International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS), 2001.

A Multiple Information Source Planner for Autonomous Planetary Exploration

Stewart J. Moorehead, Reid Simmons, William L. Whittaker
Robotics Institute
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA. 15213
USA
{sjm, reids, red}@ri.cmu.edu

Keywords exploration, information theory, mobile robots, planning, traversability maps.

Abstract

This research casts the problem of exploration into one of collecting information and enables the completion of complex exploration tasks by considering multiple sources of information. This paper describes a methodology to solve exploration tasks which computes the expected information gain, for each information source, from a sensor reading. It also proposes a greedy search algorithm to collect information while accounting for costs such as driving, sensing and planning. An instantiation of the methodology to solve the exploration task of creating traversability maps is presented. This task uses three information sources: frontier detection, increase map certainty and determine reachability. These sources are explained in detail. Finally, the traversability map creation task is performed in simulation.

1 Introduction

As planetary exploration mission goals become more ambitious, communication limitations require that exploration robots be more autonomous. They must manage finite resources efficiently to discover and record valuable information and they must make decisions based on this information to gather additional information.

When exploring, a robot will have many sources of information adding to its knowledge of an area. For example, a robot searching for water ice on the Moon would need to consider lighting information (to identify permanently dark regions where water can exist [11]), drilling conditions, age of crater and the location within a crater [2]. These information sources are different from a robot searching for signs of life on Mars, which may consider the presence of water, both present and historical, the geology of rocks and the presence of organic material. Thus the number and types of information sources are determined by the exploration task. The exploring robot must balance these competing information sources to collect the maximum amount of information with its limited resources.

Another challenge for robotic explorers is deciding where to go. Unlike traditional path planning where robots make decisions to get to a specified destination, robotic explorers must make decisions that help them collect information from varied sources in an efficient manner. They do not have a destination to reach but rather tasks to complete. For example, the explorer searching for ice on the Moon must decide if it should go to point A where it can learn a lot about lighting conditions, or go to point B and drill for water or to go to point C and learn a little about lighting conditions and crater age. Thus an explorer must decide amongst many, perhaps even all, possible destinations to further its task. To complete its exploration it will have to repeat this process many times.

This paper presents a methodology for solving complex exploration problems that is able to consider multiple, competing sources of information when planning exploration actions. The methodology computes the expected gain in information, for each information source, that results from a robot action. Using these expected information gains, a path is planned which maximizes the information collected, while minimizing costs such as driving and sensing times. The method has been applied to an indoor exploration problem [10] and can be adapted to many different exploration tasks by changing the sources of information.

2 Related Work

A frontier is defined as a boundary between known and unknown regions in a map. Yamauchi introduces a technique for exploring based on frontiers [15][16]. The exploration strategy used is to identify all the frontier regions in the current map and then drive to the nearest frontier. This process is repeated until the region has been explored. This method has two shortcomings for the exploration tasks considered here. First, it treats all frontiers equally. Rather than looking at the potential information to be gained at each frontier versus the time it will take to get there, it simply chooses the closest frontier. In some situations it may be better to drive further to gain more information. Secondly, it is limited to one source of information, finding new terrain.

Simmons et al. [14] look at using multiple robots to explore using a similar frontier strategy. However, each frontier is evaluated based on the expected number of unknown map cells the robot can see from the frontier as well as the distance from the robot. Thus the exploring robots choose the frontier which will provide the highest utility (information gain minus driving cost) rather than simply the closest frontier. The multiple robots also coordinate their efforts to reduce overlap while exploring.

In [3][4], Elfes presents an integrated approach to robot navigation which incorporates task specific information needs, perception sensor capabilities and robot knowledge into the motion planning process. An Inference Grid is used to represent the robot's knowledge and task needs. The integrated architecture allows the robot to plan both motor and perceptual actions to solve a given task. This approach is different than the two discussed above in that it is able to consider more than one source of information to solve the exploration task.

Roy et al. [13] look at including the ability to localize the robot, using natural landmarks, when planning paths. This produces paths which remain close to walls. While [13] was not trying to explore new regions it is considering two criteria, shortest distance and localization ability, when planning paths. Further, the ability of a robot to localize itself could be considered as an information source while exploring.

Estlin et al. [5] present a planning system for science exploration with multiple robots. The system clusters spectrometer readings to determine rock types. Goal points are set at the two mutually most distant points in physical space for each rock type seen so far. The planner is thus biased to explore towards the edges of its known world. The approach of Estlin et al. differs from that of this paper in that they only consider one type of information, rock types, in their goal selection. Further, they do not explicitly consider how much information will be gained when selecting goal points, choosing instead to use a heuristic based on relative rock positions.

3 Exploration

Robots record information as they explore, creating a world model. An explorer will encounter multiple sources of information and must be able to estimate the potential information to be gained from multiple sources in all locations. The explorer must also be able to use this world model to plan routes which allow it to explore efficiently.

This research models the world as a uniform grid, $\boldsymbol{M(t)}$, which takes the robot's pose, $\boldsymbol{X}$, and maps it to a unique cell. The map, $\boldsymbol{M(t)}$, is a function of time because it changes as sensor data is collected and added to the world model. Each cell in the map contains two vectors: $\boldsymbol{A}$ a vector of cell properties or attributes and $\boldsymbol{G}$ a vector of expected information gains. The size and composition of $\boldsymbol{A}$ and $\boldsymbol{G}$ depend on the exploration task being considered. A specific implementation of $\boldsymbol{A}$ and $\boldsymbol{G}$ can be found in section 4.

The cell attribute vector $\boldsymbol{A}$, contains the information the explorer knows about that cell. Some typical types of information stored in the elements of $\boldsymbol{A}$ are cell height and traversability. Each element of $\boldsymbol{A}$ represents a different type of information. For binary valued variables this number is the probability of the state being true. This is the approach used in Inference Grids [3]. For non-binary valued variables, such as cell height, a pair of numbers are used. The first indicates the property value and the second is the explorer's certainty in that property value. Typically the certainty is based on the number of sensor readings received in that cell. While it would be more rigorous to compute and store the probability distribution of these non-binary variables, the method proposed here is easier to implement. A similar strategy has been successfully employed in outdoor navigation [9].

For example a robot exploring the surface of Mars might be interested in mapping the terrain, finding traces of present or past water and discovering signs of life. Thus, the explorer would be interested in recording the height and traversability of a cell. It would also be interested in recording the probability of the cell containing sedimentary rock (more likely to contain fossils), water and living organisms. The attribute vector would have five elements: height, traversability, probability of sedimentary rock, probability of water and probability of living organisms. The first two elements of $\boldsymbol{A}$ are pair elements having a number for the quantity (height or traversability) and a certainty. The last three are binary valued attributes.

The expected information gain vector, ***G***, has one element for each source of information being considered by the explorer. Each element in ***G*** has a value from 0 to 1 and represents the expected amount of information to be gained by taking a sensor reading in that cell for a particular information source. It is the expected information gain since it is predicting how much new information will be received by taking a sensor reading in this cell.

The information sources in ***G*** generally correspond to the type of data being recorded in the attribute vector, ***A***. For example, the Martian explorer described above might have information sources for detecting new terrain (increasing its information of height and traversability), identifying sedimentary rock, detecting water and finding life.

It is important to note that ***G*** measures how much more information the robot will receive by taking a sensor reading in the cell. The sensor provides information over a region and thus ***G*** depends on what the robot currently knows about this region. Taking a sensor reading in one cell will modify the values of ***G*** in other map cells. In other words ***G*** is based on the knowledge contained in the *current* map, $\boldsymbol{M(t_p)}$, where t_p is the present time. This means that the ***G*** vectors in different map cells are dependent on each other.

The total expected information to be gained by taking a sensor reading in map cell *m*, is computed by performing a weighted sum of the elements in cell *m*'s ***G*** vector:

$$E[I(m, M)] = \sum_i \alpha_i g_i^m \qquad (1)$$

where the α's are chosen to weight the relative importance of the various information sources.

3.1 Planning Exploration Paths

An autonomous explorer must decide where to drive and where to take sensor readings to maximize the amount of information it collects. At the same time it must minimize the costs of collecting this information such as driving time, sensing time and planning time. The goal of the exploration planner is to find a path, *p*, which maximizes the utility to the explorer. The path, *p*, is an ordered set of cells that the explorer must drive through or take sensor readings in. The utility of a path, *p*, is defined as:

$$U(p) = k \sum_{m \in p} S(m)E[I(m, M)] - \sum_{m \in p} C_c(m) - C_G \quad (2)$$

$E[I(m, M)]$ is the expected information gain of cell *m* as computed in equation (1). $S(m)$ is 1 if a sensing action is performed in cell *m* and 0 otherwise - indicating that no information is gained unless a sensor reading is taken. $C_c(m)$ is the per cell cost which is the amount of time, in seconds, spent in cell *m*. This includes the driving time and sensing time. C_G is also in seconds and includes any global costs such as planning time. Finally, *k* is the *value of information* and is used to set the relative importance of information and cost.

Unlike traditional path planning problems, an explorer does not have a destination cell to plan a path to. Instead, the exploration planner must maximize $U(p)$ over all possible paths. This is similar to the prize collecting travelling salesman problem (PCTS) where each map cell is a city and the information to be gained is the prize. Since the PCTS problem is NP-complete [1] it is unlikely that an optimal *p* can be found for the exploration problem.

Unlike the PCTS problem the $E[I]$'s in the exploration problem are not independent but depend on ***M(t)***, the current map of world knowledge. Every time the planner decides to take a sensor reading, ***M(t)*** changes, and so do the $E[I]$'s in the neighborhood of the sensing action. This means that the $E[I]$'s in equation (2) depend not only on the cell location but also on the path used to get there. This greatly increases the search space required to find *p*.

Due to the dependence of $E[I]$ on path we consider a greedy planner which plans only to the next sensor reading. The planner first propagates driving costs through the map using a wave-front propagation technique [8]. Then it chooses the cell with the maximum utility as computed by equation (2). The robot will travel to that cell, take a sensor reading and then the planner will determine the next cell to visit. Since the greedy planner plans only to the next sensor reading the plans produced are in the form of a list of cells to drive through and ends with a sensor reading. This means that ***M(t)*** remains constant throughout the plan (it only changes after the sensor reading takes place). Therefore the dependence of $E[I]$ on path is eliminated. A similar planning method was used in [14].

4 Creating Traversability Maps

Section 3 presented a general framework for autonomous exploration which can be applied to numerous different exploration tasks. This section applies the methodology to a specific exploration problem - exploration for the creation of traversability maps. Traversability maps are maps which indicate how easy it to drive over an area. They tell the robot where it is safe to drive and where it is not.

The goal of the exploration robot in this task is to create a traversability map which is useful for other robots that might operate in this region at a later date. The traversability map created will have the traversability of a cell as well as how certain the robot is about that traversability. It is more important to know the traversability in rough terrain and near obstacles than in flat, benign terrain, both for path planning - if we think a narrow passage is traversable and it is not, this could drastically alter a plan - as well as robot safety - it is easier to damage the robot in rugged terrain than in flat terrain. Thus the explorer will increase the certainty in low traversability regions more so than in flat regions. It would also be useful for another robot visiting the area to know which cells are reachable and which cells are not reachable. Finally, the map will have the height and certainty of the height in each cell. This will allow robots to calculate sensor visibility regions. The height of a cell is also used to compute traversability so it is important for the explorer to record.

The exploration robot will have a sensor capable of viewing the terrain being explored. The nature of the sensor is not important for the exploration planner. The planner will know how high off the ground the sensor is mounted as well as the azimuth and elevation field of views. Currently we assume that the sensor has a 360 degree azimuth field of view. This simplifies the planning by removing the dependence on heading.

To perform the traversability map exploration task, the attribute vector, $\boldsymbol{A}$, of the map needs three elements: height ($a_{h,}$ a_{ch}), traversability ($a_{t,}$ a_{ct}), and reachability (a_r). Height and traversability as continuous valued variables are pair elements with a quantity and a certainty. The height is the maximum height that the robot has perceived in the cell. It is relative to some global, fixed point. The height certainty is a number from 0 to 1 which is proportional to the number of sensor readings received in a cell.

The traversability of a cell is computed by fitting a plane, centered at the cell in question, to the cell height data in a region equal to the size of the robot. The traversability is determined by the roll and pitch of the plane as well as the residual from fitting the plane [9]. Using a plane the size of the robot produces a traversability score that is in configuration space. If the origin of the robot is in a cell with good traversability, this means that all parts of the robot are in good traversability. Alternately, if the cell has poor traversability some part of the robot is on dangerous terrain, perhaps a wheel would be in a deep hole. This use of configuration space traversability means that the planner can consider the robot to be a point robot [8]. The certainty in the traversability is related to the certainty in the height data used to calculate traversability. A traversability computed over a region of low height certainty would be less certain than one computed over a region with high height certainty. To maximize robot safety, the worst case scenario was chosen and the traversability certainty is set to the minimum, or worst, height certainty over the region fitted with the plane.

The reachability of a cell is a binary valued quantity - the cell is reachable or it is not. Therefore, the reachability element in $\boldsymbol{A}$, a_r, denotes the probability that a cell is reachable. A cell is reachable if the robot can drive to that cell from any other cell in the set of reachable cells. By definition a reachable cell is traversable, however, a traversable cell may not be reachable if no traversable path from the start location to it exists. The set of reachable cells is connected in that every reachable cell has at least one adjacent cell which is also reachable.

The reachability of a cell, a_r, is computed by assigning $a_r = 1$ for any cell previously visited, $a_r = 0.5$ for unknown cells and $a_r = 0$ for untraversable cells. The remaining cells are set using a decaying exponential based on the cost of driving to that cell from a cell where $a_r = 1$, the higher the driving cost the lower the value of a_r.

The exploration robot's map will also have an expected information gain vector, $\boldsymbol{G}$, in each cell with three elements: g_f, frontier information, g_c, increase certainty and g_r, determine reachability. These expected information gain quantities indicate how much information will be gained, over an entire sensor footprint, by taking a sensor reading in that cell.

Let W_m be the set of cells visible to the sensor from cell m and let $\#W_m$ be the number of cells in W_m. To compute W_m we first get the set of cells $\overline{W}_m$ which contains all the cells inside a circle centered at m with radius equal to the maximum range of the sensor. For each cell, n, in $\overline{W}_m$ a ray is traced back to cell m using the efficient Bressenham's Algorithm [6]. If the height of this ray is lower than the height in any cell it passes through then cell n is not visible and not in W_m. Otherwise cell n is added to the set W_m. Thus the shadowing effect caused by known obstacles is taken into account. Unknown cells are assumed to contain no features which obstruct sensor viewing.

The frontier information source indicates how much previously unseen terrain the explorer can expect to see. This information source is used to attract the robot explorer to the boundary of its known and unknown world and fill in the blank spots in its map. The greater

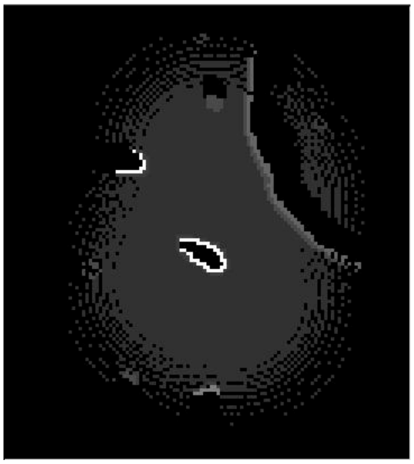

(a) Height

(b) Height Certainty

(c) Frontier

(d) Increase Certainty

(e) Determine Unreachable

Figure 1 (a) Height values of the robot's map. The higher a cell's height the whiter its color. Black is unknown. (b) The certainty in the height values. Black is zero, white is one. (c, d, e) The expected information gains for the three information sources. Black is zero, white is one.

the number of unknown cells expected to be viewed from cell *m*, the greater the expected frontier information gain. For a given cell, *m*, in the map the expected frontier information gain is calculated as:

$$g_f^m = \begin{cases} \dfrac{\sum_{\forall n \in W_m} \mathrm{UNKNOWN}(a_{ch}^n)}{\# \mathrm{W}_m} & ;\ \text{if m is a frontier cell} \\ 0 & ;\text{otherwise} \end{cases} \quad (3)$$

where a_{ch}^n is the height certainty in cell *n* and UNKNOWN(a_{ch}^n) is 1 if a_{ch}^n is less than a fixed threshold and 0 otherwise. A frontier cell is one which is known and traversable and has at least one cell adjacent (in an 8 connected sense) to it which is unknown. Figure 1(c) shows the value of the expected frontier information gain for a partial map. Notice that the non-zero information gains are on the edge of the known and unknown world.

The increase certainty information source rewards the explorer for increasing the density of sensor readings in a cell and thus increasing the height certainty. The increase certainty information source computes the expected increase in height certainty due to a sensor reading. As mentioned earlier in this section, it is more important to have high certainty about the terrain near obstacles so the increase certainty information source weights the expected increase in height certainty by the traversability. Height certainty is used because it can be predicted with a sensor model and traversability certainty is derived from height certainty. The equation for g_c is:

$$g_c^m = \frac{\sum_{\forall n \in W_m} w(a_t^n)(E[a_{ch}^n] - a_{ch}^n)}{\#W_m} \quad (4)$$

where $w(a_t)$ is a parabola with a value of 1 for zero traversability and 0.05 for traversability of one. $E[a_{ch}^n]$ is the expected value of the height certainty in cell *n* after taking a sensor reading in cell *m*. The computation of $E[a_{ch}^n]$ depends on the sensor being used. This paper assumes a sensor which takes range measurements with fixed angular increments in both azimuth and elevation. A sensor such as a laser scanner would fit in this class. For this type of sensor the number of readings in a cell (which is proportional to the cell's height certainty) is inversely proportional to the cube of the range [7]. The function of $E[a_{ch}^n]$ used in this paper is shown in Figure 2. Since the height certainty cannot be greater than one Figure 2 has been limited to one indicating a region around the sensor yielding perfect certainty.

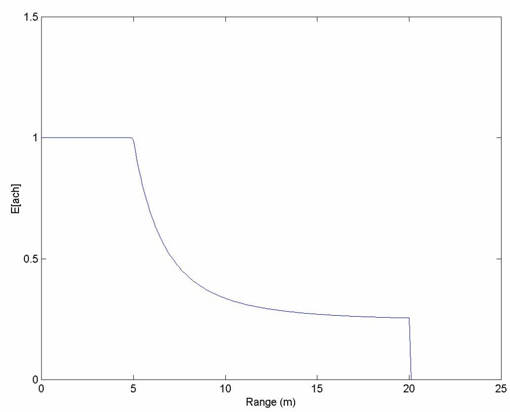

Figure 2 Graph of $E[a_{ch}^n]$ versus range from sensor.

An example of the increase certainty information source can be found in Figure 1(d). Note how the increase certainty values around the rocks and wall at the edges of the scene are high. Also note that the rock in the middle of the scene yields low increase certainty scores because it is already well known as can be seen in the height certainty map shown in Figure 1(b). Thus the increase certainty information source encourages the robot explorer to see the world with greater care, particularly in regions that could be dangerous for the robot.

The reachability information source rewards the robot for going to places which will most strongly impact its knowledge of reachability. It is calculated as:

$$g_r^m = \frac{\sum_{\forall n \in W_m} (-a_r^n \log(a_r^n) - (1 - a_r^n)\log(1 - a_r^n))}{\#W_m} \quad (5)$$

Since a_r^n is the probability that cell *n* is reachable, the numerator is the sum of the entropy over the sensor footprint [12]. This rewards the robot for viewing areas where the reachability is most uncertain (high entropy).

Using equation (1), the total expected information gain in a cell is:

$$E[I] = \frac{1}{14}g_f + \frac{12}{14}g_c + \frac{1}{14}g_r \quad (6)$$

The weight of g_c is much larger than the other two because equation (4) produces very small values. The high weight is used to bring it in line with the values computed from equations (3) and (5) and was chosen empirically. If the three weights in equation (6) were equal, then the increase certainty information source would be ignored during the planning.

5 Simulation Results

The create traversability map exploration task described in section 4 was tested in a simulation environment. The world to be explored is a 2.5D world with discrete obstacles protruding at various heights above a flat ground plane. The world is 150mx150m with 0.5m map cells and has rocks scattered about and a large crater. The environment can be seen in Figure 3(a).

The simulation assumes that the explorer robot is omnidirectional (no dependency on heading) and that it is equipped with a 3D laser sensor. The laser sensor has an azimuth field of view of 360 degrees and an elevation field of view of 10 degrees above horizontal and 80 degrees below. The simulated laser takes a reading every 1 degree in both the azimuth and elevation. The range is 23m.

The results of a simulation run can be found in Figure 3(b). The robot started in the lower left corner with no information about the world to explore. The blue line indicates the path taken by the robot and the green crosses are where it took sensor readings. In general the robot was going to its frontier to take new sensor readings. This is evident in the spacing between successive sensor readings. However the increase map certainty information source caused the robot to deviate and take those sensor readings near obstacles. Further, it is inter-

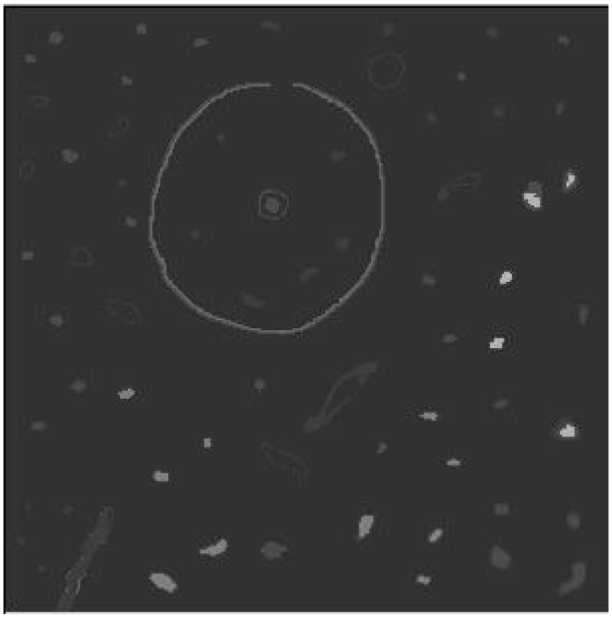

(a) Actual Environment

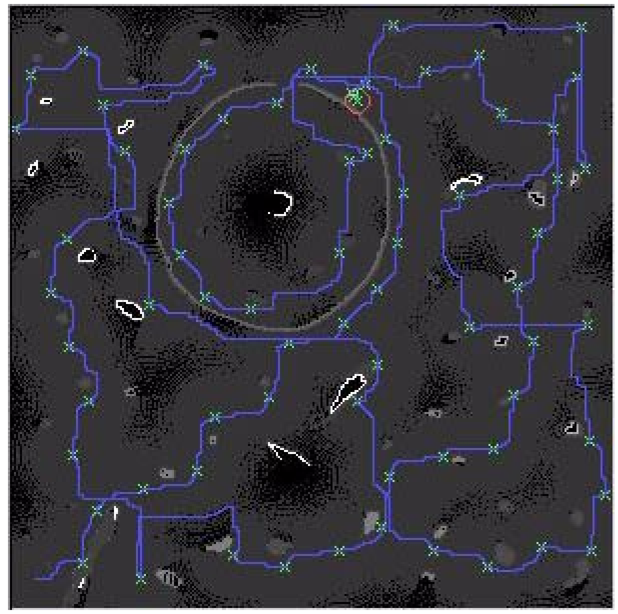

(b) Exploration Results

Figure 3 (a) Map of world to explore with a crater and rocks. The whiter a pixel is the higher its height. (b) Results of the exploration. The blue line is the path taken by the exploring robot and the green crosses are where sensor readings were performed.

esting to see how the robot followed the right edge of the crater until it found to way inside at the top. This behavior was due in part to the increase map certainty, because the crater wall was untraversable, but also to the determine reachability information source. The wall of the crater was low enough that the robot could see the traversable terrain on the other side. The reachability of this terrain was very uncertain and so the determine unreachable information source tried to increase this certainty by eventually getting inside the crater.

6 Summary

This paper presents a methodology for exploring that considers multiple sources of information. The method casts the exploration problem as one of collecting information and allows the explorer to consider more than one source of information. By selecting the appropriate information sources this method can be applied to many different exploration tasks.

The approach presented here is particularly applicable to complex exploration tasks such as planetary exploration. In these tasks the robot must balance competing criteria (information sources) with the its limited resources and try to maximize the information return of its mission.

The exploration method was applied to the task of creating a traversability map of a previously unknown region. Simulation results of this task are presented and show that the robot is able to balance the needs of discovering new terrain, viewing obstacles more closely and determining which cells are reachable.

Currently we are working on the information sources to explore a world with a cliff or a canyon in such a way that the robot tries to see the cliff face or canyon wall. This exploration task is inspired by evidence on some Martian cliffs of past water flow down the face.

Acknowledgements

This work was performed at the Field Robotics Center, Carnegie Mellon University and is supported by NASA under the grant NAG9-1256.

References

[1] Awerbuch, B., Y. Azar, A. Blum, S. Vempala, “New Approximation Guarantees for Minimum-Weight k-Trees and Prize-Collecting Salesmen”, *SIAM Journal of Computing*, 28(1), 1999.

[2] Deans, M., S. Moorehead, B. Shamah, K. Shillcutt, W. Whittaker, “A Concept for Robotic Lunar South Pole Exploration”, *Space ‘98*, 1998.

[3] Elfes, A., “Robot Navigation: Integrating Perception, Environmental Constraints and Task Execution Within a Probabilistic Framework”, *Reasoning with Uncertainty in Robotics*, L. Dorst, M. van Lambalgen and F. Voorbraak (eds), Amsterdam, 1995.

[4] Elfes A., “Dynamic Control of Robot Perception Using Stochastic Models”, *Information Processing in Mobile Robots*, G. Schmidt (ed.), Springer Verlag, 1991.

[5] Estlin, T., A. Gray, T. Mann, G. Rabideau, R. Castano, S. Chien, E. Mjolsness, “An Integrated System for Multi-Rover Scientific Exploration”, *Proc. of the National Conf. on Artificial Intelligence (AAAI-99)*, 1999.

[6] Foley, J., A. vanDam, S. Feiner, J. Hughes, *Computer Graphics Principles and Practice*, Addison-Wesley, 1990.

[7] Kelly, A., “An Intelligent Predictive Control Approach to the High-Speed Cross-Country Autonomous Navigation Problem”, *Ph.D. Thesis*, Carnegie Mellon University, Tech Rep. CMU-RI-TR-95-33, 1995.

[8] Latombe, J.C., *Robot Motion Planning*, Kluwer Academic Publishers, 1991.

[9] Moorehead, S.J., R. Simmons, D. Apostolopoulos, W. Whittaker, “Autonomous Navigation Field Results of a Planetary Analog Robot in Antarctica”, *Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, Noordwijk, Netherlands, 1999.

[10] Moorehead, S.J., R. Simmons, W.L. Whittaker, “Autonomous Exploration Using Multiple Sources of Information”, *Proc. of the 2001 IEEE Int. Conf. on Robotics and Automation*, Seoul, South Korea, 2001.

[11] Nozette, S., C.L. Lichtenburg, P. Spudis, R. Bonner, W. Ort, E. Malaret, M. Robinson, E.M. Shoemaker, “The Clementine Bistatic Radar Experiment”, *Science*, v. 274, pp. 1495-1498, 1996.

[12] Reza, F.M., *An Introduction to Information Theory,* Dover Publications, 1994.

[13] Roy, N., W. Burgard, D. Fox, S. Thrun, “Coastal Navigation - Mobile Robot Navigation with Uncertainty in Dynamic Environments”, *Proc. of the 1999 IEEE Int. Conf. on Robotics and Automation*, Detroit, MI., 1999.

[14] Simmons, R., D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, H. Younes, “Coordination for Multi-Robot Exploration and Mapping”, *Proc. of the National Conf. on Artificial Intelligence*, Austin TX, August 2000.

[15] Yamauchi, B., “A Frontier-Based Approach for Autonomous Exploration”, *Proc. of CIRA ‘97*, Monterrey, CA., July 1997.

[16] Yamauchi, B., A. Schultz and W. Adams, “Mobile Robot Exploration and Map-Building with Continuous Localization”, *Proc. of the 1998 IEEE Int. Conf. on Robotics and Automation*, Leuven, Belgium, May 1998.