DOI 10.31256/UKRAS19.11

An Information Theoretic Approach to Path Planning for Frontier Exploration

Callum Rhodes, Cunjia Liu and Wen-Hua Chen

***Abstract*— Efficient path planning for autonomous robots to explore unknown areas is a critical area of research due to the requirements on mission times and problems that dictate as quick a solution as possible e.g. search and rescue scenarios. It also proves to be a difficult problem to solve due to its inherent NP-hard nature, which requires that an optimal (albeit not necessarily perfect) path is defined based on a set of defined principles. This gives rise to a wide variety of logical solutions. This paper proposes an information theoretic addition to the well-established Frontier Exploration in order to build a 2D spatial map of an area of interest as intelligently as possible. The informative method is then compared with a greedy approach toward information gain, as well as the traditional 'nearest frontier' approach to frontier exploration. The proposed method is shown to outperform other methods in terms of the total number of actions required to resolve the map, as well as being consistently the quickest method of reducing map entropy throughout the mapping procedure. We also discuss how, by exploiting an information theoretic framework, other quantities of interest can be mapped efficiently alongside a spatial map.**

I. Introduction

Autonomous exploration is an increasingly useful area of research as industry pushes for the autonomy of tasks perceived to be trivial or too dangerous for human undertaking. Environmental surveillance after dangerous incidents, such as that seen in the Fukushima nuclear power plant explosion, are a prime use case for autonomous exploration as the environment has become uncertain physically and is too dangerous for human intervention. In this scenario, the spatial mapping of the environment must be attained in cases where the localisation within the environment is unknown. Therefore, mapping and localisation must be performed simultaneously, commonly referred to as SLAM (Simultaneous Localisation and Mapping). However, trajectory planning must also be considered to reduce the uncertainty of the map, whether the map produced is spatial in nature or of another form (e.g. gas or radiation concentration mapping). An efficient trajectory is desirable as limitations may be imposed on the mission time or the number of actions capable of being performed by the robot. This paper seeks to explore the planning of the most efficient routes in typical spatial mapping missions.

For a robot to successfully and *efficiently* map an area autonomously, it must complete the *active SLAM* problem (coined by Leung et al. [1]). The SLAM portion of active SLAM is a vast area of research and there are many approaches to the issue which transcend the scope of work for this paper. The SLAM review by Cadena et al. [2] concludes that SLAM is *not* solved and therefore trying to address the path planning problem independently allows for evolving solutions to SLAM to not supersede work on the path planner. Therefore, we assume that the Localisation and Mapping are accurately handled by an appropriate SLAM algorithm and that the path planning portion can be dealt with independently of this (acknowledging that the performance of the path planner is entirely dependent on the SLAM algorithm). This assumption requires that the robot does not need to perform specific loop closure activities to provide high fidelity mapping. This is, however, dependent on the SLAM package used and the environment it is operating in, as shown in [3]. The independence assumption becomes insignificant as the performance of various SLAM packages improves.

Frontier Exploration, proposed by Yamauchi [4], has been chosen as the goal setting technique due to its simplistic nature and successful application within the exploratory mapping domain. Frontier Exploration functions by defining naive targets for a robot to navigate by stating that any unknown areas of a map that border free space are points of interest. Traditionally, Frontier Exploration simply chooses the closest frontier of a minimum size and plans an obstacle free trajectory to resolve said frontier. The process then eventually terminates when all frontiers have been resolved. While effective, this approach does not consider the possible information gain at the target and therefore may not choose the most efficient route. By considering information gain at points of interest, the path planner will make informed decisions on where to navigate and therefore provide an overall mapping trajectory that is more efficient than that of an uninformed technique.

Recent sampling-based path planning methods that have exploited information gain have not used a goal setting method such as frontier exploration. They have instead used a horizon based approach wherein the prediction is either a set number of steps ahead of the current state [5] (finite horizon) or is subject to a total budget constraint [6] (infinite horizon). Using Frontier Exploration allows the horizon consideration to be ignored and narrows the focus of the path planner to specific trajectories. This reduces the total number of path considerations, reducing the computational load at the expense of planned path generality.

II. Information Theory

A. Definition of Information Gain

Information Gain when trying to resolve a variable of interest can be defined as how variable uncertainty is reduced via an action or observation. Uncertainty is mathematically characterised by the term *Entropy*. For an occupancy grid, where the occupation probability of a cell P_{cell} is defined:

$$0 < P_{cell} < 1 \tag{1}$$

$$P_{cell} = 0 \quad \text{if free} \tag{2}$$

$$P_{cell} = 1 \quad \text{if occupied} \tag{3}$$

The entropy of a single cell H_{cell} is given by:

$$H_{cell} = -(P_{cell} \log(P_{cell}) + (1 - P_{cell}) \log(1 - P_{cell})) \tag{4}$$

Due to the assumption that each cell's occupation probability is independent of each other cell [7], the total map entropy H_{map} can be defined as a simple summation of the entropy of the total number of cells, N, such that:

$$H_{map} = \sum_{i=1}^{N} H_{cell}(i) \tag{5}$$

The change in entropy between two states can be considered as the information gain over a step. When comparing multiple states to the same reference state, this provides a suitable measure of the information gain. However, the value returned is predicated on the reference state entropy and therefore cannot be used to compare the change in entropy over time. The relative entropy between consecutive distributions shows how different the two distributions are and therefore can be thought of as the relative information gain. The Kullback-Liebler divergence [8] is a method of calculating the relative information gain between two distributions. For calculating the information gain between the current map state (P_1) and a comparison map state (P_2) in this manner, the

$$D_{KL}(P_1||P_2) = -\sum_{i=1}^{N} P_1(i) \log\left(\frac{P_2(i)}{P_1(i)}\right) \tag{6}$$

following equation is used:

where i denotes the i-th cell, N is the total number of the cells and both map occupancy probability matrices are discrete probability distributions of the same probability space.

B. Predicting Information Gain

In order to predict the information gain at a point of interest using an occupancy grid, an inverse sensor model [9] can be employed which takes key parameters of the sensor that is being used by the robot (in this case a RPLIDAR A2 lidar scanner). The inverse sensor model casts rays in an occupancy grid at a given location as if the scanner were to be deployed at said location. It then returns any obstacle hits by the rays, based on the current occupancy status of the cells within range. If no obstacles are hit by the beam, then a blank distance reading is returned. The occupancy status of cells within range are then updated based of the parameters of the inverse model. In this case a hit cell is updated with probability 0.7 and any free cells along the path with 0.3.

If unreturned rays are ignored, then this gives the lower bound of the possible information gain at a location and is often unrealistic to the actual gain when the sensor is at the location of interest. If unreturned rays are treated as not hitting any obstacles and cells along the ray updated with the free probability until the maximum range of the sensor, then this gives the maximum possible information gain i.e. upper bound. Using the Information gain upper bound, whilst optimistic, is a better predictor of the information gain at a point and therefore is used for future map state estimation.

III. PATH PLANNING

A. Frontier Exploration

A general issue in autonomous exploration is that of setting the next goal point for the robot to navigate towards. Frontier exploration sets bounds between free space and unknown space as targets using an edge detection method. In this study, the Canny edge detection method [10] is used to abstract frontiers from the occupancy grid. The midpoint of a frontier is generally used as a goal point [4] but this possibly ignores several viewpoints along the frontier which could yield more information. This becomes increasingly important with longer frontiers. Hence, an equal distribution of viewpoints along the frontier are selected at a set distance. This gives rise to a tunable parameter that is specific to the scale and resolution of the environment being mapped.

The minimum length of frontier is also specific to the environment and is required in order to avoid needlessly attempting to resolve frontiers which are insignificant or erroneous due to noise in the scanner or mapping procedure.

B. Path Selection

The path selection problem rises once all possible obstacle free trajectories to all frontier points have been established. Traditional frontier exploration simply states that the closest frontier point should be chosen however this does not take into account any information about the future state of the map at that point.

By employing the inverse sensor model at each of the frontier points and then using the Kullback-Liebler divergence between the current map state and the predicted map state, the information gain for each path is established. By selecting only the end point of the path to employ the sensor model, rather than predicting across the entire trajectory, computational time is reduced dramatically. Further, since the robot is moving across known free space to navigate towards the frontier, the information gain during traversal to the frontier is insignificant compared to the gain at the frontier itself.

It can be naively stated that the path which exhibits the most information gain should be selected, the Next Best View (NBV). This however does not take into account the cost involved in travelling to a frontier, which is important for scenarios which either have a limited endurance e.g. UAVs, or in a time-limited scenario e.g. search and rescue operations. Therefore, a tradeoff between the path length and the information gain must be established. Whilst this can be scaled depending on which variable is more desirable, a reward function that simply calculates the path with the largest information gain per metre travelled is used for testing purposes (NBV/m).

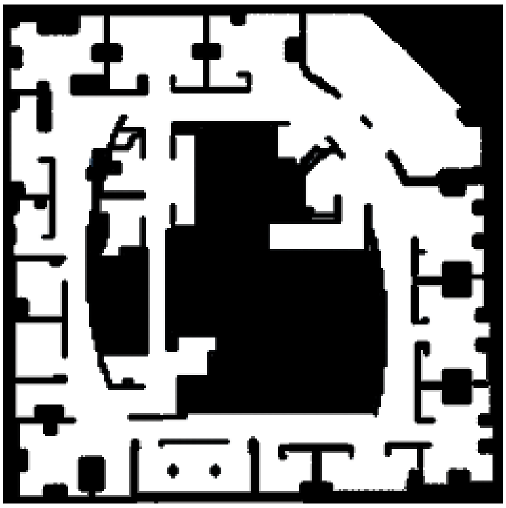

Figure 2. Intel research lab, Seattle.

$$T^* = \underset{T}{\operatorname{argmax}} \left(\frac{I_{UB}(T)}{T_{length}} \right) \qquad s.t. \quad T_{length} > r_{min} \qquad (7)$$

where T are the possible trajectories, $I_{UB}(T)$ is the upper bound information gain of a trajectory and r_{min} is a minimum radius around the robot as to avoid the robot getting stuck in local minima.

IV. Simulation Study

A. Simulation Setup

For simulating the performance of the information theoretic addition, the *Mobile Robotics Simulation Toolbox* available for *Matlab* is used with a LIDAR equipped robot. Perfect SLAM is assumed for simulation (mapping with known poses) and noise is added to the LIDAR scan readings to allow for anomalous readings in the map. The scenario used is the Intel research lab map which is often used for testing similar mapping and SLAM algorithms due to its high amount of clutter, making the mapping task more realistic and difficult than other simpler scenarios [11].

In order to test the algorithm within a full autonomous system, a simple probabilistic roadmap (PRM) is employed with a path following function that takes advantage of a differential drive system such as that seen on the Turtlebot.

For performance comparison, the 3 path selection strategies previously described (closest frontier, NBV \& NBV/m) are run under the same conditions. The rate of map entropy reduction as well as the final map entropy after 300s are logged and used as metrics to analyse the performance of each strategy. It should be noted that the simulation is set so that the robot is constantly performing actions and therefore, limiting the simulation time to 300s means that all 3 strategies will have performed the same number of movement actions.

B. Simulation Results

Figure 2 shows the Entropy of the map against time for all 3 strategies. The initial information gain for all 3 strategies is comparable as, regardless of the path chosen, the map is in a high state of uncertainty and therefore any action resolves the map significantly. As the robot moves through the environment, the difference in strategies becomes apparent as

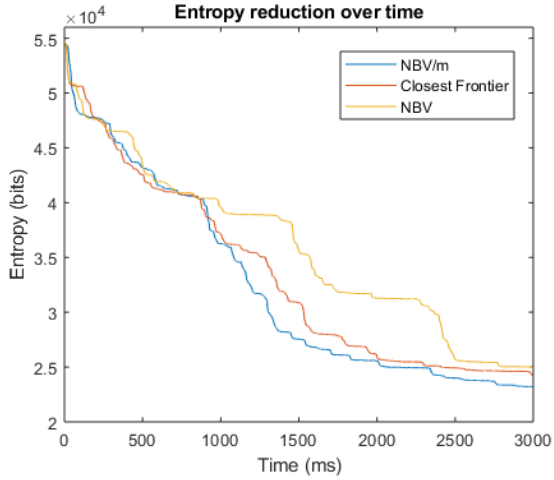

Figure 1. Entropy reduction over time for each path selection strategy.

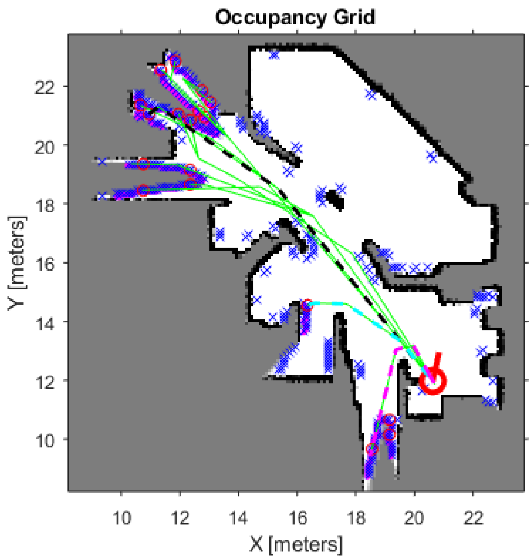

Figure 3. Comparison of path selection by NBV (black dash), NBV/m (magenta dash) and closest frontier (cyan dash) strategies. Frontier cells are represented by a blue X, frontier targets by a red O and all possible trajectories by green lines.

displayed in Figure 3. This snapshot is a prime case of the algorithm choosing the intuitively correct path. The proposed algorithm successfully ignores the closest frontier, as there is little to no information to be gained at this location and it is only deemed a frontier due to the noise in the LIDAR scanner not resolving the wall properly. It also does not greedily go to the frontier with the most information which would require a large traversal of the map. Instead, it chooses the frontier that will allow the robot to continue down the corridor thus optimising the future information gain the robot can achieve without prior knowledge of the rest of the map.

TABLE I
PERFORMANCE COMPARISON BETWEEN THE 3 STRATEGIES.

Strategy	Closest Frontier	NBV	NBV/m
Final Entropy (kbits)	24.31	24.94	23.22
Calls to function	50	27	34
Time to complete (s)	200	249	182

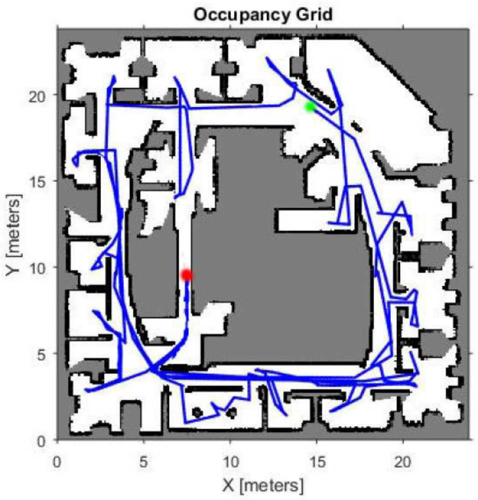

Figure 4. Completed occupancy map with total trajectory. Green dot indicates start point and red dot indicates end point.

As more of the map is resolved, Figure 2 shows that the greedy NBV approach becomes increasingly inefficient at reducing map entropy due to long traversals over areas of low uncertainty. However, the NBV/m and closest frontier approach remain relatively similar in their ability to reduce entropy, performing better than the NBV method due to their more efficient short path lengths. Where the performance of NBV/m becomes apparent is in the later stages of the mapping process where large areas of uncertainty are less common. In this instance, whereas the closest frontier approach may attempt to resolve very small frontiers (which add little information) purely because they are close by, the NBV/m algorithm intelligently deduces areas where uncertainty can still be significantly reduced despite the larger path cost.

Final mapping performance is seen when the relative entropy is reduced to a significantly low value, that the map can be deemed complete. For the intel map this was judged to have occurred when:

$$argmax\left(D_{KL}(T)\right) \leq 400bits \tag{8}$$

Table 1 shows the numerical performance of each strategy. Analysing the time taken to complete, NBV/m shows an improvement of 27% over the greedy approach and 9% improvement over the traditional approach. The final entropy value of the NBV/m approach is also lower, which accentuates the fact that it performs significantly better than the traditional approach in the latter stages of the mapping procedure.

Figure 4 shows the final occupancy grid created by the simulated robot in the NBV/m case and it is clear that all major structures of the environment have been adequately resolved, thus showing that the stopping criterion of 400 bits is suitable for the Intel Lab environment. The overall trajectory consisted of one anticlockwise loop of the main corridor followed by a second half loop to suitably resolve some of the smaller rooms that weren't resolved adequately during the first loop.

V. CONCLUSIONS AND DISCUSSION

The addition of an intelligent information theoretic path planner adds significant benefits to the traditional frontier exploration approach to autonomous exploration. The use of a greedy approach is proven to be disadvantageous compared to the traditional approach and provides evidence that efforts must be made so that the reward function takes into account the path cost effectively. In this regard, whilst the reward function presented provides better performance for a simulated SLAM mission of an office environment, it may not be optimal for all mapping scenarios and should be investigated further.

Using uncertainty as the primary consideration for path planning also adds flexibility to its application, as uncertainty is a universal metric when mapping any unknown quantity (such as temperature, gas concentration etc.). This means that not only spatial mapping can be planned efficiently using this technique but theoretically, given an appropriate sensor prediction model, any quantity that requires the reduction in its uncertainty can be used. Further to this, if multiple quantities are in the same information framework, they can be mapped simultaneously with weightings on whichever quantity is of greater interest to the mission.

REFERENCES

[1] C. Leung, S. Huang, and G. Dissanayake, “Active SLAM using model predictive control and attractor based exploration,” IEEE International Conference on Intelligent Robots and Systems (IROS), pp. 5026–5031, 2006.
[2] C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” IEEE Transactions on Robotics, vol. 32, no. 6, pp. 1309–1332, 2016.
[3] W. Hess, D. Kohler, H. Rapp, and D. Andor, “Real-time loop closure in 2D LIDAR SLAM,” Proceedings - IEEE International Conference on Robotics and Automation, vol. 2016-June, pp. 1271–1278, 2016.
[4] B. Yamauchi, “A Frontier Based Approach for Autonomous Exploration,” Proceedings of IEEE CIRA, pp. 146–151, 1997.
[5] S. S. Geoffrey Hollinger, “Proofs and Experiments in Scalable, Near-Optimal Search by Multiple Robots,” Proceedings of Robotics: Science and Systems IV, pp. 206–213, 2008.
[6] G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” International Journal of Robotics Research, vol. 33, no. 9, pp. 1271–1287, 2014.
[7] S. Thrun, “Learning occupancy grids with forward models,” Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No.01CH37180), vol. 3, pp. 1676–1681, 2015.
[8] S. Kullback and R. A. Leibler, “On information and sufficiency,” The annals of mathematical statistics, vol. 22, no. 1, pp. 79–86, 1951.
[9] A. Elfes, “Sonar-Based Real-World Mapping and Navigation,” IEEE Journal on Robotics and Automation, vol. 3, no. 3, pp. 249–265, 1987.
[10] J. Canny, “A Computational Approach To Edge Detection,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 8, no. 6, pp. 679–714, 1986.
[11] S. Owald, M. Bennewitz, W. Burgard, and C. Stachniss, “Speeding-up robot exploration by exploiting background information,” IEEE Robotics and Automation Letters, vol. 1, no. 2, pp. 716–723, July 2016.