



# An Architecture for Distributed Environment Sensing with Application to Robotic Cliff Exploration

VIVEK A. SUJAN AND STEVEN DUBOWSKY

*Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*  
vasujan@alum.mit.edu

TERRY HUNTSBERGER, HRAND AGHAZARIAN, YANG CHENG AND PAUL SCHENKER

*Mobility System Concept Development Section, Jet Propulsion Lab, California Institute of Technology, Pasadena, CA 91109, USA*

**Abstract.** Future planetary exploration missions will use cooperative robots to explore and sample rough terrain. To succeed robots will need to cooperatively acquire and share data. Here a cooperative multi-agent sensing architecture is presented and applied to the mapping of a cliff surface. This algorithm efficiently repositions the systems' sensing agents using an information theoretic approach and fuses sensory information using physical models to yield a geometrically consistent environment map. This map is then distributed among the agents using an information based relevant data reduction scheme. Experimental results for cliff face mapping using the JPL Sample Return Rover (SRR) are presented. The method is shown to significantly improve mapping efficiency over conventional methods.

**Keywords:** cooperative robots, visual exploration, information theory, data fusion, robot communication

## 1. Introduction

To date planetary robots missions have been limited to moving over rather benign terrain (Schenker, 1998). These systems are not capable of exploring highly irregular terrain such as cliff surfaces that are potentially geologically rich and hence very interesting for planetary science (Baumgartner, 1998; Huntsberger, 2000). To succeed robot teams working cooperatively to acquire and share data have been proposed (Pirjanian, 2001; Schenker, 2001; Sujan, 2002; Trebi-Ollenu, 2002; Huntsberger, 2001). Here an efficient cooperative multi-agent algorithm for the visual exploration of unknown environments is proposed. This algorithm repositions the systems' sensors using an information theoretic approach and fuses available sensory information from the agents using physical models to yield a geometrically consistent environment map while minimizing the motions of the robots over the hazardous surfaces. This map is distributed among the agents using

an information based relevant data reduction scheme. Thus, the experiences (measurements) of each robot can become a part of the collective experience of the multi-agent team.

The algorithm has been applied in this study to a team of four robots to cooperatively explore a cliff surface. Figure 1 shows schematically four cooperative robots working in an unstructured field environment to lower one robot down a cliff face that is not accessible by a single robot alone. One robot (Cliff-bot) is lowered down a cliff face on tethers. Two robots (Anchorbots) act as anchor points for the tethers. A fourth robot, RECON-bot (REmote Cliff Observer and Navigator) provides mobile sensing. All the robots are equipped with a limited sensor suite, computational power and communication bandwidths. The Cliff-bot, usually the lightest system, may be equipped with primarily a science sensor suite, and limited sensors for navigation. The RECON-bot, serves to observe the environment to be traversed by the Cliff-bot and communicates the data

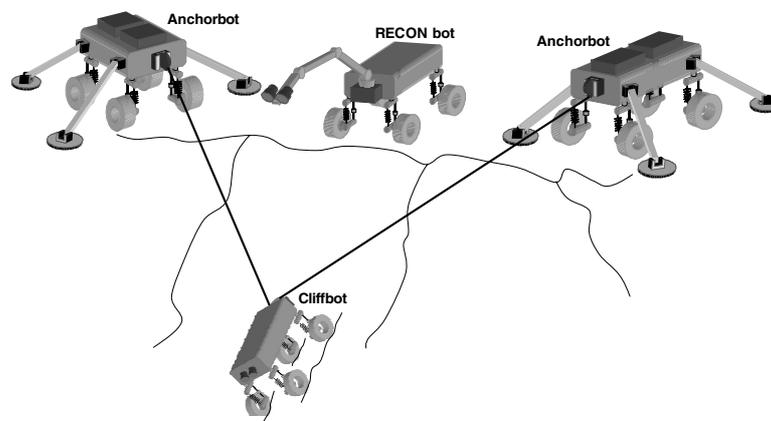


Figure 1. Schematic for a cooperative robot cliff descent.

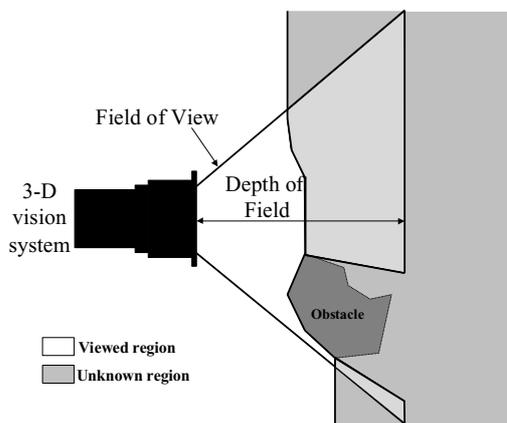


Figure 2. Sensing limitations due to occlusions.

relevant for navigation to the Cliff-bot. The RECON-bot has an independently mobile camera and other on-board sensors to map and observe the environment. Rocks, outcroppings, other robots, etc. limit sensing and sensor placement resulting in uncertainties and occlusions (see Fig. 2). There is significant uncertainty the robots' locations and poses with respect to the environment. Due to these limitations and uncertainties it is difficult or impossible for all robots to independently measure the environment to control the system.

Environment mapping by mobile robots falls into the category of Simultaneous Localization and Mapping (SLAM). In SLAM a robot is localizing itself as it maps the environment. Researchers have addressed this problem for well-structured (indoor) environments and have obtained important results (Anousaki,

1999; Asada, 1990; Burschka, 1997; Castellanos, 1998; Choset, 2001; Kruse, 1996; Kuipers, 1991; Leonard, 1991; Thrun, 2000; Tomatis, 2001; Victorino, 2000; Yamauchi, 1998). These algorithms have been implemented for several different sensing methods, such as camera vision systems (Castellanos, 1998; Hager, 1997; Park, 1999), laser range sensors (Tomatis, 2001; Yamauchi, 1998), and ultrasonic sensors (Anousaki, 1999; Choset, 2001; Leonard, 1991). Sensor movement/placement is usually done sequentially (raster scan type approach), by following topological graphs or using a variety of *greedy* algorithms that explore regions only on the extreme edges of the known environment (Anousaki, 1999; Choset, 2001; Kuipers, 1991; Leonard, 1991; Rekleitis, 2000; Victorino, 2000; Yamauchi, 1998). Geometric descriptions of the environment is modeled in several ways, including generalized cones, graph models and voronoi diagrams, occupancy grid models, segment models, vertex models, convex polygon models (Choset, 2001; Kuipers, 1991). The focus of these works is accurate mapping. They do not address mapping efficiency. Researchers have addressed mapping efficiency to a limited amount (Kruse, 1996). However, sensing and motion uncertainties are not accounted for. They also generally assume that the environment is effectively flat (e.g. the floor of an office or a corridor) and readily traversable (i.e. obstacles always have a route around them) (Anousaki, 1999; Thrun, 2000; Choset, 2001; Kuipers, 1991; Lumelsky, 1989; Yamauchi, 1998) and have not been applied to robot teams working in rough planetary environments. Also, prior work has not addressed optimizing the communication between agents for both multi-agent planning and cooperative map-building.

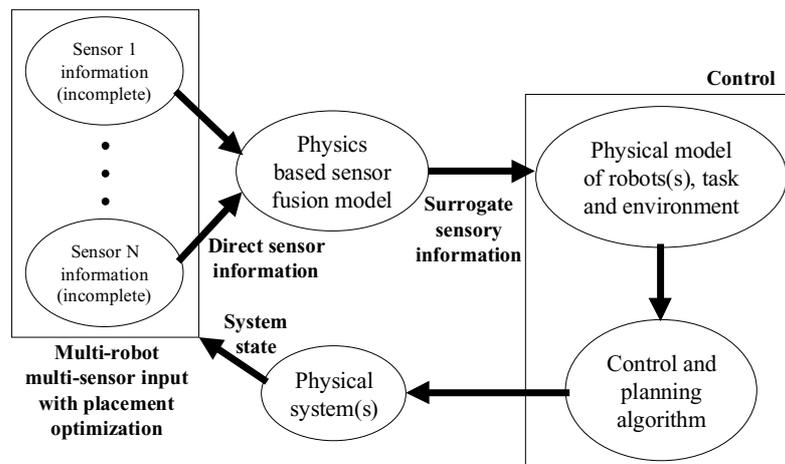


Figure 3. An architecture for multi-robot cooperative sensing.

To achieve the localization function landmarks and their relative motions are monitored with respect to the vision systems. Several localization schemes have been implemented, including topological methods such as generalized voronoi graphs and global topological maps (Choset, 2001; Kuipers, 1991; Tomatis, 2001; Victorino, 2000), extended Kalman filters (Anousaki, 1999; Leonard, 1991; Park, 1999), and robust averages (Park, 1999). Although novel natural landmark selection methods have been proposed (Hager, 1997; Simhon, 1998; Yeh, 1995), most SLAM architectures rely on identifying landmarks as corners or edges in the environment (Anousaki, 1999; Kuipers, 1991; Castellanos, 1998; Victorino, 2000; Choset, 2001; Leonard, 1991). This often limits the algorithms to structured indoor-type environments. Others have used human intervention to identify landmarks (Thrun, 2000).

This paper presents a cooperative multi-agent algorithm for the visual exploration of an unknown environment. The basic approach to the algorithm is shown in Fig. 3 (Sujan, 2002). This algorithm fuses sensory information from one or multiple agents using physical sensor models, robot models, and environment maps to yield geometrically consistent surrogate information in lieu of missing data due to the environment, task, robot and sensor uncertainties. The algorithm falls into the general category of SLAM. The mapping and localization process is as follows. First, each agent efficiently repositions its sensors using an information theoretic approach so as to optimally fill in uncertain/unknown regions of the environment map, based on maximiz-

ing the expected new information obtained. Next, each agent fuses the data to its known environment model by localizing itself with respect to a global fixed reference frame. Finally, each agent shares its known environment map with its team which is then integrated by the other agents into their own environment maps. The information is then used by the control and planning architecture to plan further movements of the sensors for each agent. Thus, a common environment map is built by fusing the data available from the individual robots, providing improved accuracy and knowledge of regions not visible by all robots. The algorithm is unique in that it utilizes the quantity of information of the environment that it currently has to predict high information yielding viewpoints from which to continue exploring the environment. This results in a significantly more efficient exploration process. This algorithm is general in that it is directly applicable to a team of robots as it is to a single explorer. This is achieved through a unique information sharing and fusion architecture.

In this paper the algorithm is applied to a cliff surface exploration robot team as described above. In this application the sensing agent is the JPL Sample Return Rover (SRR) which optimally surveys the cliff surface and transmits the information to other agents. Experimental results compare the Model and Information Theory based Sensing And Fusion ExploreR (MIT-SAFER) architecture to conventional raster (or sequential) sensing schemes. The algorithm details are developed and results presented below.

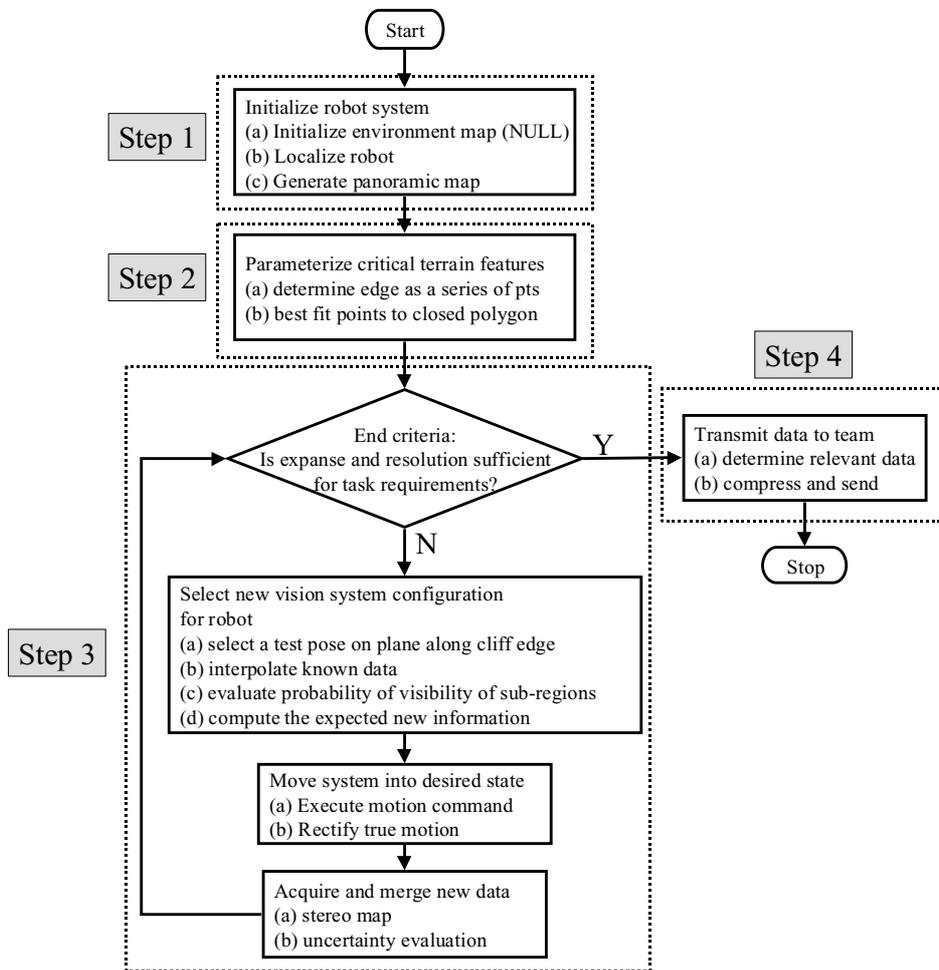


Figure 4. Multi robot environment sensing and distribution flow diagram.

## 2. MIT-SAFER Analytical Development

In general for each sensing agent the algorithm consists of four steps (See Fig. 4).

*Step 1. System initialization:* Here the environment map is initialized, the robots are localized, and a first map is generated. The environment is mapped to a 2.5D elevation grid, i.e., the map is a plane of grid cells where each grid cell value represents the average elevation of the environment at that cell location. This map is built in a fixed reference frame defined by a well-known landmark measurable by all the sensing agents. All robots contributing to or requiring use of the map are localized with respect to the initial map. For the cliff exploration team, the RECON-bot contributes to and uses the en-

vironment map, while the Cliff-bot only uses the environment map. Localization may be achieved by either:

- (a) Absolute localization—is achieved by mapping a common environment landmark that is visible by all robots or
- (b) Relative localization—is done by mapping fiducials on all robots by other robot team members where one robot is selected as the origin. Relative localization is used in this application, with the RECON-bot localizing the Cliff-bot with respect to itself (the origin—see Fig. 5). Then, each agent initially senses the environment.

*Step 2. Critical terrain feature identification:* In some applications, certain regions of the terrain may be

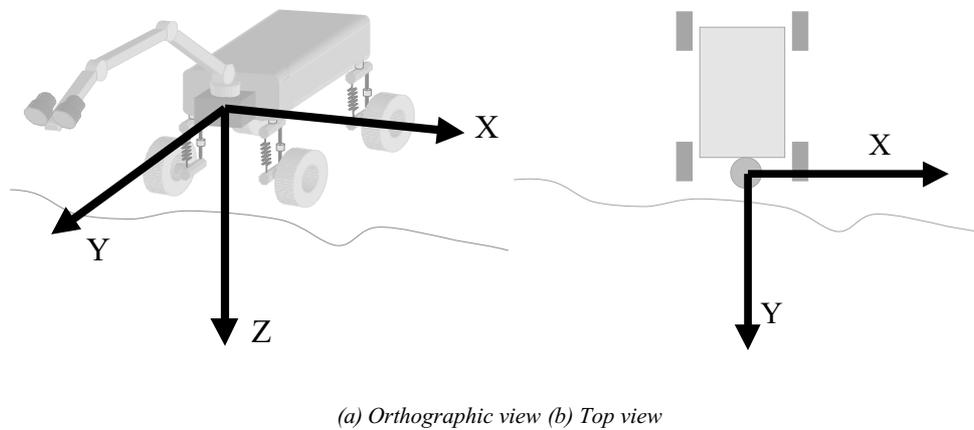


Figure 5. Initial environment map coordinate frame.

critical, requiring early identification and mapping. An example is determining regions of safe travel for the sensing agents.

In this application, identification of the cliff edge by the RECON-bot is critical. The edge is parameterized by the edge of a best-fit non-convex polygon of the local terrain. This permits the RECON-bot to move along the cliff edge without falling over it. In cliff edge parameterization, the surface currently in contact with the RECON-bot is identified in the environment model. This surface is then approximated by a best-fit polygon. The tolerance of the fit is limited by the known rover wheel diameter, i.e., fit tolerance = wheel characteristic length/length per pixel. For this process the incomplete environment model is temporarily completed by a Markovian approximation for unknown grid cells. For all unknown points a worst case initial guess is assumed. This value is the lowest elevation value currently in the known model. A nearest measured neighbor average is performed and iterated till convergence. An example of this is shown in Fig. 6.

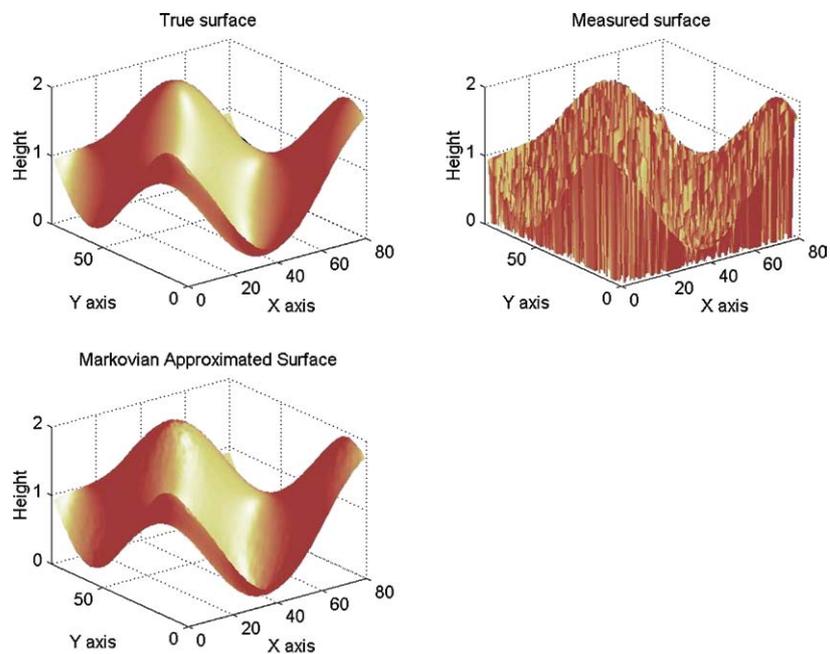
Using the Markovian approximation of the environment, the current rover contact surface (called the plateau) is first identified. This is achieved by setting a height threshold bound to the environment model and projecting the resulting data set onto the XY plane. This is followed by a region growing operation around the current known rover coordinates. Next, the binary image is smoothed by an image closing operation (dilation + erosion). Plateau edge pixels are easily identified at this stage. However, to remove small holes in the plateau, an edge following operation is performed. At this stage we are left with a single closed loop of

boundary pixels. Finally, this set of points is parameterized by a closed polygon. This is initiated by fitting the full set of boundary pixels to a straight line. For any given sub set of boundary pixels that is currently fit to a line, if the error bound on this fit exceeds the prescribed tolerance, then the pixel set is divided into two, and process is repeated. However, before error bound evaluation, line segments fit to each sub set of boundary pixels, are joined to form a closed polygon. The cliff edge parameterization algorithm is outlined in Fig. 7.

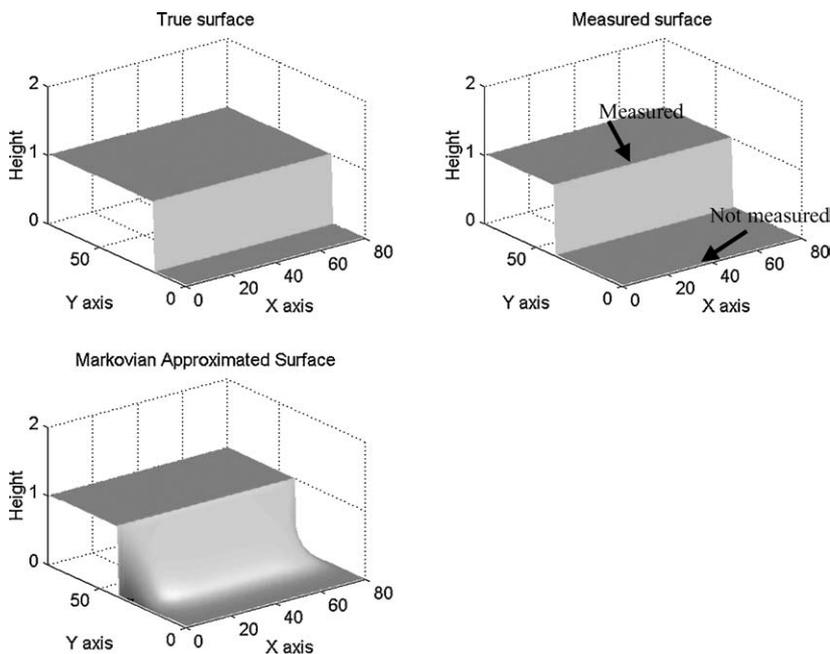
An example of the process is shown in Fig. 8 on a simulated Mars-type environment (note—simulated environment is based on Viking I/II Mars lander rock distribution statistics).

*Step 3. Optimum information gathering pose selection:* A rating function is used to determine the next location (position and orientation) of the sensing agent from which to explore the unknown environment. The objective is to acquire as much new information about the environment as possible with every sensing cycle, while maintaining or improving the map accuracy. Hence, minimizing the exploration time. The process is constrained by selecting goal points that are not occluded and that can be reached by a collision free feasible path.

The new information (NI) is equal to the expected information of the unknown/partially known region viewed from the sensor pose under consideration. In the case of the cliff surface exploration application, the sensors are CCD stereo cameras. This is based on the known obstacles from the current environment map, the field of view of the sensor and



(a) *sample 1-Sinusoidal surface with random unknown cells*



(b) *sample 2-Step surface with half unknown cells*

Figure 6. Markovian interpolation of unknown regions.

a framework for quantifying information. Shannon showed that the information gained by observing a specific event among an ensemble of possible events may be described by the following function (Shannon,

1948):

$$H(q_1, q_2, \dots, q_n) = - \sum_{k=1}^n q_k \log_2 q_k \quad (1)$$

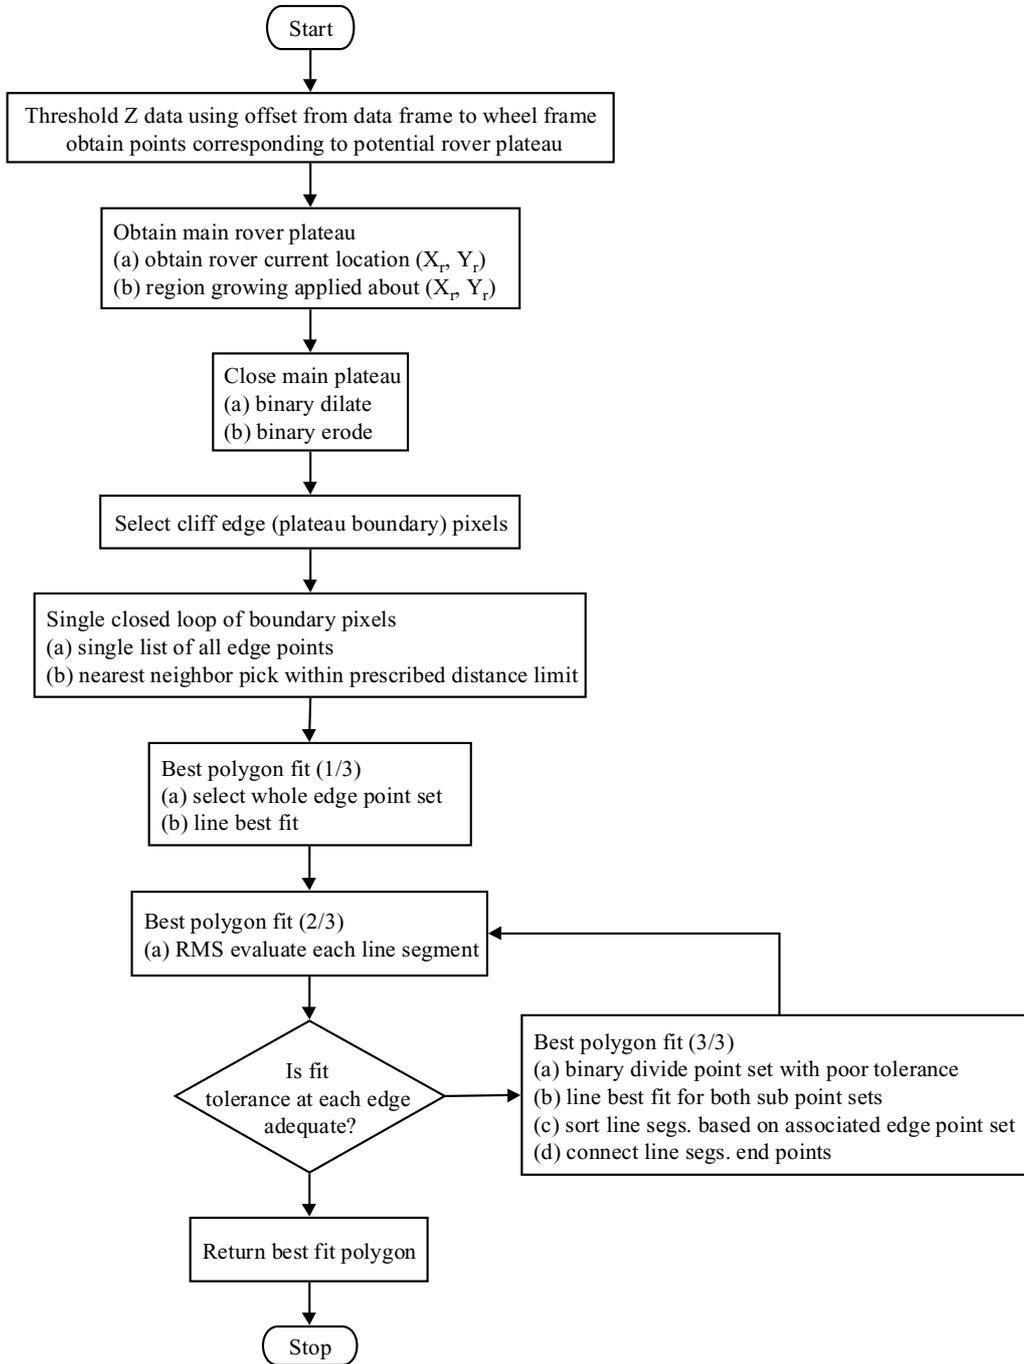


Figure 7. Cliff edge parameterization algorithm flow diagram.

where  $q_k$  represents the probability of occurrence for the  $k$ th event. This definition of information may also be interpreted as the minimum number of states (bits) needed to fully describe a piece of data. Shannon's emphasis was in describing the information content

of 1-D signals. In 2-D the gray level histogram of an ergodic image can be used to define a probability distribution:

$$q_i = f_i/N \quad \text{for } i = 1 \dots N_{\text{gray}} \quad (2)$$

where  $f_i$  is the number of pixels in the image with gray level  $i$ ,  $N$  is the total number of pixels in the image, and  $N_{\text{gray}}$  is the number of possible gray levels. With this definition, the information of an image for which all the  $q_i$  are the same—corresponding to a uniform

gray level distribution or maximum contrast—is a maximum. The less uniform the histogram, the lower the information.

It has been shown that it is possible to extend this idea of information to a 3-D signal (Sujan, 2002). In this

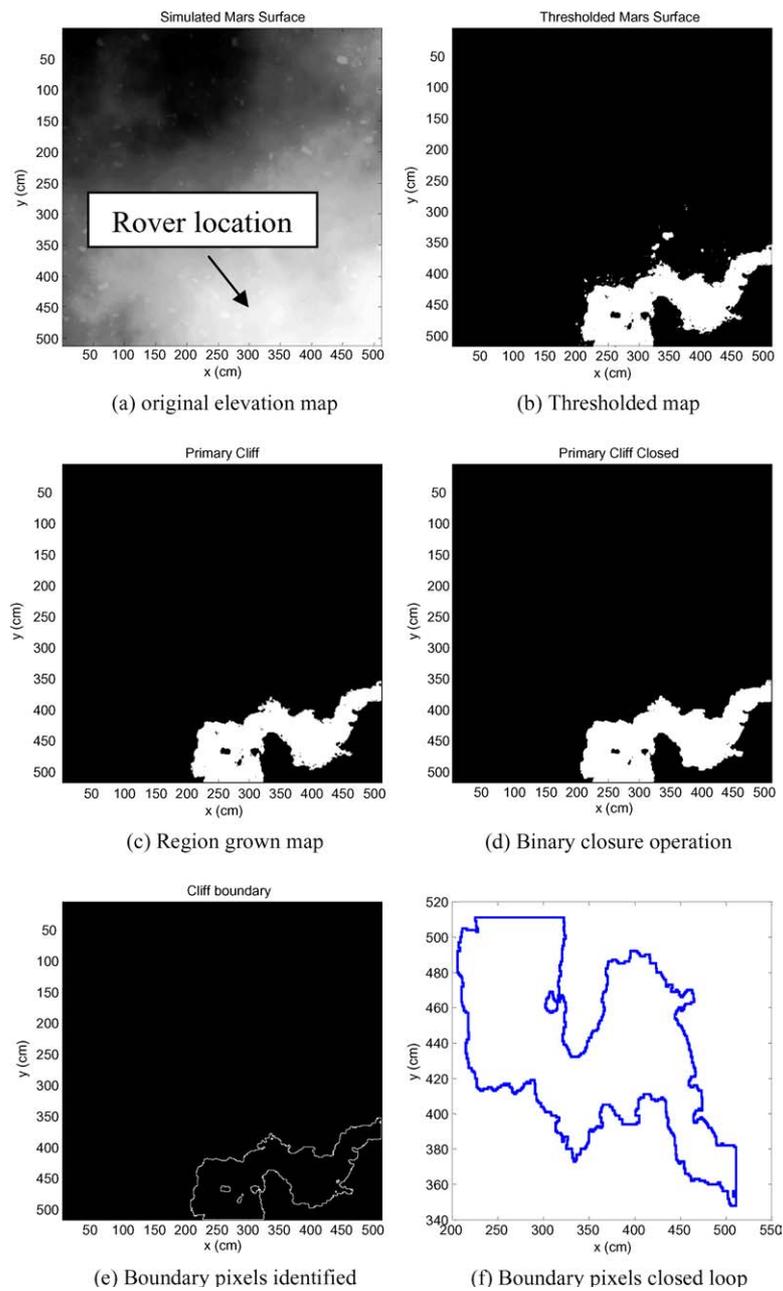


Figure 8. Example of cliff edge parameterization.

(Continued on next page.)

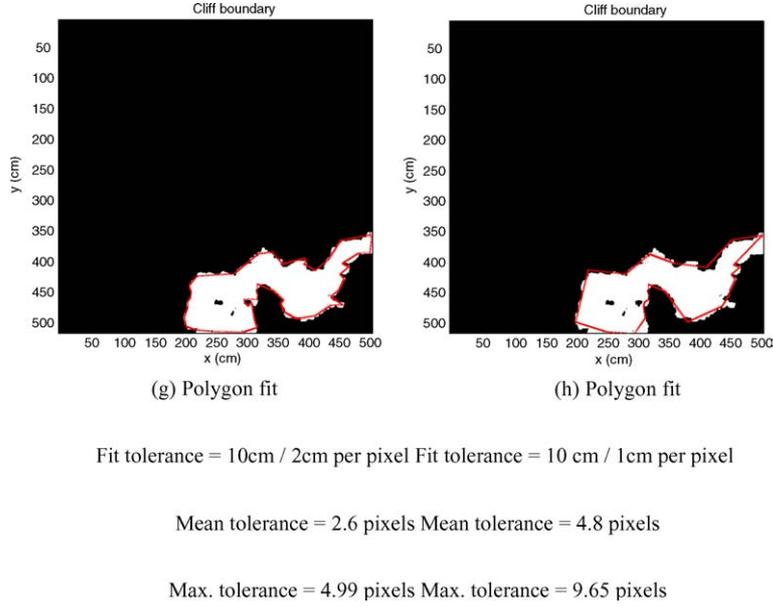


Figure 8. (Continued).

paper this idea is extended to a 2.5D signal—environment elevation map. The new information content for a given sensor (camera) view pose is given by:

$$\begin{aligned}
 H(\text{cam}_{x,y,z,\theta_p,\theta_y}) &= \sum_i \frac{n_{\text{grid}}^{\max} - n_{\text{grid}}^i}{n_{\text{grid}}^{\max}} \left\{ \left( \frac{P_V^i}{2} \log_2 \frac{P_V^i}{2} \right) \right. \\
 &\quad \left. + \left( 1 - \frac{P_V^i}{2} \log_2 \left( 1 - \frac{P_V^i}{2} \right) \right) \right\} \quad (3)
 \end{aligned}$$

where  $H$  is summed over all grid cells,  $i$ , visible from camera pose  $\text{cam}_{x,y,z,\theta_p,\theta_y}$ ;  $n_{\text{grid}}^i$  is the number of environment points measured and mapped to cell  $i$ ;  $n_{\text{grid}}^{\max}$  is the maximum allowable mappings to cell  $i$ ; and  $P_V^i$  is the probability of visibility of cell  $i$  from the camera test pose.

A single range observation of a point ( $\bar{x}$ ) is modeled as a 3-D Gaussian probability distribution centered at  $\bar{x}$ , based on two important observations. First, the use of the mean and covariance of a probability distribution function is a reasonable form to model sensor data and is a second order linear approximation (Smith, 1986). This linear approximation corresponds to the use of a Gaussian (having all higher moments of zero). Second, from the central limit theorem, the sum of a number of independent variables has a Gaussian distribution regardless of their individual distributions. The standard

deviations along the three axes of the distribution correspond to estimates of the uncertainty in the range observation along these axes. These standard deviations are a function of intrinsic sensor parameters (such as camera lens shape accuracy) as well as extrinsic sensor parameters (such as the distance to the observed point or feature). For most range sensing systems, this model can be approximated as (Sujan, 2002):

$$\begin{aligned}
 \sigma_{x,y,z} &= f(\text{extrinsic parameters, intrinsic parameters}) \\
 &\approx S \cdot T_{x,y,z} \cdot L^n \quad (4)
 \end{aligned}$$

where  $S$  is an intrinsic parameter uncertainty constant,  $T_{x,y,z}$  is an extrinsic parameter uncertainty constant,  $L$  is the distance to the feature/environment point, and  $n$  is a constant (typically 2).

$P_V^i$  is evaluated by computing the likelihood of occlusion of a ray  $\text{ray}_{x,y,z}$  using the elevation,  $\text{Ob}_{x,y,z}$ , and the associated uncertainty,  $\sigma_{x,y,z}$ , at all cells lying along this ray path shot through each position in the environment grid to the camera center. From Fig. 9, if grid cell  $i$  falls within the camera field of view, then its average elevation,  $\text{Pt}_{x,y,z}$  (obtained either as an average of all measured points mapped to cell  $i$ , or as the Markovian approximation of its neighborhood if no points have currently been mapped to cell  $i$ ) traces

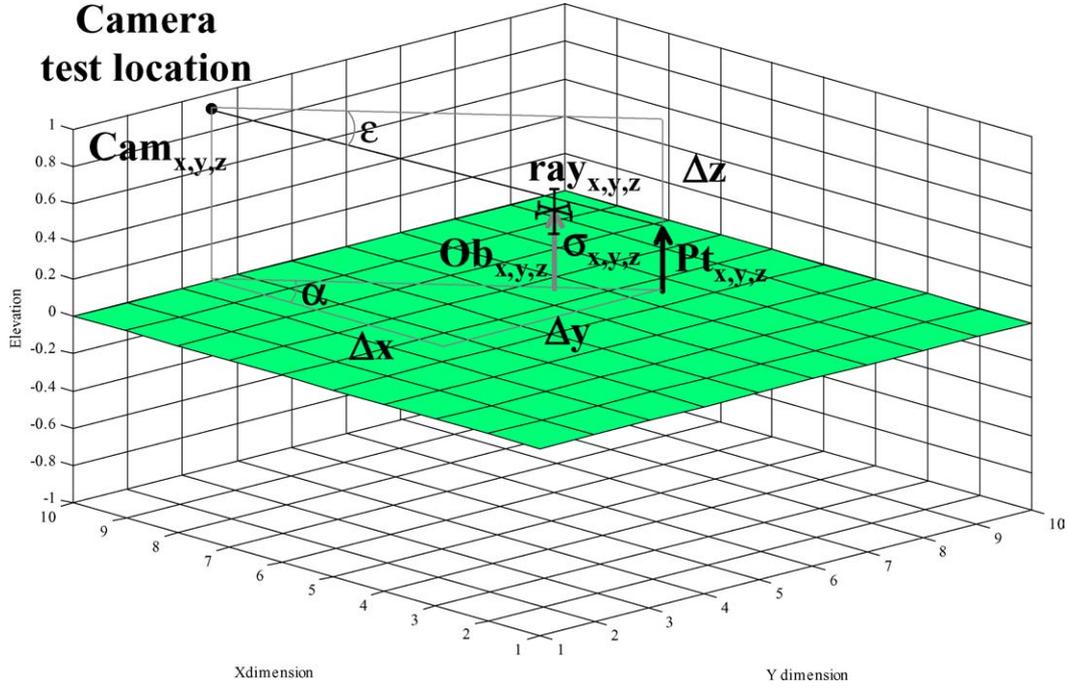


Figure 9. Ray tracing to determine probability of visibility of a grid cell from a given camera configuration.

a ray to the camera center,  $\mathbf{Cam}_{x,y,z}$ .  $P_V^i$  is given by:

$$P_V^i = \prod_{\Delta x} \left\{ \text{sgn}(\mathbf{ray}_z - \mathbf{Ob}_z) \cdot \int_0^{(\mathbf{ray}_z - \mathbf{Ob}_z)} \frac{1}{\sigma_z \sqrt{2\pi}} \exp\left(-\frac{z^2}{2\sigma_z^2}\right) dz + 0.5 \right\} \quad (5)$$

This definition for NI has an intuitively correct form. Regions with higher visibility and associated higher level of unknowns yield a higher expected NI value. Higher occlusions or better known regions result in lower expected NI values.

During the mapping process some regions that are expected to be visible may not be, because of sensor characteristics (e.g., lack of stereo correspondence due to poor textures or lighting conditions), and inaccuracies in the data model (e.g., expected neighboring cell elevations and uncertainties—occlusions). However, after repeated unsuccessful measurements of cells expected to be visible, it becomes more likely that sensor characteristics are the limitation. This is represented as a data quality function that reduces as the number of unsuccessful measurements of the visible cell increases. The probability of visibility of the cell  $i$ ,  $P_V^i$ ,

is pre-multiplied by a “interest function,” I.F., for the cell  $i$  given at the  $k$ th unsuccessful measurement by:

$$\begin{aligned} \text{I.F.}_i^0 &= 1 \\ \text{I.F.}_i^k &= \frac{1}{e^{\beta P_V^i}} \cdot \text{I.F.}_i^{k-1} \end{aligned} \quad (6)$$

where  $\beta$  is a scaling constant determined empirically—larger values result in faster decrease of I.F. Note that cells with low  $P_V^i$  resulting in an unsuccessful measurement are not as severely penalized as cells with high  $P_V^i$ . Hence, occluded regions do not translate to low data quality regions. This permits future “interest” in such regions that may be explored later.

*Data fusion:* A final step in environment map building is to fuse the newly acquired data by each agent with the environment model currently available to that agent. Each agent only fuses its own newly acquired data to the environment map stored in its memory. Thus as the environment map develops on an individual agent level, it needs to be shared and integrated among the team to keep each agent updated. Optimal map sharing protocols for multi agent systems is currently work in progress, i.e., decentralized protocols instructing the team members when and how to share their individual

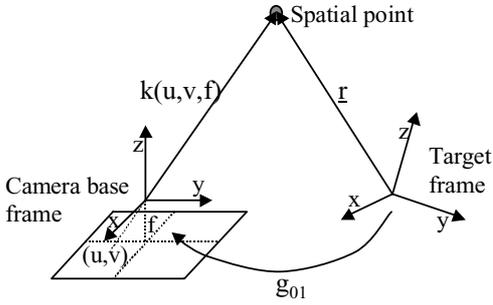


Figure 10. Relationship of camera and target frames.

environment maps. However, once an agent shares its map, the other agents fuse this shared map into their own environment maps using the same method to fuse directly measured data, as described below.

Since the environment model has been developed in a fixed frame (see Step 1), all agents contributing to the environment map require identification of their vision system motion with respect to the fixed coordinate frame, i.e., the agents require global localization. This process eliminates robot positioning errors, such as camera motion errors, and vehicle suspension motions, and allows for accurate data fusion from multiple sources. The process for data fusion is as follows. A single spatial point in the fixed reference frame,  $\bar{r}_i$ , is related to the image point  $(u_i, v_i)$  in the sensor frame by the  $4 \times 4$  transformation matrix  $\mathbf{g}_{01}$  (see Fig. 10). Spatial points are selected and tracked based on a Forstner interest operator and a homography transform (Huntsberger, 2001).

For motion calibration of a camera  $\mathbf{g}_{01}$  needs to be identified:

$$\begin{bmatrix} k_i u_i \\ k_i v_i \\ k_i f \\ 1 \end{bmatrix} = \mathbf{g}_{01} \cdot \bar{r}_i = \begin{bmatrix} [\mathbf{R}_{01}]_{3 \times 3} & \bar{X}_{3 \times 1} \\ \bar{0} & 1 \end{bmatrix} \cdot \begin{bmatrix} r_i^x \\ r_i^y \\ r_i^z \\ 1 \end{bmatrix} \quad (7)$$

where  $\mathbf{R}_{01}$  is the rotational matrix,  $\bar{X}$  is the translation vector,  $f$  is the camera focal length, and  $k_i$  is a scaling constant. For computational reasons it is more convenient to treat the 9 rotational components of  $\mathbf{R}_{01}$  as independent, rather than a transcendental relation of 3 independent parameters. Each spatial point gives three algebraic equations, but also introduces a new variable,  $k_i$ —multiplicative constant to extend the  $i$ th image point vector  $(u, v, f)_i$  to the  $i$ th spatial point in the camera coordinate frame.  $k_i$  may be found from

the disparity pair of the stereo images. For  $n$  points we have:

$$\begin{aligned} u = \mathbf{g}_{01} r &\Rightarrow \begin{bmatrix} k_1 u_1 & k_2 u_2 & \dots & k_n u_n \\ k_1 v_1 & k_2 v_2 & \dots & k_n v_n \\ k_1 f & k_2 f & \dots & k_n f \\ 1 & 1 & \dots & 1 \end{bmatrix} \\ &= \mathbf{g}_{01} \begin{bmatrix} r_1^x & r_2^x & \dots & r_n^x \\ r_1^y & r_2^y & \dots & r_n^y \\ r_1^z & r_2^z & \dots & r_n^z \\ 1 & 1 & \dots & 1 \end{bmatrix} \end{aligned} \quad (8)$$

This set of linear equations can be readily solved using conventional techniques. A least mean square error solution is given by:

$$\mathbf{g}_{01} = \mathbf{u}(\mathbf{r}^T \mathbf{r})^{-1} \mathbf{r}^T \quad (9)$$

The rotation matrix,  $\mathbf{R}_{01}$ , and the translation vector,  $\bar{X}$ , of the camera frame with respect to the base frame are extracted directly from this solution of  $\mathbf{g}_{01}$ . However, for real measured data and associated uncertainty, a larger number of spatial points are required to more correctly identify the geometric transformation matrix,  $\mathbf{g}_{01}$ . Given the  $(i + 1)$ st spatial and image point, from Eq. (9),  $\mathbf{R}_{i+1}$  and  $\bar{X}_{i+1}$  can be obtained. A recursive method is used to determine the mean and covariance of  $\bar{X}$  and  $\mathbf{R}_{01}$  based on the previous  $i$  measurements as follows:

$$\begin{aligned} \hat{X}_{i+1} &= \frac{(i \hat{X}_i + \bar{X}_{i+1})}{i + 1} \\ C_{i+1}^{\bar{X}} &= \frac{i C_i^{\bar{X}} + [\bar{X}_{i+1} - \hat{X}_{i+1}][\bar{X}_{i+1} - \hat{X}_{i+1}]^T}{i + 1} \\ \hat{\mathbf{R}}_{i+1}^{(l,m)} &= \frac{(i \hat{\mathbf{R}}_i^{(l,m)} + \mathbf{R}_{i+1}^{(l,m)})}{i + 1} \\ C_{i+1}^{\mathbf{R}^{(l,m)}} &= \frac{i C_i^{\mathbf{R}^{(l,m)}} + [\mathbf{R}_{i+1}^{(l,m)} - \hat{\mathbf{R}}_{i+1}^{(l,m)}][\mathbf{R}_{i+1}^{(l,m)} - \hat{\mathbf{R}}_{i+1}^{(l,m)}]^T}{i + 1} \end{aligned} \quad (10)$$

This essentially maintains a measure on how certain the camera motion is with respect to its original configuration (assuming the original configuration is known very precisely with respect to the common reference frame). This camera pose uncertainty must be accounted for to obtain an estimate on the position uncertainty of a measured point in the environment. Let

the measurement  $\bar{z}$  be related to the state vector (actual point position)  $\bar{x}$  by a non-linear function,  $h(\bar{x})$ . The measurement vector is corrupted by a sensor noise vector  $\bar{v}$  of known covariance matrix,  $\mathbf{R}$ .

$$\bar{z} = h(\bar{x}) + \bar{v} \quad (11)$$

Assume that the measurement of the state vector  $\bar{x}$  is done multiple times. In terms of the current measurement, a Jacobian matrix of the measurement relationship evaluated at the current state estimate is defined as:

$$H_k = \left. \frac{\partial h(\bar{x})}{\partial \bar{x}} \right|_{\bar{x}=\bar{x}_k} \quad (12)$$

The state (or position) may then be estimated as follows:

$$\begin{aligned} K_k &= P_k H_k^T [H_k P_k H_k^T + R_k]^{-1} \\ \hat{\bar{x}}_{k+1} &= \hat{\bar{x}}_k + K_k [\bar{z}_k - h(\bar{x}_k)] \\ P_{k+1} &= [1 - K_k H_k] P_k \end{aligned} \quad (13)$$

This estimate is known as the Extended Kalman Filter (Gelb, 1974). Using this updated value for both the measured point  $\bar{x}$  and the absolute uncertainty  $P$ , the measured point may then be merged with the current environment model.

Provided two observations are drawn from a normal distribution, the observations can be merged into an improved estimate by multiplying the distributions. Since the result of multiplying two Gaussian distributions is another Gaussian distribution, the operation is symmetric, associative, and can be used to combine any number of distributions in any order. The canonical form of the Gaussian distribution in  $n$  dimensions depends on the standard distributions,  $\sigma_{x,y,z}$ , a covariance matrix ( $C$ ) and the mean ( $\bar{x}$ ) (Stroupe, 2000; Smith, 1986):

$$\begin{aligned} p(\bar{x}' | \bar{y}) &= \frac{1}{(2\pi)^{n/2} \sqrt{|C|}} \exp\left(-\frac{1}{2}(\bar{y} - \bar{x}')^T C^{-1}(\bar{y} - \bar{x}')\right) \end{aligned}$$

where

$$C = \begin{bmatrix} \sigma_x^2 & \rho_{xy}\sigma_{xy}\sigma_{xy} & \rho_{zx}\sigma_{zx}\sigma_{zx} \\ \rho_{xy}\sigma_{xy}\sigma_{xy} & \sigma_y^2 & \rho_{yz}\sigma_{yz}\sigma_{yz} \\ \rho_{zx}\sigma_{zx}\sigma_{zx} & \rho_{yz}\sigma_{yz}\sigma_{yz} & \sigma_z^2 \end{bmatrix} \quad (14)$$

where the exponent is called the Mahalanobis distance. For un-correlated measured data  $\rho = 0$ . The formulation in Eq. (14) is in the spatial coordinate frame. However, all measurements are made in the camera (or sensor) coordinate frame. This problem is addressed through a transformation of parameters from the observation frame to the spatial reference frame as follows:

$$C_{\text{transformed}} = R(-\bar{\theta})^T \cdot C \cdot R(-\bar{\theta}) \quad (15)$$

where  $R(\theta)$  the rotation matrix between the two coordinate frames. The angle of the resulting principal axis can be obtained from the merged covariance matrix:

$$C_{\text{merged}} = C_1(I - C_1(C_1 + C_2)^{-1}) \quad (16)$$

where  $C_i$  is the covariance matrix associated with the  $i$ th measurement. Additionally, a translation operation is applied to the result from Eq. (14), to bring the result into the spatial reference frame. To contribute to the probabilistic occupancy environment model, all measured points corresponding to obstacles are merged. That is, all measured points falling in a particular grid cell, contribute to the error analysis associated with that voxel.

Note that adding noisy measurements leads to a noisier result. For example, the camera pose uncertainty increases as the number of camera steps increase. With every new step, the current uncertainty is merged with the previous uncertainty to get an absolute uncertainty in camera pose. However, by merging (multiplying) redundant measurements (filtering) leads to a less noisier result (e.g., the environment point measurements).

In addition to maximizing information acquisition, it is also desirable to minimize travel distance and maintain/improve the map accuracy, while being constrained to move along feasible paths. A Euclidean metric in configuration space, with individual weights  $\alpha_i$  on each degree of freedom of the camera pose  $\bar{c}$ , is used to define the distance moved by the camera:

$$d = \left( \sum_{i=1}^n \alpha_i (c_i - c'_i)^2 \right)^{1/2} \quad (17)$$

where  $\bar{c}$  and  $\bar{c}'$  are vectors of the new and current camera poses respectively. Here  $\alpha_i$  is set to unity. In general this parameter reflects the ease/difficulty in moving the vision system in the respective axis. Map accuracy is based on the accuracy of localization of each sensing agent. This may be obtained by adding the localization

error of the agent along the path to the target. Paths containing more promising fiducials for localization result in higher utility in determining both the goal location and the path to the goal. The new information, the travel distance and the net improvement of map accuracy is combined into a single utility function that may be optimized to select the next view pose.

*Step 4. Map distribution:* As each agent maps and fuses an environment section to the environment map, it needs to distribute this updated map among the other agents. This is required so that each agent may optimally plan its next move and add information to the map. Once completed, the environment map needs to be distributed to the team. For example, to explore the cliff, after the RECON-bot has developed the geometrical cliff surface map, it needs to transfer this to the Cliff-bot for task execution (e.g., science instrument placement) (Huntsberger, 2001; Huntsberger, 2003).

Due to communication bandwidth limitations of NASA/JPL present and near-term rovers, an appropriate data transfer algorithm needs to be developed. For example, during the 1997 Mars Sojourner mission, both the lander and rover carried 9600 baud radio modems, with an effective data rate of 2400 bps ([http://mars.jpl.nasa.gov/MPF/rover/faqs\\_sojourner.html](http://mars.jpl.nasa.gov/MPF/rover/faqs_sojourner.html)). For the 2003 Mars Exploration Rover (MER) mission the data transfer rates of MER-to-Earth is expected to vary from 3 Kbps to 12 Kbps and MER-to-orbiter is expected to stay constant at 128 Kbps ([http://mars.jpl.nasa.gov/mer/mission/comm\\_data.html](http://mars.jpl.nasa.gov/mer/mission/comm_data.html) <x-html>). These communication limitations may be further exacerbated with multiple cooperating agents. Thus successful communication requires the reduction of the data set into relevant data, i.e., only communicate data that is necessary for task execution.

The data reduction algorithm used here breaks down the environment map into a *quadtree of interest regions*. This is achieved by first reducing the entire elevation map with adaptive decimation. This removes highly insignificant objects, such as small pebbles. The resulting data set is divided into four quadrants. The information content of each quadrant is evaluated using Eqs. (1) and (2). This information content reflects the amount of variation in the terrain quadrant (where higher information content signifies higher variation in the terrain). Quadrants with high information content are further divided into sub-quadrants and the evaluation process is continued. Once it is determined that a quadrant does

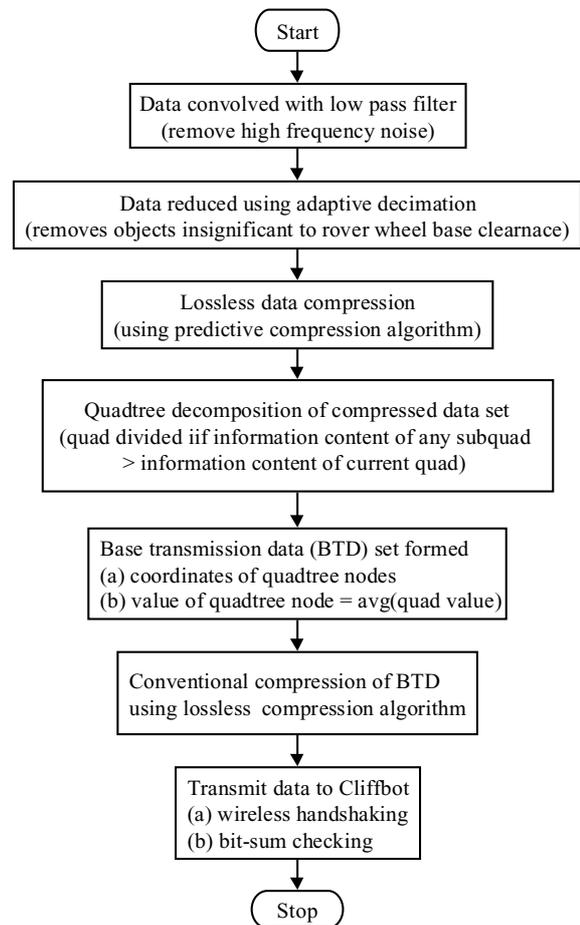


Figure 11. Inter-robot communication flow diagram.

not require further subdivision, an average elevation value of the particular quadrant is used for transmission (rather than the elevation of all grid cells within that quadrant). This cutoff threshold of information is based on a critical robot physical parameter (e.g., the wheel diameter). This results in a significantly reduced data set known as the *quadtree of interest regions*. Conventional lossless compression schemes may then be applied to the reduced data set to further reduce the number of transmission bits. The flow diagram of this process is given in Fig. 11.

### 3. Experimental Results

The basic MIT-SAFER algorithm was applied to the cooperative exploration of cliff surfaces by a team of four robots. The JPL Sample Return Rover (SRR) served

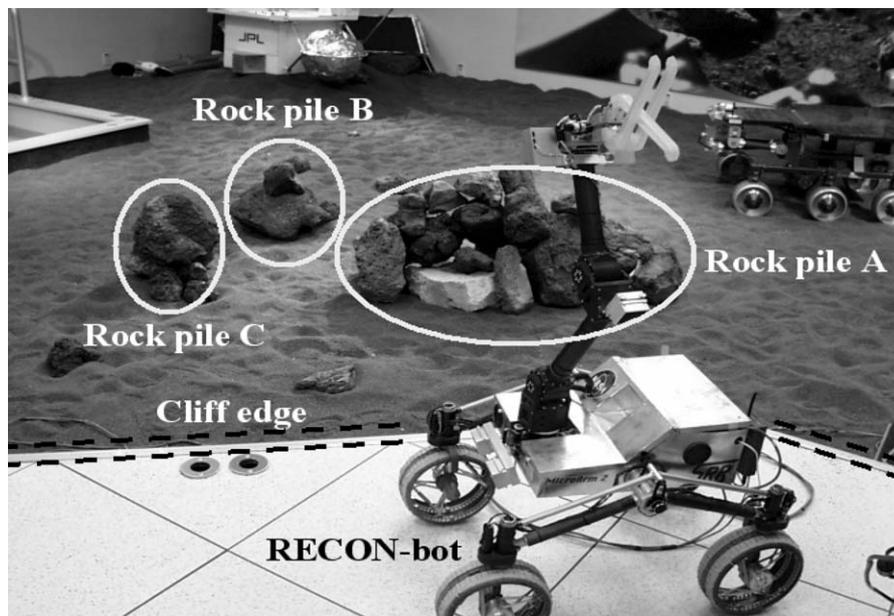


Figure 12. Experimental laboratory setup.

as the RECONbot for this application. The SRR is a four-wheeled mobile robot with independently steered wheels and independently controlled shoulder joints. It carries a stereo pair of cameras mounted on a three DOF articulated manipulator. The SRR is equipped with a 266 MHz PC-104 computer platform, operating with VX-Works. Five mapping techniques, including the one developed above, were implemented. These were: These include:

*Method 1.* Raster scanning without yaw.

*Method 2.* Raster scanning with yaw.

*Method 3.* Information based environment mapping with cliff edge assumed to be a straight line segment.

*Method 4.* Information based environment mapping with cliff edge approximated as a non-convex polygon.

*Method 5.* Information based environment mapping with interest function and cliff edge approximated as a non-convex polygon.

The first two methods reflect commonly used environment mapping schemes (Asada, 1990; Burschka, 1997; Castellanos, 1998; Choset, 2001; Kuipers, 1991; Rekleitis, 2000; Victorino, 2000). The latter three reflect with increased complexity the algorithm developed here.

The experimental setup for the first study in the Planetary Robotics Lab (PRL) at JPL is shown in Fig. 12. A recessed sandpit containing several rock piles is mapped. The edge of the sandpit, a vertical drop, acts as the cliff edge. This limits the motion of the RECON-bot to lie in the flat plane behind the cliff edge (see Fig. 12). Figure 13 shows the number of environment grid cells explored as a function of the number of stereo imaging steps. From this experimental study, the improved efficiency of the method presented in this paper over conventional raster scanning methods can be seen, with an order of magnitude more points being mapped by Method 5 over those returned from Method 1 for the same number of stereo imaging steps. A significant improvement in efficiency can be seen while progressing from Method 3 to Method 5. In Method 4, by parameterizing the cliff edge, the rover is able to follow the edge more aggressively, thus covering a larger variety of view points.

Figure 14 shows a top view of the environment points mapped using Methods 3 and 5. It is seen that Method 5 takes approximately half the number of steps to map a qualitatively similar region. Further, it is observed that the left region of the sandpit in Fig. 12 yields poor data (due to lack of stereo correspondence). Since this region is expected have high information content (due to lack of occlusions), the algorithm in Method 3 tends to converge to view points looking in that direction.

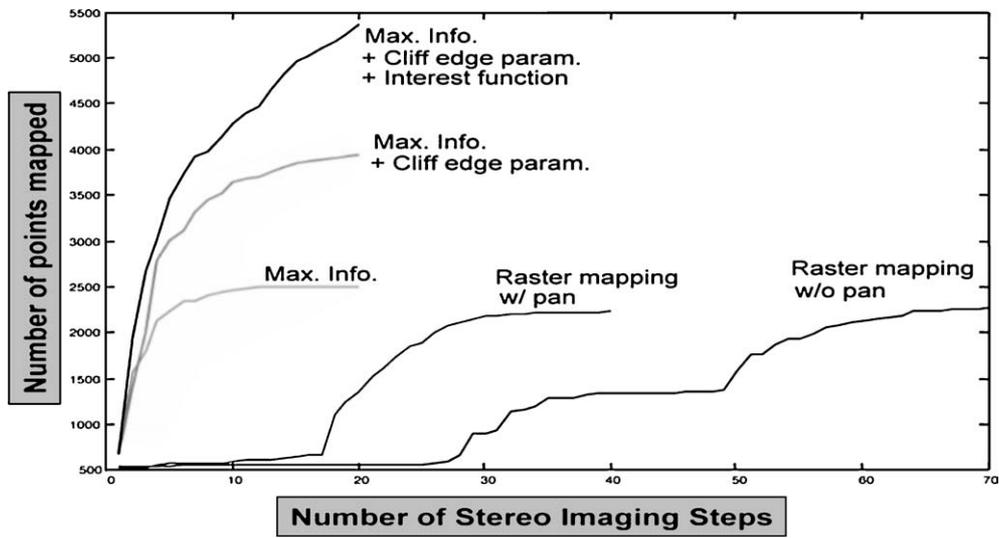


Figure 13. Amount of environment explored.

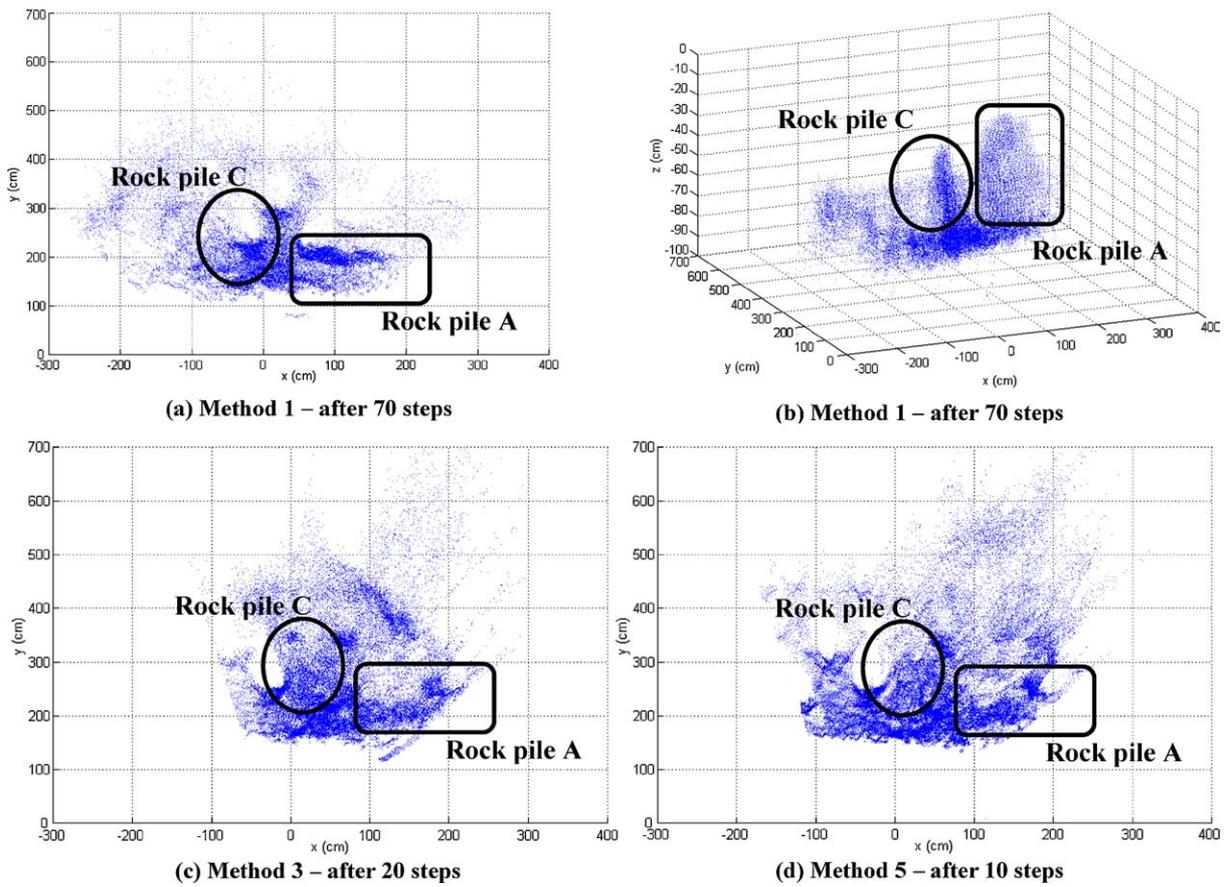


Figure 14. Top view of mapped points.

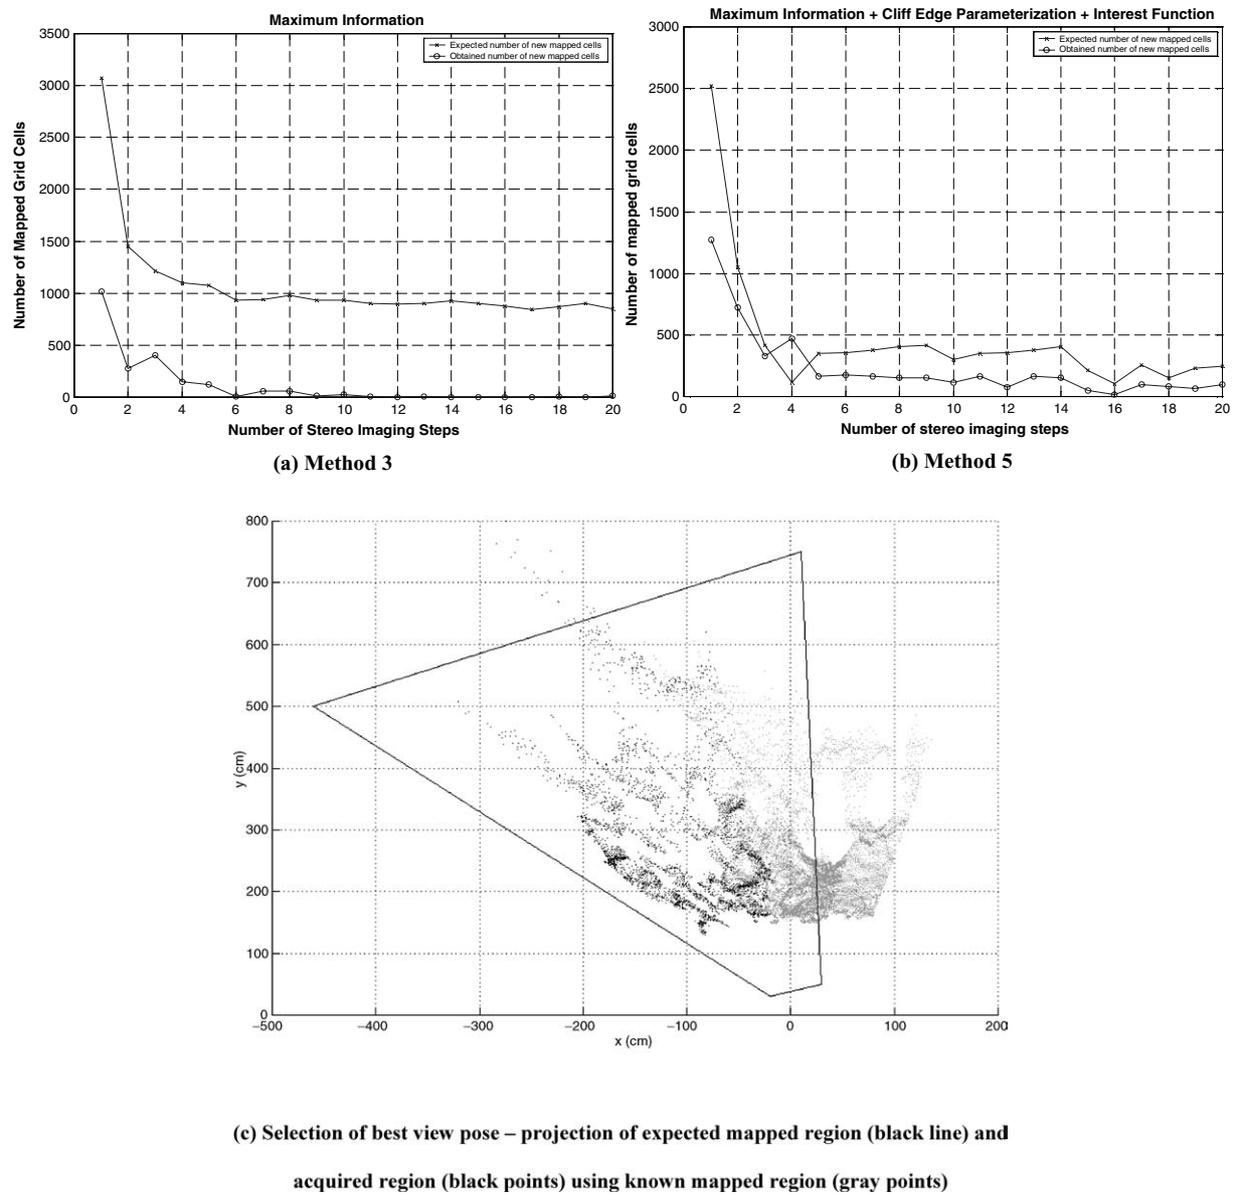


Figure 15. Comparison of the number of expected new mapped cells versus the number obtained.

However, in Method 5, the algorithm concludes that the data quality is poor and eventually loses interest in this region. This is seen in Fig. 15 that shows the number of expected environment grid cell measurements as opposed to the number obtained. In Method 5, there is reasonable agreement. However, in Method 3, while the expected number of measurements is significant, the obtained number of grid cell measurements drops off to zero. Figure 15(c) shows an example of the

projected mapped area as opposed to the true mapped area. Differences exist primarily due to poor imaging (stereo correspondence). However, occlusions and inaccuracies in projected area from local slope variations also contribute to this difference. Figure 16 shows the interest function value obtained in Method 5 for each environment grid cell using Eq. (6). It is seen that regions to the left rapidly lose their interest values with time since they yield low quality data.

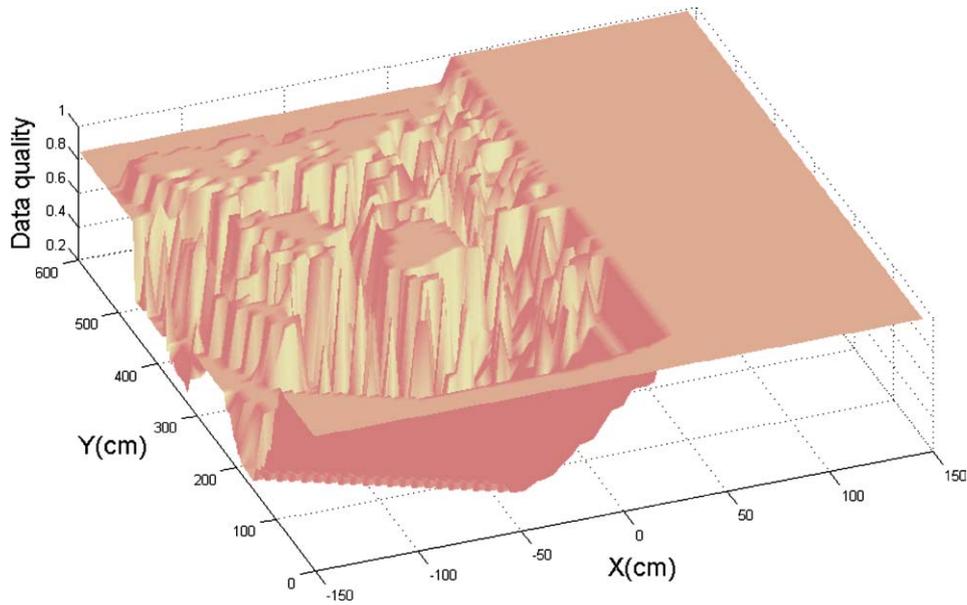


Figure 16. Interest function value after 10 steps using Method 5.

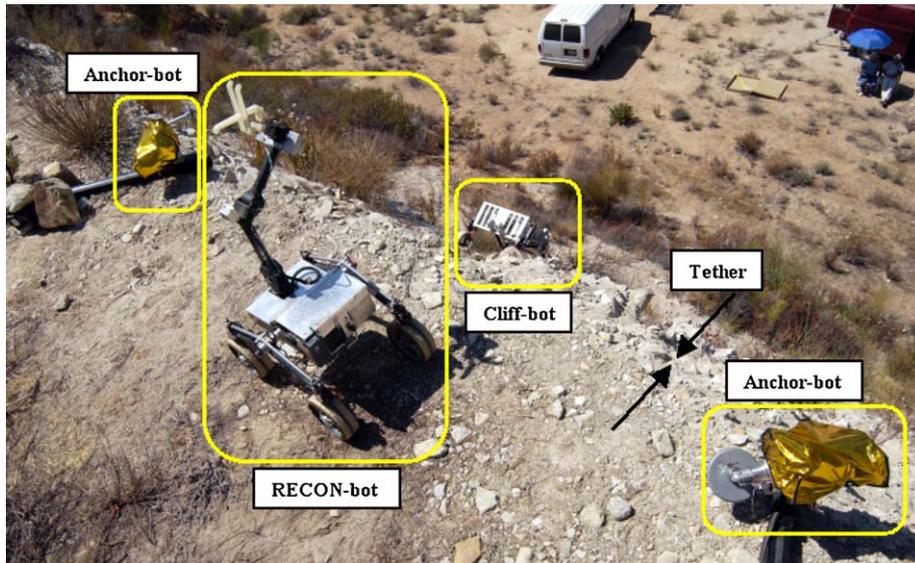


Figure 17. Experimental field system setup.

Field tests were conducted near the Tujunga Dam in Tujunga, CA on a natural cliff face with a vertical slope of  $\sim 75^\circ$ . This setup seen in Fig. 17. This is the physical realization of the conceptual description provided in Section 1 of a team of four cooperating robots exploring a cliff surface. Due to time constraints, ex-

perimental tests could only be run for Method 3 using the maximum information content and Method 5 using the maximum information content with interest function. The results of the study for 10 imaging steps are shown in Fig. 18. Figure 19 shows part of the cliff surface and its corresponding map. Of particular interest

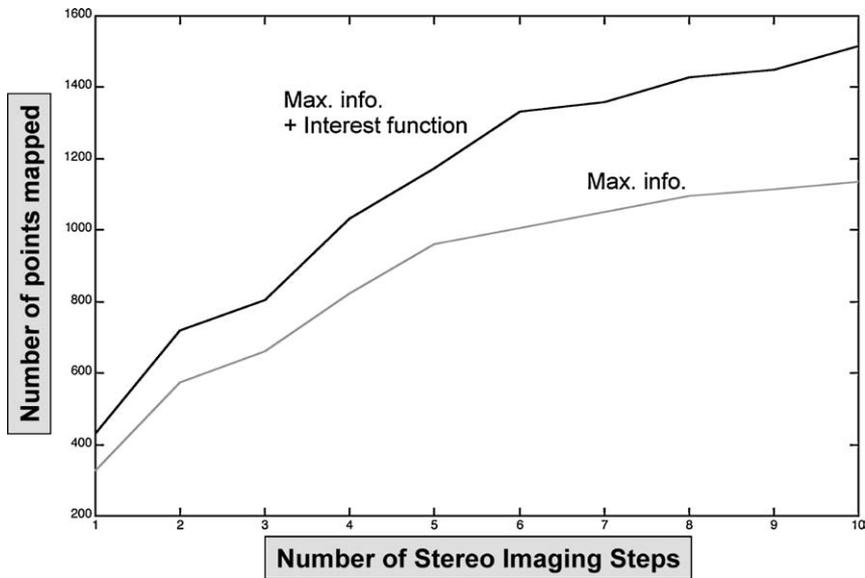
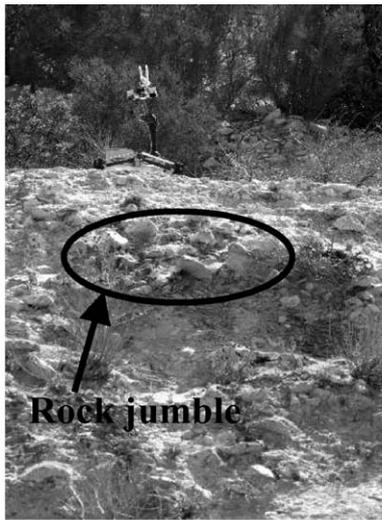
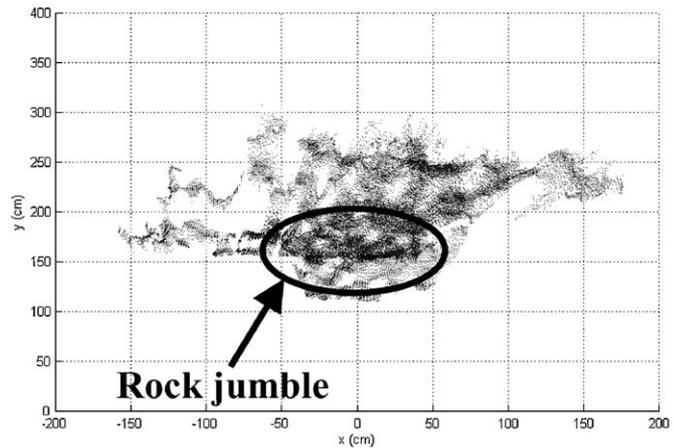


Figure 18. Amount of environment explored.



(a) View of RECONbot mapping the cliff



(b) Overhead view of surface mapped

Figure 19. Tujung dam cliff site.

is the rock jumble to the Cliffbot, which may choose to avoid it during traversal.

Figure 20 compares the number of expected environment grid cell measurements and the number obtained for the two methods. Method 5 shows reasonable agreement, while Method 3 results in a large discrepancy. Once again, differences exist primarily due to poor imaging (stereo correspondence). However, occlusions

and inaccuracies in projected area from local slope variations also contribute to this difference. Finally, Fig. 21 shows the interest function value obtained in Method 5 for each environment grid cell using Eq. (6).

These results demonstrate the effectiveness of the multi-agent environment mapping architecture developed in this paper. To demonstrate the effectiveness of the map reduction and distribution algorithm for robots

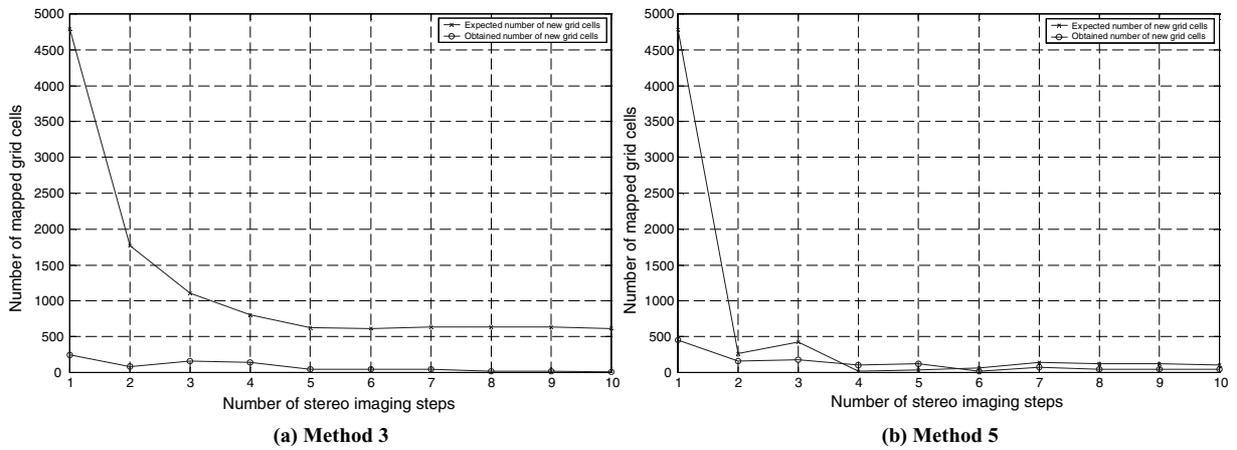


Figure 20. Comparison of the number of expected new mapped cells versus the number obtained.

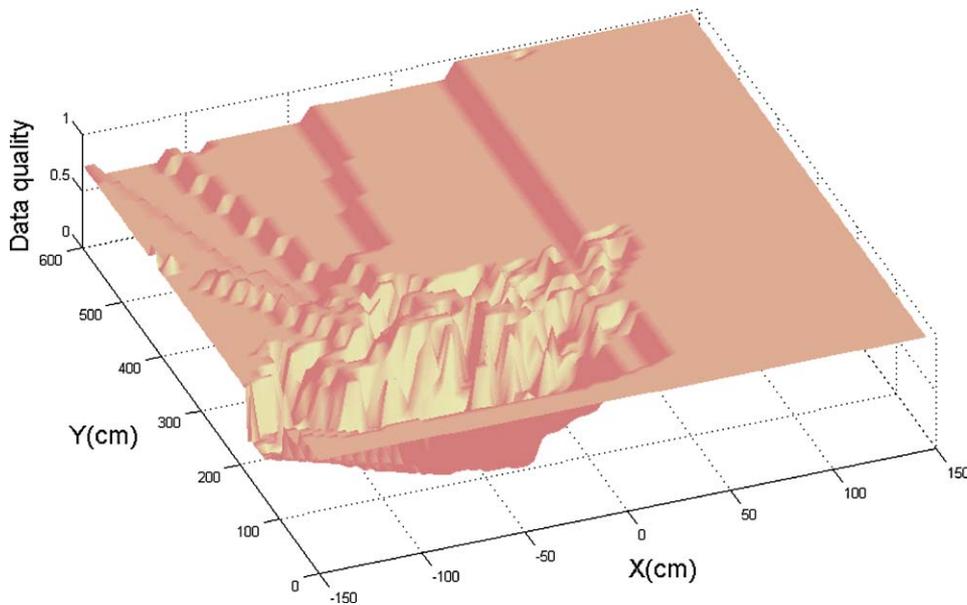


Figure 21. Interest function value after 10 steps using Method 5.

in real Mars field environments, 32 different elevation maps of fixed dimensions, based on the statistics of Viking I/II Mars lander data were tested. The data of each elevation map was reduced with respect to a robot with varying wheel diameter. To compare the data reduction, a terrain variation parameter,  $dH$ , is defined as the terrain elevation variation normalized by the robot wheel diameter. Thus, it is expected that robots with smaller wheel diameters (higher  $dH$ ) require a greater amount of terrain detail for navigation, than those with larger wheel diameters for the same terrain map. Figure 22 confirms this expectation. It shows the

data reduction factor as a function of  $dH$  using the algorithm described above (without conventional lossless compression added at the end). The variation at each data point represents the variation in data reduction expected for a given elevation map.

An example of this data reduction process is shown in Figs. 23–26. It compares the grayscale elevation map before and after the data reduction process—lighter regions indicate higher elevations. Figure 23 shows a contour map of the simulated environment. Figure 24 shows the quad-tree decomposition of the environment. Note that higher divisions of regions indicate more

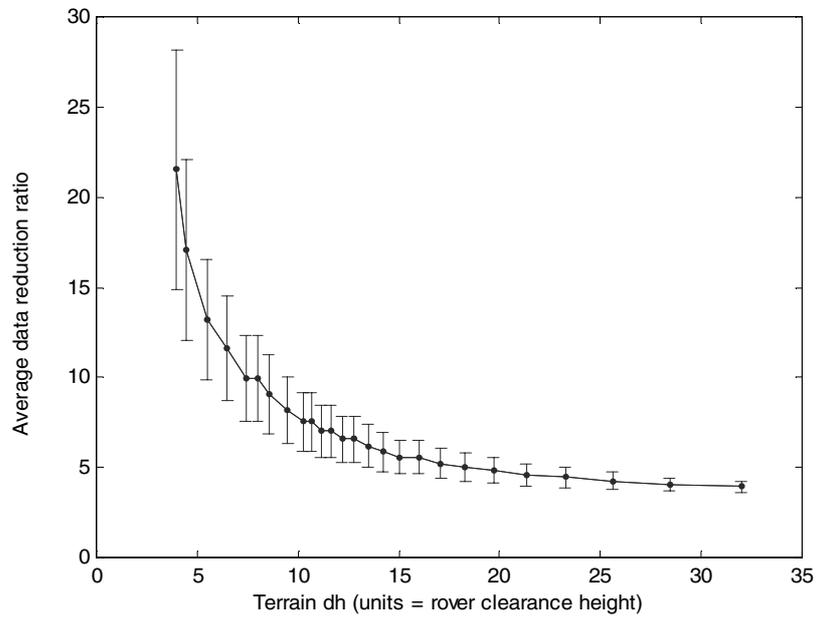


Figure 22. Elevation map data reduction for transmission as a function of terrain elevation range.

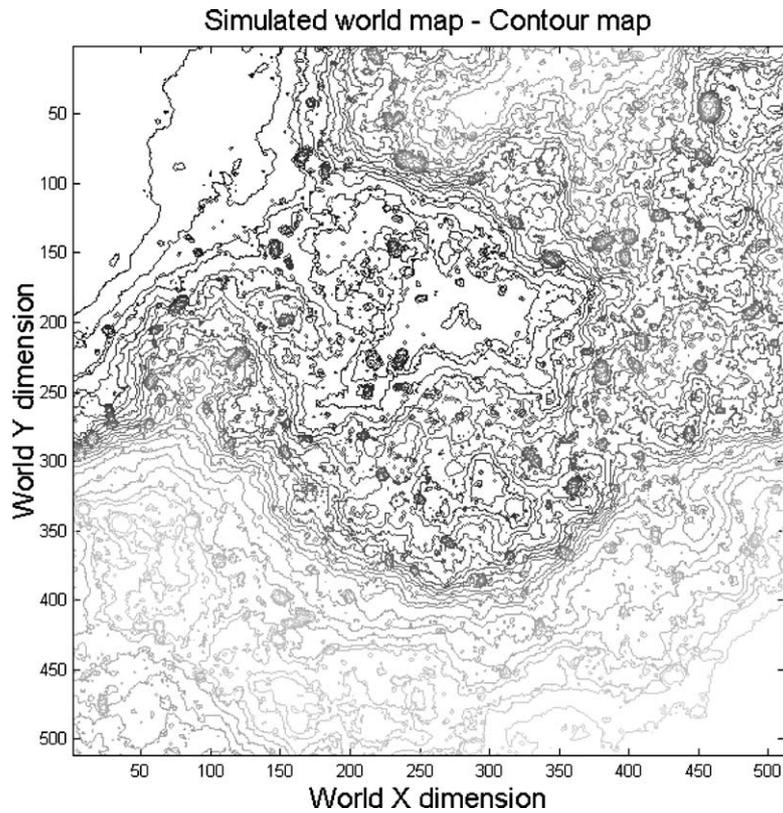


Figure 23. Simulated world contour map—gray color code indicates darker regions to be of lower elevation.

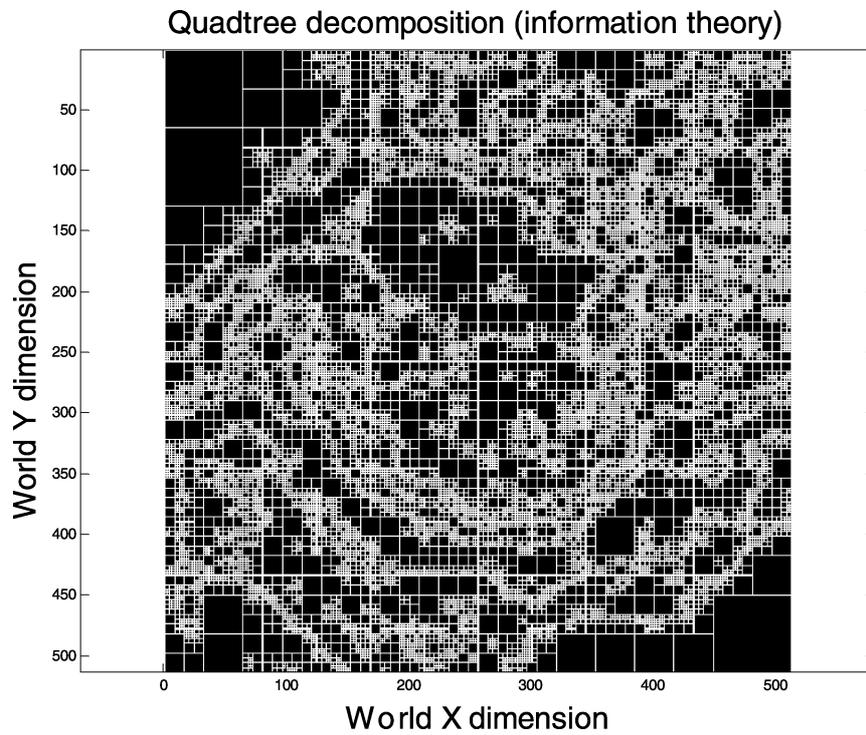


Figure 24. Quadtree decomposition of elevation map.

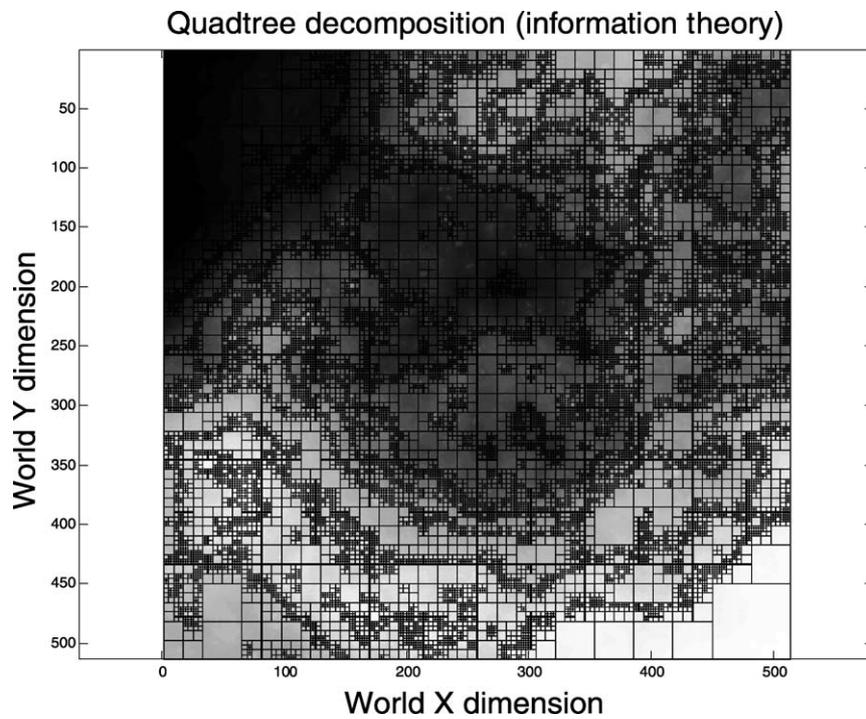


Figure 25. Quadtree decomposition overlaid on original elevation map.

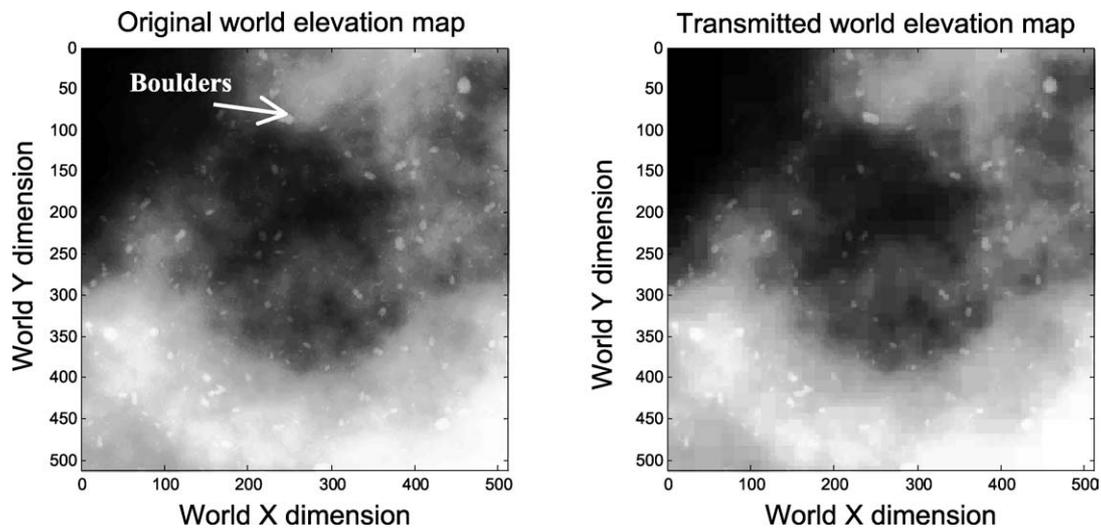


Figure 26. The original (left) and the process/transmitted (right) environment elevation maps.

complex terrain and consequently higher relevant information content. Figure 25 shows the contour map overlaid on the original grayscale elevation map. Finally, Fig. 26 compares the elevation map before and after the data reduction process.

For this example, a data reduction factor of approximately 10 was achieved with a  $dH = 8$ . Although visually the left and right images of Fig. 26 may appear the same, closer inspection reveals regions in the transmitted image (such as the bottom right corner) to contain very little information content. This indicates that the region in the original elevation contained insignificant terrain variation with respect to the particular wheel diameter. However, other regions such as the boulders, indicated in the original elevation map, that are critical with respect to the wheel diameter, are faithfully transmitted. It is seen that using this method, significant data reduction can be achieved while maintaining the overall map structure. Although, this is applied to a 2.5D environment elevation map here, the algorithm is directly applicable to 3D maps.

#### 4. Conclusions

This paper has presented a cooperative multi-agent distributed sensing architecture. This algorithm efficiently repositions the systems' sensors using an information theoretic approach and fuses sensory information from the agents using physical models to yield a geometrically consistent environment map. This map is then

distributed using an information based relevant data reduction scheme for communication. The architecture is proposed for a team of robots cooperatively interacting to explore a cliff face. Experimental results using the JPL Sample Return Rover (SRR) have been presented. This single rover acts as a surveyor, optimally generating a map of the cliff face. The method is shown to significantly improve the environment mapping efficiency. The algorithm shows additional mapping efficiency improvement when an interest function is included. This function measures the data quality in the environment. Future work includes implementation and testing of the inter-robot communication algorithm on the experimental platform.

#### Acknowledgments

The authors would like to acknowledge the support of the Jet Propulsion Laboratory, NASA under contract number 1216342.

#### References

- Anousaki, G.C. and Kyriakopoulos, K.J. 1999. Simultaneous localization and map building for mobile robot navigation. *IEEE Robotics & Automation Magazine*, 6(3):42–53.
- Asada, M. 1990. Map building for a mobile robot from sensory data. *IEEE Transactions on Systems, Man, and Cybernetics*, 37(6).
- Baumgartner, E.T., Schenker, P.S., Leger, C., and Huntsberger, T.L. 1998. Sensor-fused navigation and manipulation from a planetary rover. In *Proceedings of SPIE Symposium on Sensor Fusion*

- and *Decentralized Control in Robotic Systems*, Boston, MA, vol. 3523.
- Burschka, D., Eberst, C., and Robl, C. 1997. Vision based model generation for indoor environments. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, Albuquerque, New Mexico, vol. 3, pp. 1940–1945.
- Castellanos, J.A., Martinez, J.M., Neira, J., and Tardos, J.D. 1998. Simultaneous map building and localization for mobile robots: A multisensor fusion approach. In *Proceedings of 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, vol. 2, pp. 1244–1249.
- Choset, H. and Nagatani, K. 2001. Topological simultaneous localization and mapping (SLAM): Toward exact localization without explicit localization. *IEEE Transactions on Robotics and Automation*, 17(2):125–137.
- Gelb, A. 1974. *Applied Optimal Estimation*. MIT Press: Cambridge, Massachusetts, U.S.A.
- Hager, G.D., Kriegman, D., Teh, E., and Rasmussen, C. 1997. Image-based prediction of landmark features for mobile robot navigation. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, Albuquerque, New Mexico, vol. 2, pp. 1040–1046.
- Huntsberger, T.L., Rodriguez, G., and Schenker, P.S. 2000. Robotics: Challenges for robotic and human Mars exploration. In *Proceedings of ROBOTICS2000*, Albuquerque, New Mexico, pp. 299–305.
- Huntsberger, T.L., Pirjanian, P., and Schenker, P.S. 2001. Robotic outposts as precursors to a manned Mars habitat. In *Proceedings of the 2001 Space Technology and Applications International Forum (STAIF-2001)*, Albuquerque, New Mexico, pp. 46–51.
- Huntsberger, T.L., Suján, V.A., Dubowsky, S., and Schenker, P.S. 2003. Integrated system for sensing and traverse of cliff faces. In *Proceedings of Aerosense'03: Unmanned Ground Vehicle Technology V. SPIE*, Orlando, Florida, vol. 5083.
- Kruse, E., Gutsche, R., and Wahl, F.M. 1996. Efficient, iterative, sensor based 3-D map building using rating functions in configuration space. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, Minneapolis, Minnesota, vol. 2, pp. 1067–1072.
- Kuipers, B. and Byun, Y. 1991. A robot exploration and mapping strategy based on semantic hierarchy of spatial representations. *Journal of Robotics and Autonomous Systems*, 8:47–63.
- Leonard, J.J. and Durrant-Whyte, H.F. 1991. Simultaneous map building and localization for an autonomous mobile robot. In *IEEE 1991 International Workshop on Intelligent Robots and Systems*, Osaka, Japan, vol. 3, pp. 1442–1447.
- Lumelsky, V., Mukhopadhyay, S., and Sun, K. 1989. Sensor-based terrain acquisition: The “sightseer” strategy. In *Proceedings of the 28th Conference on Decision and Control*, Tampa, Florida, vol. 2, pp. 1157–1161.
- Osborn, J.F., Whittaker, W.L., and Coppersmith, S. 1987. Prospects for robotics in hazardous waste management. In *Proceedings of the Second International Conference on New Frontiers for Hazardous Waste Management*, Pittsburgh, Pennsylvania.
- Park J., Jiang, B., and Neumann, U. 1999. Vision-based pose computation: Robust and accurate augmented reality tracking. In *Proceedings of the 2nd IEEE/ACM International Workshop on Augmented Reality*, San Francisco, California, pp. 3–12.
- Pirjanian, P., Huntsberger, T.L., and Schenker, P.S. 2001. Development of CAMPOUT and its further applications to planetary rover operations: A multirobot control architecture. In *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems IV*, Newton, Massachusetts, vol. 4571, pp. 108–119.
- Rekleitis, I., Dudek, G., and Milios, E. 2000. Multi-robot collaboration for robust exploration. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, California, vol. 4, pp. 3164–3169.
- Schenker, P.S., Baumgartner, E.T., Lindemann, R.A., Aghazarian, H., Ganino, A.J., Hickey, G.S., Zhu, D.Q., Matthies, L.H., Hoffman, B.H., and Huntsberger, T.L. 1998. New planetary rovers for long-range Mars science and sample return. In *Proceedings of SPIE Symposium on Intelligent Robots and Computer Vision XVII: Algorithms, Techniques, and Active Vision*, Boston, Massachusetts, vol. 3522.
- Schenker, P.S., Pirjanian, P., Huntsberger, T.L., Trebi-Ollennu, A., Aghazarian, H., Leger, C. JPL, Dubowsky, S. MIT, McKee, G.T., University of Reading (UK). 2001. Robotic intelligence for space: Planetary surface exploration, task-driven robotic adaptation, and multirobot cooperation. In *Proceedings of SPIE Symposium on Intelligent Robots and Computer Vision XX: Algorithms, Techniques, and Active Vision*, Newton, Massachusetts, vol. 4572.
- Shaffer, G. and Stentz, A. 1992. A robotic system for underground coal mining. In *Proceedings of 1992 IEEE International Conference on Robotics and Automation*, Nice, France, vol. 1, pp. 633–638.
- Shannon, C.E. 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423.
- Simhon, S. and Dudek, G. 1998. Selecting targets for local reference frames. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, vol. 4, pp. 2840–2845.
- Smith, R.C. and Cheeseman, P. 1986. On the representation and estimation of spatial uncertainty. *International Journal of Robotics Research*, 5(4):56–68.
- Stroupe, A.W., Martin, M., and Balch, T. 2000. Merging gaussian distributions for object localization in multi-robot systems. In *Proceedings of the Seventh International Symposium on Experimental Robotics, ISER '00*, Hawaii.
- Sujan, V.A. and Dubowsky, S. 2002. Visually built task models for robot teams in unstructured environments. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, Washington, DC, vol. 2, pp. 1782–1787.
- Thrun, S., Burgard, W., and Fox, D. 2000. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3D mapping. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, California, vol. 1, pp. 321–328.
- Tomatis, N., Nourbakhsh, I., and Siegwar, R. 2001. Simultaneous localization and map building: A global topological model with local metric maps. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, Maui, Hawaii, vol. 1, pp. 421–426.
- Trebi-Ollennu, A., Das, H., Aghazarian, H., Ganino, A., Pirjanian, P., Huntsberger, T., and Schenker, P. 2002. Mars rover pair cooperatively transporting a long payload. In *Proceedings of 2002 IEEE International Conference on Robotics and Automation*, Washington, DC, vol. 4, pp. 3136–3141.
- Victorino, A.C., Rives, P., and Borrelly, J.-J. 2000. Localization and map building using a sensor-based control strategy. In *Proceedings*

*of 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Takamatsu, Japan, vol. 2, pp. 937–942.

Yamauchi, B., Schultz, A., and Adams, W. 1998. Mobile robot exploration and map-building with continuous localization. In *Proceedings of 1998 IEEE International Conference on Robotics and Automation*, Leuven, Belgium, vol. 4, pp. 3715–3720.

Yeh, E. and Kriegman, D.J. 1995. Toward selecting and recognizing natural landmarks. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Pittsburgh, Pennsylvania, vol. 1, pp. 47–53.



**Vivek Sujan** received his B.S. in Physics and Mathematics *summa cum laude* from the Ohio Wesleyan University in 1996, his B.S. in Mechanical Engineering, with honors from the California Institute of Technology, 1996, and his M.S. and Ph.D. in Mechanical Engineering from the Massachusetts Institute of Technology in 1998 and 2002 respectively. He is currently a Post-Doctoral Associate in Mechanical Engineering at M.I.T. Research interests include the design, dynamics and control of electromechanical systems; mobile robots and manipulator systems; optical systems and sensor fusion for robot control; image analysis and processing. He has been elected to Sigma Xi (Research Honorary Society), Tau Beta Pi (Engineering Honorary Society), Phi Beta Kappa (Senior Honorary Society), Sigma Pi Sigma (Physics Honorary Society), Pi Mu Epsilon (Math Honorary Society).



**Steven Dubowsky** received his Bachelor's degree from Rensselaer Polytechnic Institute of Troy, New York in 1963, and his M.S. and Sc.D. degrees from Columbia University in 1964 and 1971. He is currently a Professor of Mechanical Engineering at M.I.T. He has been a Professor of Engineering and Applied Science at the University of California, Los Angeles, a Visiting Professor at Cambridge University, Cambridge, England, and Visiting Professor at the California Institute of Technology. During the period from 1963 to 1971, he was employed by the Perkin-Elmer Corporation, the General Dynamics Corporation, and the American Electric Power Service Corporation. Dr. Dubowsky's research has included the de-

velopment of modeling techniques for manipulator flexibility and the development of optimal and self-learning adaptive control procedures for rigid and flexible robotic manipulators. He has authored or co-authored nearly 200 papers in the area of the dynamics, control and design of high performance mechanical and electromechanical systems. Professor Dubowsky is a registered Professional Engineer in the State of California and has served as an advisor to the National Science Foundation, the National Academy of Science/Engineering, the Department of Energy, and the US Army. He has been elected a fellow of the ASME and IEEE and is a member of Sigma Xi, and Tau Beta Pi.



**Terry Huntsberger** is a Senior Member of the Technical Staff in the Mechanical and Robotics Technologies Group at NASA's Jet Propulsion Laboratory in Pasadena, CA, where he is the Manager for numerous tasks in the areas of multi-robot control systems, rover systems for access to high risk terrain, and long range traverse algorithms for rovers. He is a member of the planning teams for the Mars Science Laboratory (MSL) mission planned for launch in 2009. He is an Adjunct Professor and former Director of the Intelligent Systems Laboratory in the Department of Computer Science at the University of South Carolina. His research interests include behavior-based control, computer vision, neural networks, wavelets, and bio-logically inspired system design. Dr. Huntsberger has published over 100 technical articles in these and associated areas. He received his PhD in Physics in 1978 from the University of South Carolina. He is a member of SPIE, ACM, IEEE Computer Society, and INNS.



**Hrand Aghazarian** is a Senior Member of the Engineering Staff in the Mechanical and Robotics Technologies Group at the Jet Propulsion Laboratory in Pasadena, CA. He has extensive work experience in the design and development of real-time embedded software for a variety of avionics systems. Currently he is involved in providing software solutions in the area of software architecture, low-level drivers, motor control, user interfaces for commanding planetary

rovers, and navigation algorithm for SRR and FIDO rovers. His research interests are Rover Navigation/Control and Real-Time Embedded software Design and Development. He received a dual B.S. degree in Applied Math and Computer Science and a M.S. degree in Computer Science from California State University in Northridge. He is a member of ACM and INNS.



**Yang Cheng** is a senior research staff member in the Machine Vision Group of the Jet Propulsion Laboratory. He earned his Ph.D. in remote sensing from the Geography Department of the University of South Carolina and was then a staff member at Oak Ridge National Laboratory. Since coming to JPL in 1999, he has worked on several space robotic research projects, including Landmark-Based Small Body Navigation System, the Vision and Navigation Subsystem for the FIDO rover, and Passive Imaging-Based Spacecraft Safe Landing. Yang's research interests include robotic navigational autonomy, computer vision, remote sensing, cartography, map projection, geographic information systems, parallel computing, etc.



**Paul S. Schenker** is Manager, Mobility Systems Concept Development Section, Jet Propulsion Laboratory (JPL), California Institute of Technology, Pasadena, and as such, responsible for JPL R&D in planetary mobility and robotics. His published research includes topics in robotic perception, robot control architectures, telerobotics and teleoperation, multi-sensor fusion, and most recently, multi-robot cooperation. He has led the development of robotic systems that include the Field Integrated Design & Operations Rover (FIDO), Planetary Dexterous Manipulator (MarsArm, microArm), Robot Assisted Microsurgery System (RAMS), Robotic Work Crew (RWC), and All Terrain Explorer (ATE/Cliff-bot), with resulting technology contributions to NASA missions. Dr. Schenker is active in the IEEE, Optical Society of America, and SPIE. He has served as an elected Board member and 1999 President of the last; he currently serves as an elected member of the NAS/United States Advisory Committee to the International Commission for Optics.