution where each vertex in V is uniquely matched (assigned) to a vertex in V' with the total cost minimized. Our MSMG trajectories are obtained by transforming the resultant augmenting paths back to the routing paths on G via eliminating all vertices except the goals from vertex set V'. The concept is illustrated in Fig. 1.

For multiple starting and ending vertices, multiple paths can be obtained. Since each vertex can not simultaneously be on more than one augmenting path (a property in the Hungarian algorithm), the resulting routing paths are interference-free with no shared vertex. Note that, however, these interference-free MSMG paths do not transit all nodes in V (different from routing strategies such as the Traveling Salesman Problem). A useful property of the MSMG method lies in that the paths can be tuned (adjusted) via manipulating the weights $\tilde{w}(v_i, v_i')$ of the corresponding *vertical* edges in

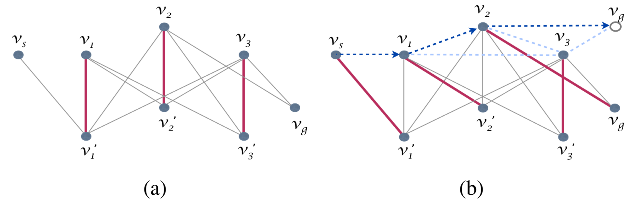

Fig. 1. Bipartite graph in the form of 3D mesh, where $V = \{v_1, v_2, v_3\}$, $V' = \{v_1', v_2', v_3'\}$. Starting nodes (i.e., AUVs' locations) are separately put in set $V_s = \{v_s\}$; similarly, the goal nodes are in set $V_g = \{v_g\}$. In this example, we have only one start node and one goal node. (a) Matched edges are in red bold, others are unmatched edges; (b) The number of matched edges increases by one after switching edge states of the augmenting path v_s—v_1'—v_1—v_2'—v_2—v_g. The projected routing path is v_s—v_1—v_2—v_g, the vertices of which are only in routing graph G. The path is illustrated by dashed arrows in the top layer.

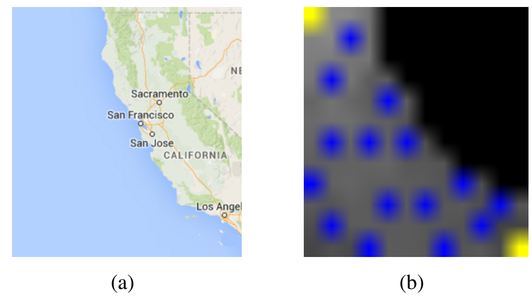

Fig. 2. (a) The sampling region along the coast of southern California. (b) The way-points (blue) with priors in the corners (yellow).

Fig. 1, as long as the value of $\tilde{w}(v_i, v_i')$ is within some feasible range [2]. Intuitively, as the weights change, the paths begin to involve more vertices and become more winding, which allows the AUVs to traverse more nearby observation points along their paths.

IV. Experimental Results

We provide initial simulation results. The robot used in simulation is an underwater glider with a simplified kinematic model. The simulation environment was constructed as a two dimensional ocean surface and we tessellated the environment into grid maps. We use salinity and ocean current data observed in the Southern California Bight region. The raw data was obtained from ROMS [5].

To control vehicles' motion, the high-level path planner described in Sect. III produces a set of MSMG paths. By setting the succeeding way-points as the short-horizon goal states, we developed a low-level Markov Decision Process based motion planner in order to generate disturbance-aware policies for the local guidance.

Fig. 2 demonstrates a set of 15 observation points that maximize mutual information in the map. In Fig. 2(b), the black regions represent lands while the gray areas denote oceans. The yellow dots and blue blobs represent prior sampling points and resultant observation points, respectively.

Following the evaluation design of prior works [1, 3], we compare our proposed method with the myopic greedy framework. We first evaluated the performance of single robot planning. Fig. 3 shows statistical comparisons between the myopic greedy approach and this proposed nonmyopic method. In the figures, the x-axis corresponds to the number

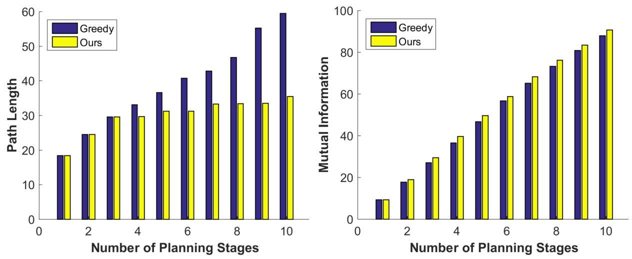

(a) 2 observations (b) 2 observations

Fig. 3. (a)(b) Comparisons in terms of path length and mutual information given two prior observations.

of planning stages, which is equivalent to the number of way-points to be generated. Fig. 3(a) shows that our approach generates paths with much shorter path lengths. We then compare the informativeness. Fig. 3(b) reveals the total mutual information gained after completing the whole paths. We can observe that our method is slighter superior to the greedy strategy. Therefore, our method performs better than the myopic greedy method in terms of the information gain, and much superior from the perspective of saving traveling time and energy.

With the navigation way-points output from the information-driven planner, the AUV follows local decisions represented by the MDP's optimal policy until it finishes the current batch of way-points. Since the GP based ocean model is temporally varying, i.e., the uncertainty of an observed point is increasing after an AUV finishes its observation at this location and moves to somewhere else, therefore, we simulate such a property by incorporating the fact that the observation uncertainty increases as time elapses. A resultant trajectory of the AUV is illustrated in Fig. 4. The colormap in these figures denote the information-gain (informativeness), which implies that the proposed method produces informative path that explores and covers most of uncertain regions. We can also see that the uncertainty of explored regions in fig. 4(a) gradually increases as the AUV reaches the later time frames in fig. 4(b) and fig. 4(c). Such a time-varying uncertainty model is exactly the motivation for the long-term monitoring.

We then tested the MSMG routing strategy for scenarios with multiple AUVs. As described in Sect. III-C, adjusting weights on the vertical edges (v_i, v_i') of the 3D matching graph will "perturb" the paths to involve differing number of way-points. Formally, for each vertex v_i, we re-initialize the weight, $\tilde{w}(v_i, v_i')$, as the least cost among all outgoing edges from v_i, and use a scaling/tuning parameter λ to scale the weight $\tilde{w}(v_i, v_i')$. Fig. 5(a)–5(c) show different results when different λ values are set. We can observe that, with a larger λ, the paths are more winding and allow the AUVs to transit more observation way-points.

V. Conclusions

In this paper, we presented an informative path planning method for long-term multi-AUV ocean monitoring. By taking into account the spatio-temporal variations of ocean phenomena, we developed an information-driven approach that computes the most informative observation way-points

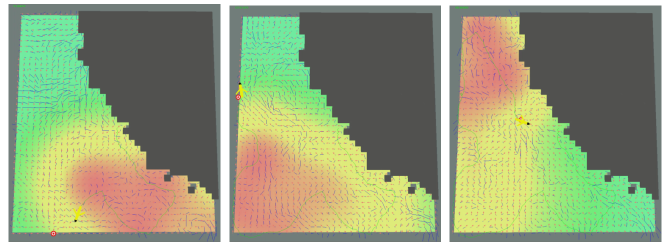

(a) (b) (c)

Fig. 4. Demonstration of a planning result with a total of 10 way-points and one robot. (a)–(c) the AUV follows a series of way-points which cover the most uncertainty regions. The colormap depicts the variance of phenomenon prediction, where a warmer color represents a smaller variance. Blue lines are ocean currents and red arrows denote the MDP action policy. Uncertainty of sampled points increases as the time evolves. Such temporally varying uncertainty model requires repetitive sampling operations in the long term.

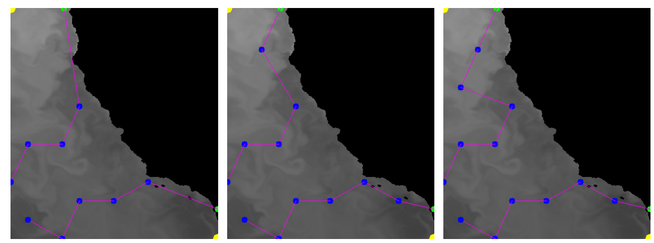

(a) (b) (c)

Fig. 5. Demonstration of planning results for 2 robots. (a)–(c) Twos paths connecting the way-points (blue) generated from our multi-robot routing planner with different λs (0, 0.5, 1, respectively). The green dot represents AUV's starting location.

that aim at minimizing the ocean model and prediction uncertainty. The sampling paths of AUVs are then formulated and solved through a matching graph based routing method, which allows the vehicles to transit the obtained informative way-points in an efficient and non-conflict way. We provide preliminary simulation results to validate the proposed method. In future, we plan to further investigate and assess the information-gain performance for different routing results tuned with differing path parameters.

References

[1] N. Cao, K. H. Low, and J. M. Dolan. Multi-robot informative path planning for active sensing of environmental phenomena: A tale of two algorithms. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, pages 7–14, 2013.

[2] L. Liu and D. A. Shell. Physically Routing Robots in a Multi-robot Network: Flexibility through a Three Dimensional Matching Graph. *ijrr*, 32(12):1475–1494, 2013.

[3] K. H. Low, J. M. Dolan, and P. Khosla. Active markov information-theoretic path planning for robotic environmental sensing. In *Proceedings of the 10th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS-11)*, pages 753–760, May 2011.

[4] R. Ouyang, K. H. Low, J. Chen, and P. Jaillet. Multi-robot active sensing of non-stationary gaussian process-based environmental phenomena. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, pages 573–580, 2014.

[5] A. F. Shchepetkin and J. C. McWilliams. The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Modelling*, 9(4):347–404, 2005.

[6] A. Singh, A. Krause, C. Guestrin, W. Kaiser, and M. Batalin. Efficient planning of informative paths for multiple robots. In *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, pages 2204–2211, 2007.