Shortest Path Exploration with Fast Marching

Nicholas R.J. Lawrance, Jen Jen Chung and Geoffrey A. Hollinger

***Abstract*— In this work we consider the problem of exploring an initially unknown region to find the minimum cost path from a fixed start position to a fixed goal position. We propose a sequential sampling algorithm to successively select observation locations that maximize the probability of improving the best path. We use Gaussian Process (GP) regression to generate a smooth estimate of the underlying cost field based on past observations. The best path based on the current estimate is then found using a Fast Marching (FM) approach. We evaluate a set of proposed observation locations using a novel FM update method and select the best location based on the likelihood of improving the current best path. Simulated results show that by sequentially selecting observation locations in this way, the resulting path estimate approaches the true optimal path.**

I. INTRODUCTION

This work is motivated by the problem of using exploration to find the lowest-cost path through an environment with an initially unknown or partially-known cost structure (for example across variable-cost terrain potentially containing obstacles). Consider the problem of a mobile robot that needs to connect a cable, such as an undersea communications cable [1], between two known positions on a map with partial prior information. Given a limited travel or sampling budget, the robot should try to find the lowest-cost path between the fixed start and goal locations for the cable. Figure 1a shows an example where the goal is to find the shortest path across the (initially unknown) cost map. This work proposes a method for evaluating the likelihood that taking an observation will reduce the best path cost. We propose a replanning method, a variant of BiFM [2], for computing fast and accurate estimates of the minimum-cost path to estimate the utility of new observations. The BiFM method was shown to significantly reduce the number of node expansions compared to restarting the search or using the dynamic-replanning E* FM method [3].

The work presented here draws on a number of existing methods for planning in continuous cost spaces. We draw primarily on the Fast Marching (FM) method proposed in [4] as a solution to approximating the cost-to-come function of continuous spaces with a discrete grid. To adapt the FM method for the minimum-cost replanning problem we suggest two primary features. Firstly, we propose the use of a bi-directional FM search. This allows computation of the minimum cost path on a new map and is shown to improve performance for incorporating map updates to find the new best global path. Secondly, we present an allocation algorithm that sequentially selects observations with the goal of continually improving the lowest-cost path estimate.

*This work was supported in part by ONR grant N00014-14-1-0509 and NASA grant NNX14AI10G.

[1]Authors are with the School of Mechanical, Industrial, and Manufacturing Engineering, Oregon State University, Corvallis, OR, 97331, USA, {nicholas.lawrance, jenjen.chung, geoff.hollinger}@oregonstate.edu

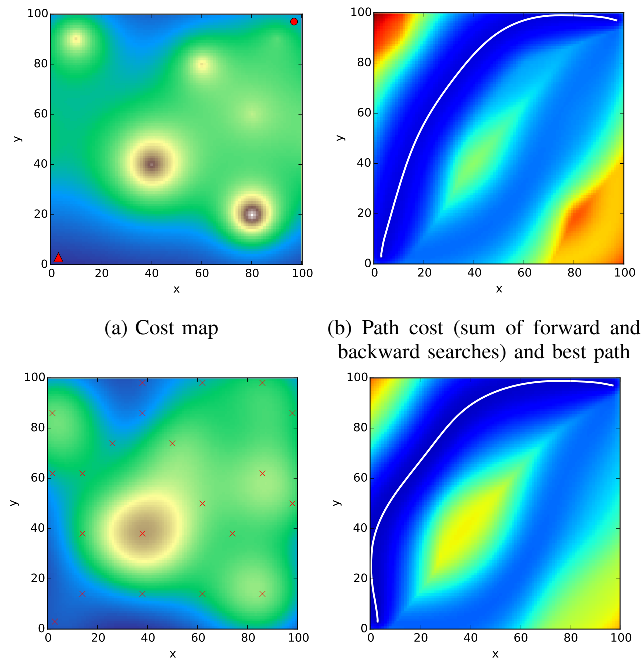

(a) Cost map

(b) Path cost (sum of forward and backward searches) and best path

(c) Estimated cost map

(d) Estimated path cost and current best path

Fig. 1: Exploration search problem. The original cost map is shown in 1a. The field is isotropic, so traversal cost only depends on location and not on the direction of travel. The goal is to find the minimum-cost path from the triangle at $(3, 3)$ to the circle at $(97, 97)$. The lower two subfigures show the estimated cost map and path costs after 20 exploration actions. The model started with a single data point at (3,3). Observations were selected sequentially from a 9×9 grid.

II. METHOD

We structure the shortest path exploration problem as a sequential process with three main steps (Algorithm 1). Firstly, when a new observation is recorded, the cost map is estimated using Gaussian Process (GP) regression. Second, the cost map is searched using FM to evaluate traversal costs over the whole map and to find the current best path. Finally, a set of proposed observation locations are compared by estimating the likelihood of an observation at a location reducing the cost of the best path. The best observation location is selected, the robot samples the cost field at that location and repeats the process.

Algorithm 1 Shortest path exploration with Fast Marching.

```
procedure FMExplorer
    while Remaining observation budget > 0 do
        Estimate cost map with GP using observation set
        Bi-directional FM search to find best path
        Evaluate and rank proposed observation locations
        Select best location, make observation
```

A. Cost map estimation

We propose using GP regression for modeling the cost field (without obstacles) based on past observations. The GP model was selected due to the ease of handling noisy observations and because it generates smooth approximations of the cost field as well as providing variance estimates. A detailed discussion of GP regression is beyond the scope of this paper and interested readers are directed to [5]. In this work we only consider temporally static cost fields with noisy observations. We selected a stationary radial basis function, the commonly-used squared exponential covariance, in the GP regression. We selected hyperparameters for the length and noise scales at the start of the problem and did not retrain them during exploration. In ongoing work we are exploring the effect of alternative covariance functions. In particular, we are interested in functions with compact support [5] to provide a theoretical basis for our later assumptions that observations have a distance-limited effect on the cost map.

B. FM search and update

The proposed algorithm fully completes FM searches from both the start and goal nodes each time the cost map is updated by the GP. Using the results of those searches, a cost map can be constructed showing the shortest path from the start node to the goal node via any other node in the map (Fig 1b). Now, consider an update observation that changed some costs in the map in a single continuous region. Every point on the boundary of the updated region already has the cost of the shortest path to both the goal and start nodes from the existing map. If the changed nodes did not affect the calculation of the path cost (the changed nodes are 'downwind') then that path remains valid. Thus, to find the new shortest path it suffices to only update those nodes directly affected by the change. Further, since only the best path is needed, and nodes are expanded in strict ordering with lowest cost first, we only need to update one direction of the search (start to goal) until we attempt to expand a node for which an existing valid cost-to-goal still exists.

C. Observation selection

Given a probabilistic model of the cost map, we would like to find the observation location that has the highest likelihood of reducing the cost of the best path. If we consider this as a type of expected value at an observation location, then we would need to calculate the integral of the change in path cost multiplied by the likelihood over all possible observations. However, while the GP can provide the likelihood of an observation through the posterior probability density function (a Gaussian) at each location, we cannot easily formulate an analytic solution to the best path cost. We present a method to estimate this value using a finite number of proposed observations and use the FM model to calculate the new best path cost with each proposed observation.

In this paper we approximate the value function using two proposed observations at each observation location, namely the observations at one standard deviation above and below the mean estimate from the GP regression. We approximate the value of a location as the sum of the resulting best path cost with these two observations. However, to fully determine how an observation would change the cost map we would need to perform the GP regression again including the new observation and recompute the FM search, a computationally expensive operation. To avoid this, we assume that each observation will only modify the cost function within a limited distance of the observation location. To model this effect we use a polynomial compact-support covariance function [5] as a local approximation to the squared exponential used in the GP, and assume that the new cost map is the existing cost map plus the observation function. This limits the effect of the observation to a circular region around the observation location, inside which we can update the FM search.

Further, we can quickly evaluate whether an update will result in a new shortest path. The current best path cost can act as a threshold to terminate the search if at any point we pop an element from the forward search queue where the cost-to-come (from the start) plus the minimum cost-to-goal of any boundary cell is greater than the current best path cost. The attached video shows a set of updates being evaluated. Note that many are quickly rejected since even if the update reduced the traversal cost of the entire region to zero (we do not consider negative costs) then the path would still be more costly than the current best path.

III. RESULTS

We show a simulated exploration problem with a continuous underlying cost field (shown in Fig. 1a). The goal is to find the shortest path through the field with limited observations. The algorithm starts with a single observation at $(3, 3)$. Observations are drawn from a normal distribution with standard deviation 1.0 and mean of the true underlying field at that location. After each observation, the algorithm must estimate the field using the GP, complete the full bi-directional search on the updated map and then select from the proposed observations. Proposed observations are drawn from a 9×9 grid of points evenly distributed over the search space. Figures 1c and 1d show the estimated cost map, path cost and current best path after 20 exploration observations.

Calculating the updated map with the GP and completing the initial complete bi-directional FM search took approximately 1.2s. Calculating the updated best path cost for the 162 possible observations (two at each location in the grid) took approximately 2.2s. Proposed observations are evaluated sequentially in the the current implementation but the algorithm is inherently embarrassingly parallel across proposed observations after the initial search. The code was written in Python and executed on a 3.4GHz Intel processor.

REFERENCES

[1] C. Elmendorf and B. C. Heezen, "Oceanographic information for engineering submarine cable systems," *Bell System Technical Journal, The*, vol. 36, no. 5, pp. 1035–1094, Sept 1957.

[2] N. R. J. Lawrance, J. J. Chung, J. Caley, and G. A. Hollinger, "Efficient replanning with bi-directional fast marching," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2015, submitted.

[3] R. Philippsen and R. Siegwart, "An interpolated dynamic navigation function," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*. IEEE, 2005, pp. 3782–3789.

[4] J. A. Sethian, "A fast marching level set method for monotonically advancing fronts," *Proceedings of the National Academy of Sciences*, vol. 93, no. 4, pp. 1591–1595, 1996.

[5] C. E. Rasmussen and C. K. Williams, *Gaussian Processes for Machine Learning*, ser. Adaptive computation and machine learning. Cambridge, Massachusetts: The MIT Press, 2006.