ON THE UTILITY OF EXPLORATION ON TIME-CRITICAL MOBILE ROBOT MISSIONS

Kristo Heero

Data Security Laboratory
Cybernetica AS
Akadeemia tee 21, 12618 Tallinn
Estonia
kristo.heero@cyber.ee

Alvo Aabloo

Institute of Technology
Tartu University
Vanemuise 21, 51014 Tartu
Estonia
alvo.aabloo@ut.ee

Maarja Kruusmaa

Institute of Technology
Tartu University
Vanemuise 21, 51014 Tartu
Estonia
maarja.kruusmaa@ut.ee

ABSTRACT

This paper addresses the problem of the utility of exploration on time-critical mobile-robot missions. It is argued that in large environments mission-oriented mobile robot applications can become more efficient if the exploration strategy considers knowledge already gained and its applicability during the rest of the mission. This hypothesis is verified in a model test environment with Khepera robot. The conclusion is that that mission-oriented exploration heuristics could be considered in mobile robot applications that are time-critical, where the robot is operating in a large unknown environment and if this environment is hazardous.

1. INTRODUCTION

A mobile robot operating in a real-world environment faces several fundamental problems. Most important of them is the ability to build the model of the environment [1], to plan routes [2] and to follow them while avoiding unknown obstacles [3].

These problems need to be tacked in all applications of mobile robots, such as transportation, surveillance or guidance [4,5]. Vast amount of mobile robot research addresses the problem of exploration and environment mapping [6,7]. Gaining knowledge about the surrounding environment and keeping it updated is the necessary precondition of successful performance.

In practical applications, exploration of the environment is usually risky and time-consuming. The environment can be hazardous and damage the robot. Exploration of a large environment takes lots of time and computational resources.

From the utilitarian point of view, knowledge about the environment is useful only as long as it increases the performance of the robot. In many cases the application does not require the exploration of the whole environment (e.q. fetch and carry tasks) while in other applications, such as demining or search and rescue, the task definition implies a systematic search of the whole area [8,9,10].

This paper addresses the problem of the utility of exploration for time-critical mobile robot missions. It is assumed that the environment is very large and therefore exploration is time-consuming. The environment can also be hazardous and degrade the performance of the robot or slow it down. The assumption is that exploration and mapping are not goals by itself but means that permit the robot to fulfill its mission.

Related work that consider the utility of route planning usually do not address the utility of exploration but rather evaluate the risk of navigation locally, e.g. in terms of possible collisions or the characteristics of the terrain [11, 12, 13, 14].

We propose a heuristic exploration strategy that chooses between exploring new areas and exploiting knowledge about the already explored areas. The decision-making is based on the mission plan. The heuristic decision maker takes into consideration the amount of knowledge acquired so far and its applicability during the rest of the mission.

We test the strategy in a model environment with the Khepera robot. The test results show that this mission-oriented heuristic can be useful for mobile robots on time-critical missions.

2. MISSION-ORIENTED EXPLORATION

In this section we describe the problem and outline the exploration strategy of the robot.

We assume that the robot is operating in a previously unknown environment. The goal of the robot is to fulfill a mission plan. The mission plan is known in advance, consisting of target points that the robot has to reach in a predefined order (e.g. a transportation task, escorting or surveillance problem).

The target points are defined by their global coordinates. This paper is concerned about exploration strategy of time critical missions and therefore we do not address the problem of map building and localization in this context. We therefore assume that the robot is able to localize itself rather accurately (e.g. as with GPS or pseudolite localization). It is also assumed to have an environmental model in the form of a grid-based map. In the beginning of the mission the environment is unknown and the map contains only very general information (the size and shape of the environment). The robot has no knowledge about the obstacles or any other environmental factors that can degrade its performance. The robot learns the environment while traversing it and completing its mission. It updates the map when it detects obstacles with its on-board sensors.

2.1. Exploration Strategies

We verify two exploration strategies that are exploration oriented to a different extent. The bottom line of the first strategy is to always take a new route to the target (i.e. explore the environment) if it is expected to be better than the routes known so far.

The second strategy has a more conservative attitude against exploration. The difference from the first, greedy strategy is that, when gaining new knowledge it also considers the mission plan. The new knowledge is gained more probably when it is often used during the rest of the mission. Also knowledge that is used in the nearest future is gained more probably than knowledge used after a long time.

The robot has to reach predefined target points $\mathrm{G} = \{\mathrm{g}_1, \mathrm{g}_2, ..., \mathrm{g}_\mathrm{n}\}$ where g is the grid cell on the map of the robot. The mission consists of traversing the target points in a predefined order $\mathrm{M} = \mathrm{m}_1, \mathrm{m}_2, \mathrm{m}_3, ..., \mathrm{m}_\mathrm{i}, \mathrm{m}_{\mathrm{i}+1}, ..., \mathrm{m}_\mathrm{k}$, where $\mathrm{m}_\mathrm{j} = (\mathrm{g}_\mathrm{u}, \mathrm{g}_\mathrm{v})$ and $\mathrm{m}_{\mathrm{j}+1} = (\mathrm{g}_\mathrm{v}, \mathrm{g}_\mathrm{w})$ for $1 \le j < k$, if $1 \le u, v, w \le n$.

In addition to the map that is constantly updated the robot also saves the entire followed path P. A path is stored as a sequence of grid cells. For every task m_i of the mission M it can choose between using an already followed path and a new path. The new path can contain segments that traverse unexplored regions.

The traversed paths are stored together with statistics characterizing their traversability. The average time of following a path $t(P)$ is used later when the strategy chooses between exploration and exploitation of the known tracks.

Every time the robot traverses the environment the map is updated so that the knowledge about the environment accumulates during the mission and every time when a new path is planned this new knowledge is taken into account.

2.2. Greedy Exploration Strategy

The greedy exploration strategy always chooses a new path if it predicts it to be better than the best known one. The predicted average time of the new path $t(P_{new})$ is verified to the best average time of the paths stored so far $t(P_{best})$. If $t(P_{new}) < t(P_{best})$ then the new path P_{new} is chosen.

While following the planned path, the robot does not try to stay on the predefined path if the obstacles are encountered but replans a new path to the target point through the possibly unexplored regions.

2.3. Conservative Exploration Strategy

This exploration strategy makes the decision between using the best-known path P_{best} and a new path P_{new} by considering exploitation of this knowledge in the future. It chooses P_{new} if the task m_i is not encountered often in the past and if it is needed often during the rest of the mission. The sooner during the mission new knowledge will be needed the more P_{new} is preferred.

Let

$$Past = \sum_{p=1}^{i-1} \begin{cases} 1, m_p = m_i \\ 0, otherwise \end{cases}$$

denote the number of similar tasks completed in the past. Since it was assumed that the robot knows its mission plan it can be counted how many times a task similar to m_i has to be completed in the future.

$$Future = 1 + \sum_{r=i+1}^{k} \begin{cases} \dfrac{\sum_{q=i+1}^{r} \begin{cases} 1, m_q = m_i \\ 0, otherwise \end{cases}}{r-i}, m_r = m_i \\ 0, otherwise \end{cases}$$

If $\dfrac{Past}{Future} > 1$, then it is determined that the robot has gained enough knowledge about the environment and the further exploration is not beneficial. The path P_{best} will be chosen and followed. If unexpected obstacles are encountered during the path following, the robot replans the path but tries to return back to the initially chosen path P_{best} to avoid unexplored regions.

If $\dfrac{Past}{Future} \leq 1$ then the robot chooses P_{new} because very little of the environment is still explored or because the new knowledge gained can be used in the future to increase the performance.

3. EXPERIMENTAL SETUP

3.1. Test Environment

The experiments are conducted using a mini-robot Khepera. It is a differential drive miniature circular robot (with radius 26 mm) equipped with IR sensors for collision avoidance and it can be connected to a PC over a serial link.

The localization system is presented in Figure 1. A video camera is mounted to the ceiling to recognize the position and orientation of the robot. The PC processes the camera image to find robot's position and a computer algorithm controls the robot over a serial link. In this way localization errors are rather small (usually comparable to the size of the robot).

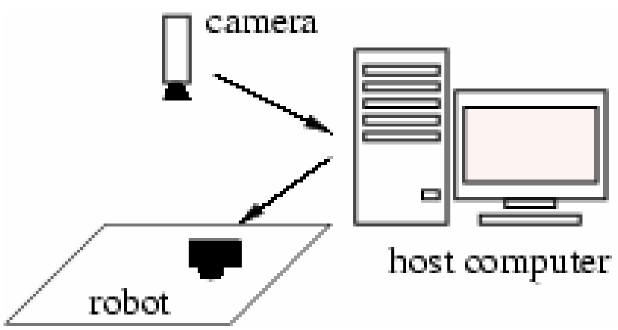

Figure 1. The experimental setup.

The test environment is represented in Figure 2. The size of the test environment is 2320mm × 1710mm.

Figure 3 shows the test environment as seen from the overhead camera. The position and the orientation of the robot are recognized with the help of 3 LEDs forming an equilateral rectangle.

Figure 4 is the graphical interface of the computer program that controls the robot and monitors its behavior. The coordinates of the target points are marked with the symbols G1, G2, etc. The thick line represents the path of the robot to the goal. Black cells represent obstacles detected with the onboard sensors. The gray and dark gray boxes are respectively unknown areas of slow and extra slow motion.

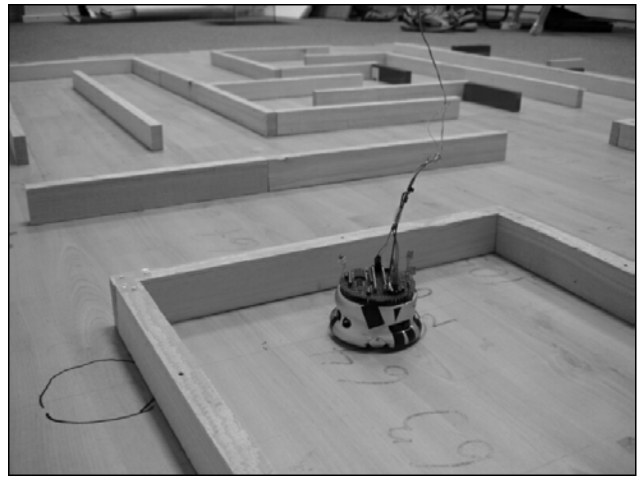

Figure 2. The test environment.

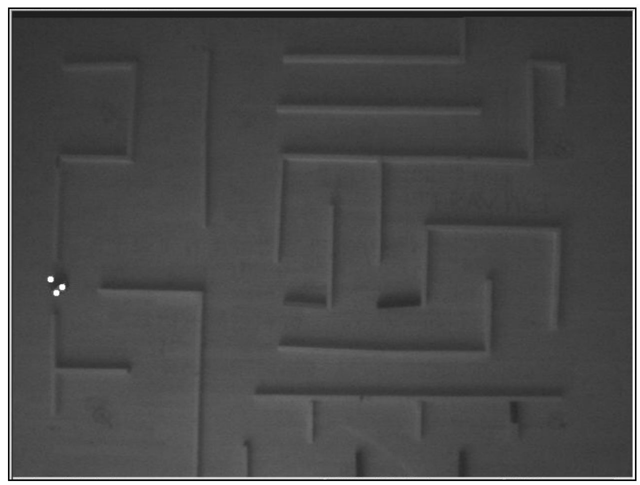

Figure 3. The test environment looked through the overview camera with the robot recognized by the 3 LEDs.

In the beginning of the mission the robot is not aware of any obstacles in the environment. While it traverses the environment it updates the map and records all the detected obstacles so that as the mission proceeds its environmental model becomes more and more complete and consistent.

In addition to the obstacles there are some regions in the environment where the motion of the robot is slowed down. These regions are introduced to simulate regions that in real robot applications are difficult to traverse (e.g. because of rough terrain). The robot is not aware of the presence and location of such regions and these are also not reflected on the map. These areas are represented in the Figure 4 as the shaded regions.

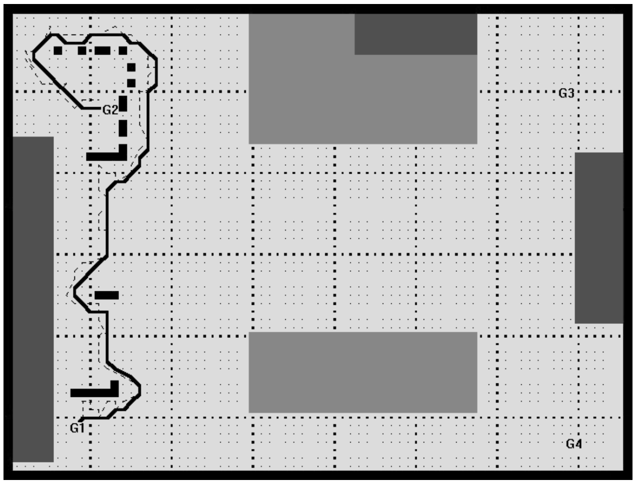

Figure 4. The control interface of the robot.

3.2. Test trials

The purpose of the experiments is to test whether the conservative exploration strategy gives better results than the greedy strategy in the presence of hazard and uncertainty.

We test both missions in equal conditions. The robot is given a mission consisting of 31 tasks and the goal of the robot is to fulfill the mission as fast as possible.

The environment, the task and the robot are identical in both cases. The test trials only differ by the exploration strategy used.

The mission plan (the sequence of target points to be traversed) is the following:

G1 G2 G1 G3 G1 G2 G4 G2 G1 G2 G1 G3 G1 G2 G4 G2 G1 G2 G1 G3 G1 G2 G4 G2 G1 G2 G1 G3 G1 G2 G4 G2

The target point G1 corresponds to the point marked with G in Figure 4 in the lower-left corner, G2 in the upper-left corner, etc. in the clockwise direction. This plan implies that the robot traverses often between G1 and G2 but quite seldom between the target points G2 and G4 or G1 and G3. Some parts of the environment, like those between G4 and G3 are not traversed at all.

Since the different regions of the map are traversed with the different intensity, exploring some regions becomes more important from the point of view of the mission than exploration of other regions.

The model environment is kept static during the mission. No dynamic obstacles are introduced. Obviously, this is not a realistic assumption on real robot missions. However, the goal of the tests is to show the advantages or disadvantages of an exploration strategy and the static environment guarantees that if one exploration strategy outperforms another then this is caused by the strategy but not by the changes in the environment.

Although the environment is kept static, the robot in the model environment still has to tackle problems caused by uncertainty of sensor readings, odometric errors and small localization errors due to the image recognition system. The generation of new paths P_{new} is stochastic and therefore the performance of the robot depends to a great extent on a stochastic algorithm. We therefore conducted several pairs of test trials to show how much the test results diverge.

The map and the memory (containing the traversed paths and their average traveling time) are stored after every task. All test data is available at http://math.ut.ee/~kristo/khepera/heuristic/

4. EXPERIMENTAL RESULTS

The goal of the robot was to fulfill the mission as fast as possible. Therefore the average time of the mission is the most important parameter indicating the efficiency of the exploration strategy.

The chart in Figure 5 shows the duration of the mission. Time of fulfilling 31 tasks is 16.74% shorter when the conservative exploration strategy is used. While both of the exploration strategies perform equally well in the beginning of the mission, the conservative learning strategy starts outperforming the greedy strategy after 10 first tasks.

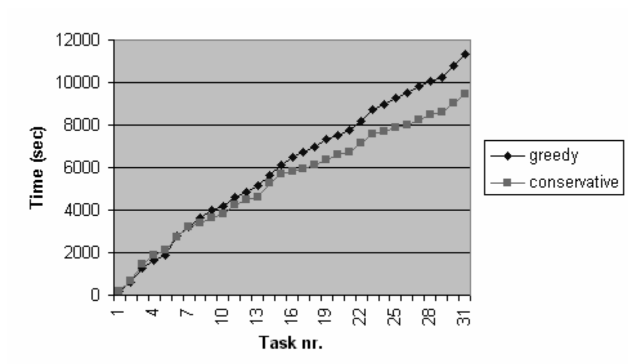

Figure 5. The duration of the mission.

The white area in Figure 6 show regions explored during the mission when the conservative exploration strategy is used. The mission plan requested frequent traversing between the lower left and upper left corner of the environment and it appears that the robot has explored these regions most extensively while the rest of the environment is searched less thoroughly.

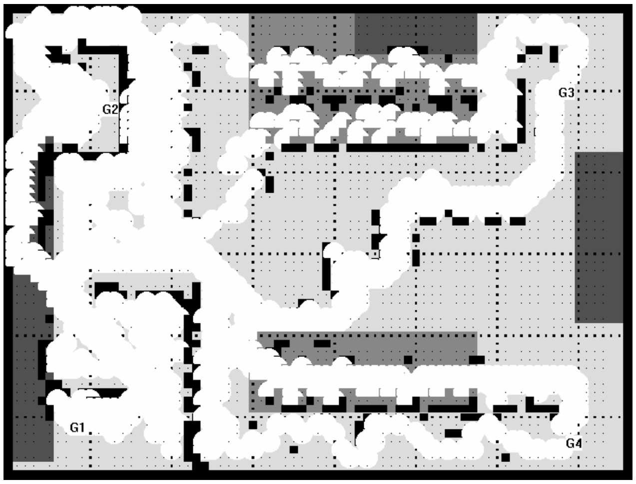

Figure 6. Explored areas with the conservative exploration strategy.

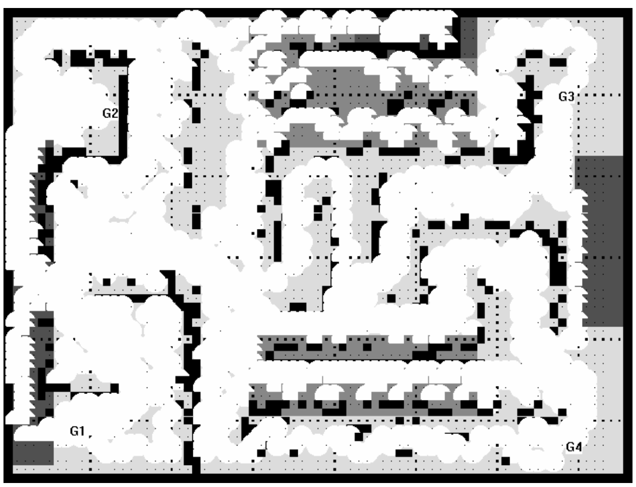

Figure 7. Explored areas with the greedy exploration strategy.

The white area in Figure 7 show regions explored with the greedy exploration strategy. It appears that almost the whole environment is searched and most of the obstacles are mapped. The robot has also been extensively traversing the areas of slow motion (marked with gray) that have eventually slowed the mission down as well as frequent replanning and maneuvering around detected obstacles.

Another indicator that shows the difference between the conservative and the greedy exploration algorithm is the number of reused paths. Traversed paths are stored in robot's memory together with statistics showing their average traversing time. Results show that the greedily exploring robot used already traversed paths in 51% of tasks while the conservative robot relied on its past experiences in 67% of the cases.

5. CONCLUSIONS

This paper presented an exploration strategy for a mobile robot in large hazardous environments. It was presumed that the robot is working under time constraints. We presented a heuristic exploration strategy that chooses between exploration and exploitation considering the amount of knowledge gained so far and the applicability of this knowledge during the rest of the mission.

We simulated a time-critical mission by conducting experiments in a model environment with a Khepera and verified the heuristics with a greedy exploration strategy. The test results showed that the robot using the conservative exploration strategy is able to fulfill the mission approximately 17% faster than the robot using the greedy exploration strategy.

These test results reveal that it can be useful to choose between exploration of the environment and the exploitation of the knowledge gained depending of the nature of the mission and the environment. Gaining as much knowledge as possible about the surrounding is not necessary beneficial if the mission time is limited.

The performance of the conservative exploration strategy undoubtedly depends on the environment where the robot is operating and on the mission assigned. We therefore are careful with generalizing these results too much. Certainly, there exist environments and assignments where the greedy exploration strategy would be more efficient. More experiments in different environments (e.g. cluttered, free space, corridor-environments, multiple rooms, etc.) and different path planning methods are required to validate the presented approach.

We conclude that mission-oriented exploration heuristics could be considered in mobile robot applications that are time-critical, where the robot is operation in a large unknown environment and when this environment is dangerous. We also suggest that this or similar heuristics can be applied for other learning problems in mobile robotics.

ACKNOWLEDGEMENTS

This paper is supported by Estonian Science Foundation grant No. 5613, and Tiger Leap Program of The Estonian Information Technology Foundation.

6. REFERENCES

[1] Thrun, S., Burgard, W., and Fox, D. A probabilistic approach to concurrent mapping and localization for mobile robots. *Machine Learning* Vol. 31, pp. 29-53, 1998.

[2] D. Ferguson and A. Stentz, "The Delayed D* Algorithm for Efficient Path Replanning." *Proc. of the 2005 IEEE Int. Conf. on Robotic and Automation, (ICRA 2005),Barcelona, Spain, April, 18-22.*

[3] J. Latombe, "Robot Motion Planning." Kluwer Academic Publishers, Boston, MA, 1991.

[4] B.Brumitt, A. Stentz, M. Herbert and The CMU UGV Group, "Autonomous Driving with Concurrent Goals and Multiple Vehicles: Mission Planning and Architecture," *Autonomous Robots*, Vol. 11, Kluwer Academic Publishers, pp. 103-115, 2001.

[5] B.Brumitt, A. Stentz, M. Herbert and The CMU UGV Group, "Autonomous Driving with Concurrent Goals and Multiple Vehicles: Experiments and Mobility Components, *Autonomous Robots,* Vol. 12, Kluwer Academic Publishers, pp. 135-156, 2002.

[6] Thrun, S. Robotics mapping: A survey. Tech. Rep. CMU-CS-02-111, School of Computer Science, Carnegie Mellon University, Pittsburg, PA 15213, February 2002.

[7] F. Amigoni and A. Gallo, "A Multi-Objective Exploration Strategy for Mobile Robots." *Proc. of the 2005 IEEE Int. Conf. on Robotic and Automation, (ICRA 2005),Barcelona, Spain, April, 18-22.*

[8] E. Garcia and P. Gonzalez de Santos, "Mobile-robot navigation with complete coverage of unstructured environments," *Robotics and Autonomous Systems*, Vol. 46, Elsevier Science, pp. 195-204, 2004.

[9] H. Choset, "Coverage of Known Spaces: The Boustrophedon Cellular Decomposition," *Autonomous Robots*, Vol. 9, Kluwer Academic Publishers, pp. 247-253, 2000.

[10] A. Poncela, E. J Perez and A. Bandera, C. Urdiales, F. Sandoval, "Efficient integration of metric and topological maps for directed exploration of unknown environments," *Robotics and Autonomous Systems,* Vol. 41, pp. 21-39, 2002.

[11] R. Manduchi, "Obstacle Detection and Terrain Classification for Autonomous Off-Road Navigation," *Autonomous Robots*, Vol. 18, Kluwer Academic Publishers, pp. 81-102, 2005.

[12] T. Belker, M. Hammel and J. Hertzberg, „Learning to Optimize Mobile Robot Navigation Based in HTN Plans," *Proc. Of the 2003 IEEE Int. Conf. on Robotic and Automation, (ICRA 2003),* Taipei, Taiwan Sept. 14-19, pp. 4136 – 4141, 2003.

[13] Z. Shung and J. Reif, „On Energy-minimizing Path on Tarrains for a Mobile Robot," *Proc. of the 2003 IEEE Int. Conf. on Robotic and Automation, (ICRA 2003),* Taipei, Taiwan Sept. 14-19, pp. 3782 – 3788, 2003.

[14] R. Philippsen and R. Siegwart, "An Interpolated Dynamic Navigation Function." *Proc. of the 2005 IEEE Int. Conf. on Robotic and Automation, (ICRA 2005),Barcelona, Spain, April, 18-22.*