Frontier Based Exploration with Task Cancellation

P.G.C.N. Senarathne and Danwei Wang

***Abstract*— Classical frontier based exploration strategies operate by iteratively selecting the next best sensing location myopically and moving to the specified location, until the entire environment is explored. And it does not consider the new information added to the map through continuous observations by the robot along the way to a selected location. This can sometimes lead to redundant traversal by the robot, such as traveling towards a dead-end when the nearby area is already mapped. In this work, we augment the classical frontier based exploration strategy to include a probabilistic decision step that decides whether further motion on the planned path is desirable or not. If the motion is not desirable, it is interrupted and a new sensing location is selected as the next sensing task. Experiments were conducted using a Pioneer 3AT robot to explore an indoor environment and is demonstrated that the proposed method on average is capable of exploring environments more efficiently.**

I. INTRODUCTION

Autonomous exploration of unknown environments is one of the most important tasks carried out during robotic missions such as environmental mapping [1], search and rescue/find [2], [3]. For a robot to autonomously explore an unknown environment, it must sequentially sense the environment at new sensing locations until complete sensor coverage of the environment is achieved. At a given instance during an exploration mission, there are multiple candidate locations where the robot could perceive the environment to expand the existing mapped area. Once the robot senses the environment through one of these caldidate locations, a new set of candidate locations appear in the map. This explosion of the candidate sensing space with each sensing action, coupled with the high uncertainties in the information gain predictions for each of these sensing locations, makes finding an optimal sequence of sensing locations for complete exploration of the unknown environment intractable. Therefore, all the exploration strategies are reduced to finding the next best sensing location as an intermediate target and moving to that target as the next task and repeating this process until complete mapping of the environment is achieved. The most popular heuristic used to generate the set of candidate locations, to find the next best sensing location (i.e. intermediate target), is to extract cells in the boundary between the mapped free and unmapped cells from an Occupancy Grid Map representation for the environment [4]. These boundary cells are called *frontier* cells and the exploration strategies based on this heuristic are called Frontier Based Exploration.

In almost all the frontier based exploration strategies, the robot is left to travel to the selected target (i.e. completion of current sensing task), before the next decision is made. While this method is simple and reduces the exploration problem to a one finding the best intermediate target to expand the map in discrete steps, it does not consider the changes to the frequently evolving frontiers, which could invalidate the usually over-estimated information gain for the intermediate target. This happens when robots select frontier cells, that are very close to obstacles which are yet to be revealed in the map, as their next intermediate targets. As a robot moves towards such a target, continuous sensing reveals the obstacle, which encloses the explored free space, fully or partially, and stops the expansion of most of the frontier cells in the neighborhood, an extreme situation of which is depicted in fig. 1. This disappearance of frontiers in the current target's neighborhood results in the further movement of the robot towards the target, largely redundant as large portion of the map does not get expanded.

Both authors are with EXQUISITUS, Center for E-City, School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore senarathne@ntu.edu.sg, edwwang@ntu.edu.sg

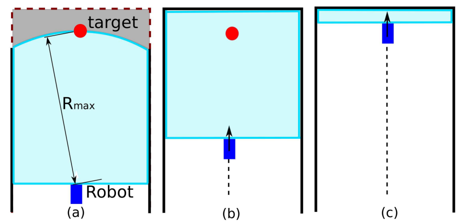

Fig. 1. (a) Illustrates a situation where a frontier cell that is close to a yet to be revealed obstacle boundary, marked in dashed lines, is selected as the next intermediate target. The white area is the already explored free space while solid black lines are the already discovered obstacle boundaries. Gray represents unmapped area. R_{max} is robot's maximum sensing range and light-blue area illustrate robot's current sensing field of view. (b) As the robot moves towards the target (as shown by the arrow), the entire area becomes fully explored (c) The robot continues to move towards the target in traditional exploration strategies, redundantly

Therefore it is important to check if the map continues to get expanded in the direction of the robot's final desired heading. This expansion is directly reflected by the distribution of the frontier cells near the current target as the robot moves towards the target. Existence of frontiers in the direction of the robot's desired heading indicate that the map gets expanded in that direction. Frontiers in other directions indicate the options available for the robot for exploration in those directions. Depending on the distribution of the desirability of these two types of frontiers, decisions can be made to either continue the current motion or to cancel it and generate a new exploratory task in another direction. The work presented in this paper augments the classical frontier based exploration strategy to include this decision step in order to conduct more efficient exploration missions.

The article begins by providing a summary of various works on autonomous exploration strategies and a description on the baseline frontier based exploration strategy. The next section describes the approach used to arrive at the task cancellation decision. It includes a probabilistic formulation for the decision event and how the parameters decide the aggressiveness or the conservativeness of the exploration process. It also details the generation of the frontier selection probability distribution used for the decision event. Details of the experiments, the results and analysis are provided in the subsequent sections.

II. RELATED WORK

Robot exploration strategies are mainly driven by the concept of iteratively selecting the next best sensing location in the environment as the intermediate target in order to expand the map. Frontier based exploration [4] selects the next best sensing location from a collection of candidate target locations generated from the boundary between the mapped-free and unknown grid cells (i.e. frontiers) in the occupancy grid map [5]. Different criteria used for the selection of the next intermediate target has resulted in various extensions of the basic frontier based exploration strategy. Selecting the closest frontier cell, balancing the information gain of frontier cells with travel cost [6], [7], [8], use of *hysteresis value* to restrict robot from frequently switching exploration tasks [9] are some of them. While the popular heuristic is to use frontier cells as candidate targets for exploration, random generation of candidate target locations have also been proposed. The works in [10] discuss the generation of candidate target locations randomly while [11] proposes the biasing of the random target generation towards the frontier boundaries.

Several exploration algorithms based on topological map representations are also proposed in the literature. These include the works of Kuipers et. al, [12], Choset et. al, [13] and Ge et. al, [14]. However, frontier based exploration strategies have become prevalent due to the ease of generation and management of occupancy grid maps compared to topological maps hence are the focus of this article. The baseline strategy used for comparison of the proposed method is a frontier based exploration strategy. The utility $U(\lambda)$ of the frontier cell λ is generated as $U(\lambda) = \alpha\mathcal{I}(\lambda) - \beta\mathcal{C}(\lambda)$ where $\mathcal{I}(\lambda)$ is the estimation of the information gain and $\mathcal{C}(\lambda)$ is the estimated travel cost to cell λ. α and β are two parameters that can be varied to decide the relative importance of the information gain and cost components. The frontier cell with the highest utility is selected as the next intermediate target. Calculation of the information gain can be done using entropy/mutual information based methods [8], [7] or by counting nearby unexplored cells, generated by thresholding the occupancy probabilities [9]. Since the second method is simpler to implement and is not considered inferior [15], it is used to calculate the information gain. The travel cost is estimated as the distance of the planned path to the frontier cell.

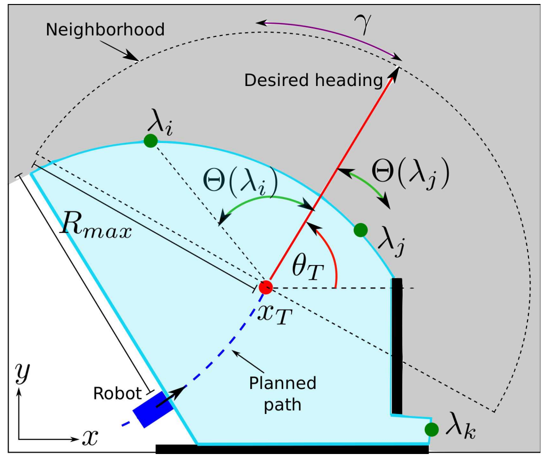

Fig. 2. x_T is the current intermediate target, θ_T is the desired heading angle. Both λ_i, λ_j are example depictions of two frontier cells neighboring x_T. Note that frontier cell λ_k is outside the neighborhood of x_T, so not considered for the cancellation decision. The neighborhood is depicted by the dashed semi-circle with a radius of R_{max}. The light-blue shaded area represents robot's current sensor field of view.

Improving the efficiency of frontier based exploration using updated frontier information has previously been mentioned by Keidar et. al, [16]. However, their work focuses only on efficiently generating frontiers in high frequency to support such improved strategies. Holz et. al, [17] discusses the drawback of continuing to move towards the intermediate target without considering the continuously updated frontier information. A repetitive rechecking approach is proposed where, the assigned intermediate target is checked for being a valid frontier cell in order to reduce redundant exploratory motion. Our work extends this approach to a more general formulation that considers the utility around the target and to decide when to cancel the current motion.

III. CANCELLING EXPLORATORY MOTION

Let us define the current intermediate target of the robot as $x_T \in \mathcal{R}^2$ and the final heading of the robot's planned motion at x_T as θ_T. Consider that the map is updated at discrete steps during the planned motion to x_T. Let $\Theta : \mathcal{R}^2 \rightarrow [0, \frac{\pi}{2}]$ be the random variable describing the absolute difference in angle between robot's desired heading θ_T and $\overrightarrow{x_T\lambda}$ where λ is any neighboring frontier cell of x_T as illustrated in fig. 2. The set of neighboring frontier cells of target x_T is denoted by $\mathcal{F}(x_T)$. At each map update step k, a probability density function $f^k_\Theta : [0, \frac{\pi}{2}] \rightarrow [0, 1]$ can be defined. This distribution describes the probability of selecting a frontier cell from $\mathcal{F}(x_T)$ in a specified absolute angle difference with θ_T, as the next intermediate target. Suppose γ to be the angle tolerance used to consider frontier cells as belonging to robot's final heading direction. Then, the robot's motion towards its current target x_T is cancelled if $P(\Theta \leq \gamma) < p_T$. The cumulative probability on the left hand side of the inequality provides a measure of desirability of the frontier cells in the 'direction' (based on angle tolerance γ) of the robot's final heading for the robot. If this desirability is less than a certain probability threshold p_T, the robot's motion

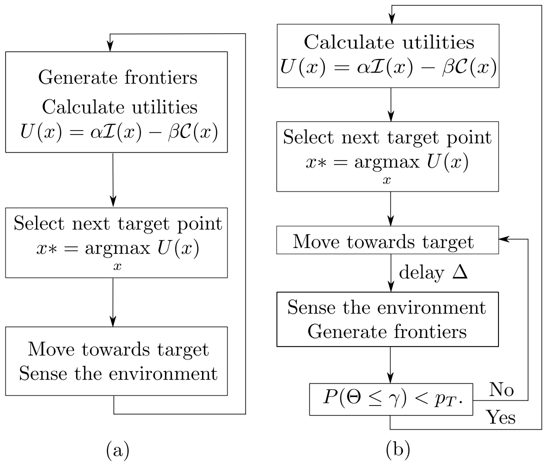

Fig. 3. (a) Major steps in the base-line exploration strategy (b) Major steps in the proposed exploration strategy. Frontiers are generated at a certain frequency and the motion is preempted to check for redundant motion. And the motion continues if the motion towards the target is still desirable. If not desirable, a new target is selected and the process continues.

towards the current target is cancelled as illustrated by the steps in fig. 3. The value of p_T can be changed according to the level of cancellation desired. A more aggressive form of exploration can be activated by setting p_T to a higher value and a more conservative form of exploration can be activated by a lower p_T value. Setting p_T to zero makes the inequality false for all scenarios and will not cancel the robot motion, hence will make the robot behave identically to the baseline strategy. $P(\Theta \leq \gamma)$ is calculated as follows.

$$\begin{aligned} P(\Theta \leq \gamma) &= \int_0^{\gamma} f_{\Theta}^k(\theta)\,\mathrm{d}\theta \\ &= \int_0^{\gamma} \sum_{\lambda \in \mathcal{X}_{\theta}^T} f_X^k(\lambda)\,\mathrm{d}\theta \end{aligned} \tag{1}$$

Here, the set $\mathcal{X}_{\theta}^T$ is defined as $\{\lambda \in \mathcal{F}(x_T) \text{ s.t. } \Theta(\lambda) = \theta\}$ and contains all the neighboring frontier cells of x_T with an absolute angle difference to θ_T equal to θ. $f_X^k : \mathcal{R}^2 \to [0, 1]$ is the probability mass function (p.m.f.) that provides a measure of desirability of selecting neighboring frontier cell $\lambda \in \mathcal{F}(x_T)$ as an alternative target.

IV. ESTIMATING NEIGHBORHOOD FRONTIER CELL SELECTION PROBABILITY

The probability of selecting a neighboring frontier cell (i.e. desirability of frontier cell as an alternative target) is calculated based on their individual utilities. Hence the p.m.f. can be defined as,

$$f_X^k(\lambda) = \frac{U(\lambda)}{\sum_{\lambda_i \in \mathcal{F}(x_T)} U(\lambda_i)} \tag{2}$$

However, it is reasonable to assume that the effect of the travel cost is negligible for the set of neighboring frontiers of x_T, $\mathcal{F}(x_T)$, as the variation among these travel costs is minimal. Hence Eq. 2 can be approximated using the information gain values $\mathcal{I}(\lambda)$ as follows.

$$f_X^k(\lambda) = \frac{\mathcal{I}(\lambda)}{\sum_{\lambda_i \in \mathcal{F}^T(x_T)} \mathcal{I}(\lambda_i)} \tag{3}$$

While crude calculation of information gain for each frontier cell, $\mathcal{I}(\lambda)$, can provide the necessary p.m.f values, it is desirable to generate these p.m.f. values with less computational burden as they need to be computed at each frontier update step. Hence the apparent correlation of information gain values among nearby frontier cells is used to quickly arrive at estimates of actual information gain values. Frontier cells generally belong to clusters of cells representing map boundary contours that share a common unknown area that is used to estimate the information gain. Corners of these contours are adjacent to already mapped area restricting the information gain, while the mid points of these contours are generally the farthest points away from mapped area allowing a higher information gain. This heuristic knowledge of information gain progression along a frontier cell contour is used to define a function that approximates the information gain for each frontier cell.

Consider a frontier contour c, then the information gain of cell $\lambda \in c$ is approximated as $\mathcal{I}(\lambda) = \mathcal{I}_c \psi_c(\|\lambda - \lambda_{\mu,c}\|)$ where $\mathcal{I}_c$ corresponds to the contour wide common information gain term. The function $\psi_c(d)$ approximates the fraction of $\mathcal{I}_c$ the robot gains by visiting a frontier cell d distance away from it's associated contour's center $\lambda_{\mu,c}$. R_c is the maximum distance to a corner of the frontier contour from $\lambda_{\mu,c}$.

$$\psi_c(d) = \begin{cases} 1 & \text{if } d + R_{max} \leq R_c \\ \frac{1}{\sqrt{2\pi}R_{max}} e^{\frac{-(d+R_{max}-R_c)^2}{R_{max}^2}} & \text{else} \end{cases} \tag{4}$$

The function value is kept at 1, indicating the highest possible information gain, when the frontier cell is more than R_{max} distance away from the frontier contour's corners, towards the center $\lambda_{\mu,c}$, thus not restricting the maximum sensor range. When the frontier cell's position makes the sensor range goes beyond the contour's corners, it restricts the sensing of the robot and the fraction of the information gained by the robot is approximated to decline according to a normal function.

V. EXPERIMENTAL RESULTS

Experiments are conducted using a Pioneer 3 AT robot equipped with a Hokuyo LRF for sensing. The exploration strategies were implemented using the navigation software layer provided by ROS [18]. The occupancy grid map update/access frequency is set to 1Hz. In all experiments, α and β parameter values are set to 0.8 and 0.3 respectively. Changes to these two values do not affect the performance of the task cancellation decision as they are used only for selection of the intermediate target, hence are kept constant throughout the experiments. Two types of experiments were conducted to compare the performance of the proposed strategy. In the first type, effect of cancellation on a single sensing task is evaluated while on the second type of experiments,

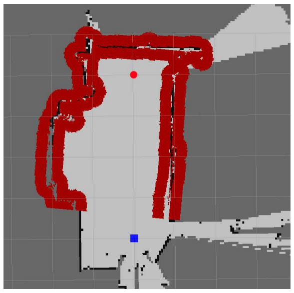

(a)

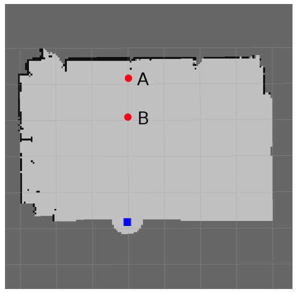

(b)

Fig. 4. (a) Map of the narrow passage. The blue box is the starting position of the robot, red circle is the target near the dead end. The red band circling the mapped space is the inflated obstacle grid cells. (b) Map of the room type environment. Two targets A and B are used

the effect of task cancellation was evaluated for complete exploration missions. Travel distance of the robot is used for comparisons between the two strategies.

A. Analyzing the effect of parameters on different environments

In the first type of experiments, the cancellation of exploration tasks is compared with traditional exploration tasks to better understand the effects of the two parameters γ and p_T. The two environments used for comparing parameter effects is shown in fig. 4. The robot's maximum sensor range, R_{max}, is set to 4m. The first environment considers the situation of a robot approaching a narrow passage with a dead end, which are often found in indoor environments. The end of the passage is 4.5m from robot's starting point. Hence, there exists unmapped area towards the dead end and robot selects a target which is 3.5m away and close to the dead-end.

Fig. 5 depicts the travel distance of the robot towards the dead-end with varying γ and p_T values. It can be seen that irrespective of the γ value, when the threshold probability is set to zero, it behaves identically to the classical exploration strategy and continues to move towards the target which is close to the dead end even though the entire passage gets fully explored by a small forward motion. In all other situations, the robot's motion gets cancelled within 0.5-1.0 m forward motion and behaves identically. This is because as the robot moves forward about 0.5m-1.0m, due to the narrowness of the passage, the entire environment is fully explored and there are no more frontiers to be selected as intermediate targets. Hence the two parameters do not affect the cancellation decision differently.

In the second environment type, two targets are used. Target A and B are placed about 0.5m and 1.5m away from the room boundary respectively as shown in fig. 4(b). In both cases, the target is at the frontier with the robot's sensor facing towards the target. In each experiment, the robot is sent to the specified intermediate target in an exploratory motion following a straight trajectory. The travel distance at the time the robot stops/gets interrupted is recorded for each run. For each target A and B, the area the robot could explore is measured by hand. This value is used to generate the expected information gain percentage function for the robot's currently travelled distance. It can be noted that, the remaining information gain exhibits a diminishing returns property with respect to the travelled distance. This is because as the robot moves towards the targets, most of the area in front of the robot gets mapped. Remaining unmapped area resides to the two sides of the robot's motion. Hence, the straight motion of the robot does not map the unknown area with high efficiency. This results in a lower net-gain with increasing travelled distance. This observation justifies cancelling the motion before the robot reaching the assigned target and reassigning a new target, when the robot's motion towards the previous target is less 'desirable'.

The two graphs in fig. 6 and fig. 7 summarizes the results from these experiments and illustrates the effect of γ and p_T in deciding the desirability of the robot's motion. For both targets, setting p_T to zero, irrespective of the γ value, makes the robot behaves exactly similar to the classical exploration approach and moves the robot to the target which is 3.5m away. And also as $\gamma = 90°$ includes the entire neighborhood as the robot's 'heading', the robot moves forward until there are no frontiers in the neighborhood, hence the long travel distances. For target A, it can be seen that, travel distances between 1.0m-1.5m explores between 70%-85% of the neighboring area approximately. Cancelling the sensing task during this travel distance interval can be considered more efficient for the exploration mission as the robot could cover the remaining unknown areas more efficiently by employing a different motion from the current one. For targets that lie very close to obstacle boundaries in a room type environment, such as A, selecting $\gamma = 75°$ and $p_T = [0.1 - 0.5]$ or $\gamma = 60°$ and $p_T = [0.1 - 0.3]$ is observed to generate motions that are neither too agressive nor too conservative in exploring according to fig. 6. For target B, the travel distance interval 2.0m-2.75m explores between 77%-90% of the neighboring area approximately. Hence, similar to target A, it can be considered that selecting $\gamma = 60°$ and $p_T = [0.3-0.5]$ or $\gamma = 45°$ and $p_T = [0.1-0.3]$ generate exploratory motions that exhibit the correct balance of agressiveness and conservativeness for targets that are not too close or far away from obstacle boundaries relative to the sensor range. It can also be observed that for $\gamma = 75°$, the generated motions are either too conservative, for $p_T = [0.1 - 0.7]$, or too agressive, for $p_T = [0.8 - 1.0]$ hence not suitable for exploration missions.

B. The effect of task cancellation on exploration missions

The previous section analyzed the effect on exploration with cancellation of a single sensing task. An exploration mission is a sequence of such sensing tasks. Experiments were conducted to measure the effect of the cancellation strategy on complete exploration missions of an indoor environment, depicted in fig. 8. The travelled distances to complete the exploration missions were recorded. Based on the results from the previous section, $\gamma = 45°$, $p_T = [0.1-0.3]$ and $\gamma = 60°$, $p_T = [0.1-0.5]$ are selected as the parameter space for full exploration missions. For each γ, P_T

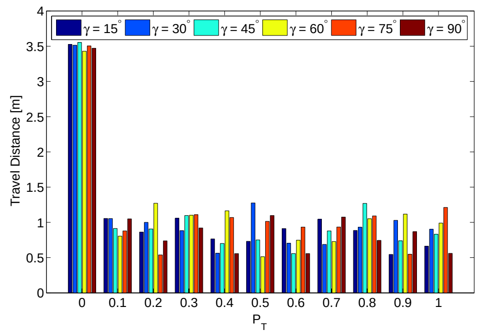

Fig. 5. Travel distances of robot during exploratory motion in the narrow passage

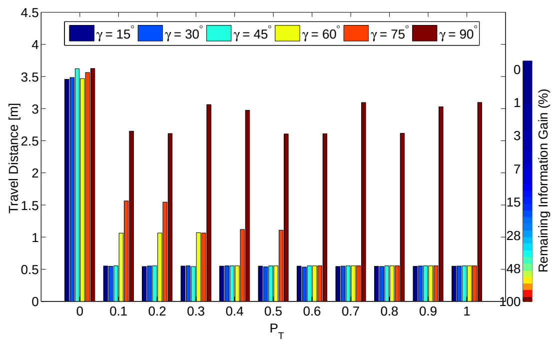

Fig. 6. Travel distance of robot and remaining information gain percentage during exploratory motion to target A

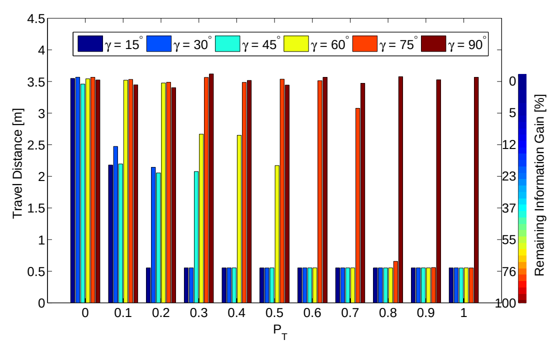

Fig. 7. Travel distance of robot and remaining information gain percentage during exploratory motion to target B

parameter pair used for experiments, 10 runs were conducted with different starting pisitions to negate any bias arising from the starting point. The numbered points in the figure illustrate these starting points. The greedy target selection strategy sometimes makes the robot move longer distances for sensing tasks during the final stages of the exploration mission. Inorder to remove any biases of the results due to this scenario, the travel distances when the environment is 95% explored are also reported. The maximum range of the sensor is set to 3.0m.

Table I summarizes the results for the various exploration missions conducted. The first row corresponds to the results from the classical exploration mission, with $p_T = 0$.

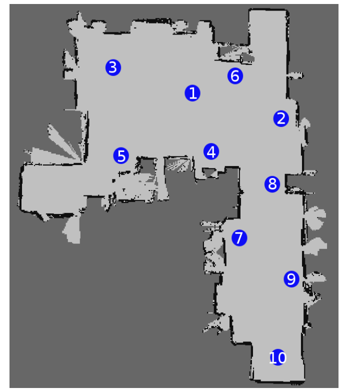

Fig. 8. The map of the indoor area used for experiments. The numbered points are the different starting points used for each (γ, P_T) parameter pair

Overall, augmenting the classical exploration process with cancellation of sensing tasks results in lower travel distances, on average for both the 100% and 95% explored scenarios. However, the standard deviation indicates that, the perceived average performance improvements are not statistically significant. Gain in travel distance during exploration occurs when the classical method selects frontier cells that are near obstacle boundaries as targets and the proposed method avoid reaching such targets. However, selection of such frontier cells by the classical method does not occur in all target selection steps during a mission. In some steps, the selected target may not be closer to any obstacle boundary and both strategies would perform identically on average. Hence, the improvements on the travelled distance by the proposed method gets averaged out over long travel distances and multiple experiments, and reduces the statistical significance of the data. However, it is observed that the number of times the task cancellation based missions having a positive gain for each parameter pair is much higher. Of the total 80 missions executed, 62 have provided a positive travel gain considering 100% exploration and 58 have provided positive travel gains considering 95% exploration. Therefore, inorder to have a better understanding on the effect of task cancellation on missions, the distribution of the travel gains should also be considered. Travel gain distributions for both 100% and 95% explored scenarios, illustrated in fig. 9, indicate that exploration with task cancellation approach generates more efficient motion with high probability. It is also observed that with high p_T values, the number of motion cancellations increase leading to more aggressive exploration missions as predicted. However, the results are inconclusive about the effect of the agressiveness of task cancellation on the efficiency in this environment. Fig. 10 qualitatively compares the efficiency of the proposed method with the classical approach. While the robot ventures in to two narrow passages out of three and moves very close to obstacles as depicted in 10(a) and taking sharp turns (top left of path) in the classical approach travelling 39.91m, the proposed approach avoids entering all the narrow passages and completes the exploration mission in 32.99m and conducts a more efficient exploration mission as expected.

TABLE I
PERFORMANCE OF EXPLORATION MISSIONS

(γ, P_T)	**Avg. Travel Dist.**		**Std-dev.**		**# Avg. Cancels**	**# +ve Gains (out of 10 experiments)**	
	100%	**95%**	**100%**	**95%**		**100%**	**95%**
(*,0)	44.48	39.94	6.53	4.99	-	-	
(45,0.1)	39.36	35.71	5.53	4.84	8	9	8
(45,0.2)	40.29	35.32	4.66	5.28	14	8	8
(45,0.3)	39.23	35.90	3.95	4.24	22	9	8
(60,0.1)	42.31	36.98	3.04	5.28	7	6	6
(60,0.2)	41.44	37.72	3.76	3.82	9	8	7
(60,0.3)	41.20	37.44	4.52	4.10	15	8	6
(60,0.4)	41.91	37.94	4.33	4.55	19	7	7
(60,0.5)	40.13	35.67	4.14	4.80	23	7	8

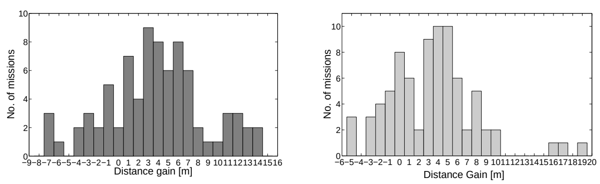

(a) (b)

Fig. 9. (a) Histogram of travelled distance gain for 100% exploration (b) Histogram of travelled distance gain for 95% exploration

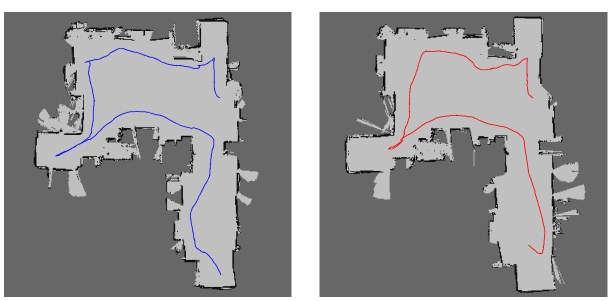

(a) (b)

Fig. 10. Qualitative comparison of classical exploration approach and the proposed approach. Robot start from starting point No. 2. (a) The path taken by robot during the classical exploration mission. Travel distance 39.91m (b) The path taken by robot with interruption. Travel distance 32.99m

VI. CONCLUSIONS AND FUTURE WORK

This work augmented the classical frontier based exploration strategy to include a decision step that cancels sensing tasks if they are no longer desirable. The check for desirability was formulated as a probabilitic decision step and the classical approach is shown to be a special case of this augmented strategy. The experiments revealed that the augmented strategy is capable of conducting efficient exploration missions than the classical approach with high probability. However, the experiments on a single environment type was not sufficient to evaluate the effect of two parameters and variances in efficiency on complete exploration missions though their effect on a single sensing task were evaluated. In future works, we expect to conduct more experiments in different environments to evaluate and find the best range of parameter values for the proposed approach.

REFERENCES

[1] M. A. Batalin and G. S. Sukhatme, "Spreading out: A local approach to multi-robot coverage," in *in Proc. of 6th International Symposium on Distributed Autonomous Robotic Systems*, 2002, pp. 373–382.

[2] I. Nourbakhsh, K. Sycara, M. Koes, M. Yong, M. Lewis, and S. Burion, "Human-robot teaming for search and rescue," *Pervasive Computing, IEEE*, vol. 4, no. 1, pp. 72–79, jan.-march 2005.

[3] D. Apostolopoulos, L. Pedersen, B. Shamah, K. Shillcutt, M. Wagner, and W. Whittaker, "Robotic antarctic meteorite search: outcomes," in *Robotics and Automation, 2001. Proceedings 2001 ICRA. IEEE International Conference on*, vol. 4, 2001, pp. 4174–4179 vol.4.

[4] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings., 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, Jul. 1997, pp. 146–151.

[5] H. P. Moravec, "Sensor fusion in certainty grids for mobile robots," *AI magazine*, vol. 9, no. 2, p. 61, 1988.

[6] H. H. Gonzalez-Banos and J.-C. Latombe, "Navigation strategies for exploring indoor environments," *The International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 829–848, Oct. 2002.

[7] F. Amigoni, V. Caglioti, and U. Galtarossa, "A mobile robot mapping system with an information-based exploration strategy," in *International Conference on Informatics in Control, Automation and Robotics*, 2004, pp. 71–78.

[8] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT Press, 2005.

[9] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, and et al., "Coordination for multi-robot exploration and mapping," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence, AAAI.*, 2000.

[10] G. Oriolo, M. Vendittelli, L. Freda, and G. Troso, "The srt method: randomized strategies for exploration," in *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, vol. 5, april-1 may 2004, pp. 4688–4694 Vol.5.

[11] L. Freda and G. Oriolo, "Frontier-based probabilistic strategies for sensor-based exploration," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*, april 2005, pp. 3881–3887.

[12] B. Kuipers and Y.-T. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *JOURNAL OF ROBOTICS AND AUTONOMOUS SYSTEMS*, pp. 47–63, 1991.

[13] H. Choset and J. Burdick, "Sensor-based exploration: The hierarchical generalized voronoi graph," *The International Journal of Robotics Research*, vol. 19, no. 2, pp. 96–125, Feb. 2000.

[14] S. S. Ge, Q. Zhang, A. T. Abraham, and B. Rebsamen, "Simultaneous path planning and topological mapping (sp2atm) for environment exploration and goal oriented navigation," *Robotics and Autonomous Systems*, vol. 59, no. 3-4, pp. 228–242, 2011.

[15] F. Amigoni, "Experimental evaluation of some exploration strategies for mobile robots," in *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, may 2008, pp. 2818–2823.

[16] M. Keidar and G. A. Kaminka, "Efficient frontier detection for robot exploration," *The International Journal of Robotics Research*, p. 0278364913494911, 2013.

[17] D. Holz, N. Basilico, F. Amigoni, and S. Behnke, "Evaluating the efficiency of frontier-based exploration strategies," in *Proceedings of Joint 41st International Symposium on Robotics and 6th German Conference on Robotics*, Jun. 2010.

[18] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.