

Task Directed Imaging in Unstructured Environments by Cooperating Robots

Vivek A. Sujan
Department of Mechanical Engineering
Massachusetts Institute of Technology
Cambridge, MA 02139
{vasujan@mit.edu} 617-253-5095

Abstract

In field environments, due to model and sensor uncertainties, it is not usually possible to provide robotic systems with optimum sensing strategies for their tasks. The robot or robot teams will need to utilize available models and sensory data to find task based optimum sensing poses. Here, an algorithm based on iterative sensor planning and sensor redundancy is proposed to enable them to efficiently position their cameras with respect to the task/target. Simulations show the effectiveness of this algorithm.

1. Introduction

An important goal of robotics research is to develop mobile robot teams that can work cooperatively in unstructured field environments, such as shown conceptually in Figure 1 [2, 8]. Potential tasks include explosive ordinance removal, de-mining and hazardous waste handling, exploration/development of space, environment restoration, and construction [2, 8, 14].




Figure 1: Representative physical system

The planning and control of such systems typically requires models of the environment and task. In unstructured field environments it is often not possible to have such a-priori models. The robot needs to construct these from available sensory information [16]. Once these models are created, the robots need to position their sensors in a task directed optimal way. A number of problems can make this non-trivial. These include the uncertainty of the task in the environment, location and orientation uncertainty in the individual robots, and occlusions (due to obstacles, work piece, other robots). If the systems are equipped with cameras mounted on articulated mounts, intelligent planning of the camera

motion can alleviate some of the problems of the occlusions. If the system consists of more than one robot, planning the behavior of these multi-information sharing systems can further improve the system performance.

Previous work in planning of visual sensing strategies can be divided into two areas [13, 17]. One of these is concerned with sensor positioning—placing a sensor so that it can best observe some feature and selecting a sensing operation which will prove the most useful in object identification and localization. Researchers have limited their work to model-based approaches, requiring previously known environments [4, 5, 7, 9]. Target motions (if any) are assumed to be known [11]. Brute force search methods divide the entire view volume into grids, octrees, constraint sets, and search algorithms for optimum sensor location, are applied [6, 7, 13]. Additionally, they require a-priori knowledge of object/target models [5, 17]. Such methods can be effective but are computationally expensive and not practical for many real field environments, where occlusions and measurement uncertainties are present.

The other direction of research in planning of sensors is sensor data fusion—combining complementary data from either different sensors or different sensor poses to get an improved net measurement [15, 17]. The main advantage of multi-sensor fusion is the exploitation of data redundancy and complementary information. For environment and target model building both areas play key roles. However, little work has been done in effectively combining the capabilities of sensor placement planning and sensory fusion, to develop a sensing strategy for model building to be used by robots and robot teams in unstructured environments.




Figure 2: Cooperative sensing by robots

An algorithm for efficient environment and task model building based on an information theoretic approach is developed and described in [16]. Using the environment/task model created by this algorithm, the

individual robots are positioned “optimally” with respect to the target. This process is described in this paper. It is assumed that the system consists of two (or more) mobile robots working in an unknown environment (such as constructing a planetary structure—see Figure 2). The target is static. Each robot has a 3D-vision system mounted on an articulated arm. Sensing and sensor placement is limited, resulting in occlusions and uncertainties. The objective is to efficiently find poses for each camera system to view the target in a task directed optimal way. The algorithm iterates on the known environment model, accounting for object motions, occlusions, and vision system characteristics.

2. Algorithm Description

2.1. Overview

Figure 3 outlines the map building and camera placement algorithm. The algorithm consists of two major parts. In the first part, the articulated cameras cooperatively scan the region around a target generating a 3D geometric model, so that the robots can locate themselves and the obstacles in the target reference frame [16]. The second stage consists of using this model to find an optimum pose for the multiple camera systems to view the target(s). The 3D map is modeled as a probabilistic discretized occupancy grid. Every voxel in the map has a value for probability-of-occupancy that ranges from 0 (empty) to 1 (occupied). A value of 0.5 indicates maximum uncertainty in occupancy of the voxel. The process is initialized by visually finding the target and robots in a common reference frame and repeated for all robots sequentially. Inter-robot collision is avoided through communication of pose solutions.

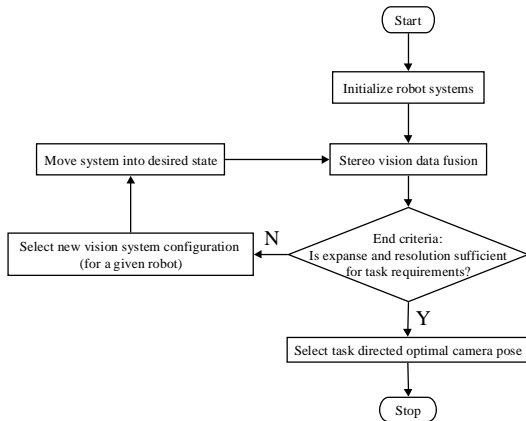


Figure 3: Outline of model building and placement algorithm

2.2. Algorithm initialization

The first step of the model building and camera placement algorithm in an unstructured environment is to identify the targets. This establishes a common reference frame for all

further environment mapping and camera placement operations. For the purposes of this paper simple circular holes in a work-piece were used as targets. These targets are detected using template matching techniques by first filtering the image using a LOG operator. For the purposes of this research, variations due to illumination conditions are neglected.

2.3. Optimum Camera Pose Identification

Given the geometric environment model with its uncertainties [16], an optimum pose for a camera to view a given target can now be developed. Based on the probabilistic geometric world map, (x,y,z) locations with $P_{x,y,z} < 0.05$ are considered as unoccupied. Such points form candidate configuration space camera pose coordinates. A rating function (RF) is defined and optimized over the known configuration space for the new camera pose. This rating function is defined as:

$$RF(x, y, z) = \text{DOF}_{\text{RF}}^{\alpha} \cdot \text{Res}_{\text{RF}}^{\beta} \cdot \text{TFV}_{\text{RF}}^{\gamma} \cdot \text{TAV}_{\text{RF}}^{\delta} \cdot (1 - P_{x,y,z}) \quad (1)$$

where DOF_{RF} , Res_{RF} , TFV_{RF} and TAV_{RF} are the contributions to the rating function due to the depth of field, resolution of the target, target field visibility and target angular visibility respectively, from the camera position (x,y,z) . These contributions are defined in the following sections. $\alpha, \beta, \gamma, \delta$ are weighting constants (set to unity for the experiments conducted in this paper). It is assumed that the camera normal vector points directly to the target. Camera position improves as the rating function value increases.

2.3.1. Depth of field (DOF)

The DOF constraint of a camera system is defined as the maximum and minimum distance from the camera-lens system between which all feature points will be sufficiently in focus. This tolerance is based on the lens aperture effects as well as the flexibility allowed by the image processing algorithms (such as range finding, feature identification, etc.). The value of $\text{DOF}_{\text{RF}}=0$ if the feature point is outside the depth of field, and $\text{DOF}_{\text{RF}}=1$ if the feature point is inside the depth of field. More complex non-binary functions may be used here that quantify how much a given feature point is in focus. However, this simple binary function gives good solutions.

2.3.2. Resolution of target

The resolution of the target along an axis from the given position of the camera is simply:

$$R = \frac{2d \tan(\alpha/2)}{n} \quad \text{and} \quad \text{Res}_{\text{RF}} = \frac{1}{R} \quad (2)$$

where d is the distance of the camera from the target, α is the camera field of view and n is the number of pixels along the detector axis. The contribution to the rating function from the resolution is given by Res_{RF} .

2.3.3. Target field visibility

The target field visibility (TFV_{RF}) at any point in space, is

defined as the largest angle of excursion the camera can traverse around a circle centered about the target, before the target is occluded. The target field visibility must account for the occlusions in the workspace and the finite size of the target. Finding the target field visibility (for finite sized targets) in a 3-D space while accounting for the occlusions can be very difficult and time consuming, growing exponentially with the number of occlusions. To reduce this difficulty, a method of occlusion expansion using convex hulls is proposed. This process consists of three steps:

(i) *Occluding object expansion and target reduction*

The finite sized target is reduced to a single point while the occlusions are expanded. Determining the field visibility of the target from any given point is thus simplified. Figure 4 outlines the idea of object expansion in 2-D. A coordinate frame is attached to the target at a reference point. This point is now placed at every vertex of the occluding object and the target projected accordingly. The new vertices of the target/feature are computed. This new set of vertices forms the expanded object.




Figure 4: Target reduction and occlusion expansion

(ii) *Convex hulls of the expanded objects*

The expanded object is simplified to occlusion region defined by the convex hull of the expanded occlusion. See Figure 4(b).

(iii) *Projection of expanded object*

The convex hull of the expanded occlusion is now projected to a sphere centered on the reduced target. The radius of the sphere is defined as the distance of the camera position to the target. This projection is seen in Figure 5. The TFV_{RF} can be directly computed from these projections.




Figure 5: Projection of expanded object

2.3.4. Target Angular Visibility

Target angular visibility is defined as the dot product of camera image normal vector and the target normal vector, or the cosine of the inter-normal angle, β . Additionally,

for most practical cases, the target may only be visible from one side. In such a situation, the contribution to the rating function due to target angular visibility is given by:

$$TAV_{RF} = \begin{cases} \cos \beta & \text{for } \frac{\pi}{2} \geq \beta \geq -\frac{\pi}{2} \\ 0 & \text{for } \beta \geq \frac{\pi}{2} \text{ or } \beta \leq -\frac{\pi}{2} \end{cases} \quad (3)$$

2.3.5. Alternate/secondary targets

In the representative scenario presented in section 1, motion of the cooperating robotic systems leads to occlusions of the target. Even with placement optimization it is possible that the target would not be visible sufficiently well to perform the task. One way to resolve such problems, is to identify additional targets and the rating function evaluated for these secondary targets. A secondary target may be a feature whose geometric relationship to the original or primary target is well known within defined tolerances. The above rating function is modified to evaluate all known secondary targets. This rating function reflects the uncertainty in the geometric relationship between the secondary and primary targets (ST_{error}) with modification to RES_{RF} :

$$R = \frac{2d \tan(\alpha/2)}{n} + ST_{error} \quad \text{and} \quad Res_{RF} = \frac{1}{R} \quad (4)$$

2.4. Camera motion correction

Once an optimum pose for the vision system is obtained the physical motion of the cameras to the desired pose is achieved by manipulator motion. However, due to positioning errors, the true camera pose needs to be identified. This process eliminates manipulator positioning errors and vehicle suspension motions, and allows for accurate data fusion. A single spatial point in the base frame, \bar{r}_i , is related to the image point (u_i, v_i) by the 4x4 transformation matrix \mathbf{g}_{01} (see Figure 6).




Figure 6: Relationship of camera and target frames

For motion calibration we need to identify \mathbf{g}_{01} :

$$\begin{bmatrix} k_i u_i \\ k_i v_i \\ k_i f \\ 1 \end{bmatrix} = \bar{u}_i = \mathbf{g}_{01} \cdot \bar{r}_i = \begin{bmatrix} [\mathbf{R}_{01}]_{3 \times 3} & \bar{\mathbf{X}}_{3 \times 1} \\ \bar{0} & 1 \end{bmatrix} \cdot \begin{bmatrix} r_i^x \\ r_i^y \\ r_i^z \\ 1 \end{bmatrix} \quad (5)$$

where \mathbf{R}_{01} is the rotational matrix, $\bar{\mathbf{X}}$ 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 3 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. Equation 5 may be generalized for n points resulting in a set of linear equations that can be readily solved using conventional techniques. A least mean square error solution is given by:

$$\mathbf{g}_{01} = \mathbf{u}(\mathbf{r}^T \mathbf{r})^{-1} \mathbf{r}^T \quad (6)$$

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 $(i+1)^{\text{st}}$ spatial and image point and its uncertainty, from Equation 6 \mathbf{R}_{i+1} and \bar{X}_{i+1} and their associated uncertainties 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:

$$\begin{aligned} \hat{\bar{X}}_{i+1} &= \frac{(i\hat{\bar{X}}_i + \bar{X}_{i+1})}{i+1} \\ C_{i+1}^{\bar{X}} &= \frac{iC_i^{\bar{X}} + [\bar{X}_{i+1} - \hat{\bar{X}}_{i+1}][\bar{X}_{i+1} - \hat{\bar{X}}_{i+1}]^T}{i+1} \\ \hat{\mathbf{R}}_{i+1}^{(l,m)} &= \frac{(i\hat{\mathbf{R}}_i^{(l,m)} + \mathbf{R}_{i+1}^{(l,m)})}{i+1} \\ C_{i+1}^{\mathbf{R}^{(l,m)}} &= \frac{iC_i^{\mathbf{R}^{(l,m)}} + [\mathbf{R}_{i+1}^{(l,m)} - \hat{\mathbf{R}}_{i+1}^{(l,m)}][\mathbf{R}_{i+1}^{(l,m)} - \hat{\mathbf{R}}_{i+1}^{(l,m)}]^T}{i+1} \end{aligned} \quad (7)$$

This essentially maintains a measure on how certain the camera motion is w.r.t. its original configuration (assuming the original configuration is known very precisely w.r.t. the common reference frame).

If the vision system is unable to maintain an original set of spatial markers within its field of view (known w.r.t. the target reference frame), then a new set of markers must be selected and the appropriate uncertainties must be accounted for. The issue of obtaining appropriate spatial points is now addressed. Spatial points are a visible set of fiducials that are tracked during camera motion. As the camera moves, the fiducials move relative to the camera, possibly moving out of the camera view. This requires methods to identify and track new fiducials. Fiducials are selected from the environment model [16] and the vision system view based on three criteria: the degree of certainty with which a sampled point is known, the visual contrast of the sampled point with its surroundings, and depth contrast of the sampled point with its surroundings. These are combined into a single fiducial evaluation function:

$$F.E.F. = f(P(\mathbf{x})) + g(C(u,v)) + h(H(\mathbf{x})) \quad (8)$$

(i) Fiducial certainty: $f(P(\mathbf{x})) \sim P(\mathbf{x})/r$, where r is the radius of a sphere centered at the potential fiducial within which neighboring voxels have decending certainty levels. Outside this sphere voxel certainty levels increase. Lower values for r suggest that the region surrounding a potential fiducial is well known—a desirable property.

(ii) Fiducial visual contrast: $g(C(u,v)) \sim \text{contrast} (C) * \text{window size} (w)$. Contrast is defined as:

$$C(u,v) = \frac{I(\mathbf{x}) - \bar{I}_w}{\bar{I}_w} \quad (9)$$

where $I(\mathbf{x})$ is the 2D image intensity value of the potential fiducial at \mathbf{x} , \bar{I}_w is the average intensity of a window centered at the potential fiducial in the 2D image, and w is the maximum window size after which the contrast starts to decrease.

(iii) Fiducial depth contrast: $h(H(\mathbf{x})) \sim H(\mathbf{x}) * \text{window size} (w)$, where $H(\mathbf{x})$ is the maximum spatial frequency (from a 3D Fourier transform) at the potential fiducial within a window, and w is the maximum window size after which the power spectrum (of the 3D Fourier transform) starts shifting to higher frequencies.

Additionally, a penalty is added if a potential fiducial is too close to other identified fiducials. Using the identified fiducials, camera motion can be identified. Fiducials can be tracked with simple methods such as region growing or image disparity correspondence.

To obtain an estimate on the absolute position and the uncertainty of a fiducial in the environment, the camera pose uncertainty must be accounted for when the fiducial is measured. 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 random noise vector \bar{v} of known covariance matrix, \mathbf{R} .

$$\bar{z} = h(\bar{x}) + \bar{v} \quad (10)$$

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:

$$H_k = \left. \frac{\partial h(\bar{x})}{\partial \bar{x}} \right|_{\bar{x}=\bar{x}_k} \quad (11)$$

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

$$\begin{aligned} K_k &= P_k H_k^T [H_k P_k H_k^T + R_k]^{-1} \\ \hat{\bar{x}}_{k+1} &= \hat{\bar{x}}_k + K_k [\bar{z}_k - h(\bar{x}_k)] \\ P_{k+1} &= [1 - K_k H_k] P_k \end{aligned} \quad (12)$$

This estimate is also known as the Extended Kalman Filter.

3. Simulation Results

Results using the RF to define an optimal camera pose given the probabilistic geometric world map are shown

here. Regions where the probability of occlusion < 0.05 are considered empty, and form candidates for optimal camera placement locations. The rating function (RF) cannot be optimized analytically. In practice, finding an optimum value for RF requires exhaustive searching through the known configuration space—a process that takes $O(n)$ time, where n is the number of discrete points in the configuration space. Methods to reduce the search time include: (i) increasing the environment grid “coarseness”, (ii) bounding the evaluation of RF by distance to the target, (iii) employ a finite random selection of goal configurations to evaluate. Thus, while the best goal configuration would be the one maximizing RF, any configuration with a high value for RF will suffice. Such a configuration can be found with reasonable effort.

In the first simulation study, a planar environment is set up (see Figure 7a). The primary target center is located at world coordinate (50,50). A secondary target is located at world coordinate (50,75). Figure 7b shows the evaluation of the RF (Equation 1) over the entire environment, for two positions of a potential occlusion. It is assumed that the environment is known in both instances. Accounting for actual motions of objects will be demonstrated later. Note that the RF value increases as the pixel intensity increases i.e. the darker the intensity the better the camera location. In the first instance, the optimum location is found by viewing the primary target. However, in the second instance, the optimum location is found by viewing the secondary target.

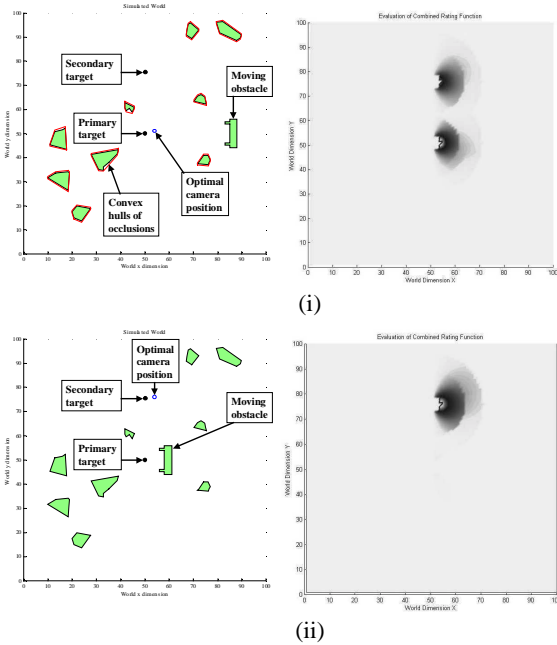


Figure 7: Optimal camera placement

For most practical situations, it is expected that there would be movement of some objects in the environment, since often the task would involve motions by one or more

of the cooperating robots. For example, the cooperative assembly of a panel into a mating slot. Clearly, the optimum camera location would change as a function of the panel position. This is analogous to a human mounting a clock or picture frame on a wall. As the object is brought toward the target (e.g. a hook), the human repositions his or her head to continue monitoring the approach of the object to the target.

The optimal positioning of a vision system in the presence of moving objects is addressed in the second simulation study. For cooperative robots, it is assumed that moving objects are well known (e.g. CAD models are given). However, the measured uncertainty associated with their position must be accounted for. As mentioned earlier, every measured environment point has an associated uncertainty. Rather than remap the environment every time an object (with a known model) moves, the algorithm simply updates the environment model using the object CAD model. This is achieved as follows.

1. Grid points/voxels in the environment model belonging to the moving object are identified. This is achieved using conventional image processing approaches such as template matching, Hough transforms, etc. The CAD model of the object is fit to the mapped points. These grid points are removed from the environment model and assumed to be unoccupied space.
2. Points in the current field of view of the vision system corresponding to the moving object are identified. The identified points are fit to the object CAD model (as in step 1).
3. The current position of the moving object is identified. This is achieved using the known vision system pose and the identified object pose relative to the vision system.
4. Measurements from step 2 and the object CAD model are used to update the environment model.

Table 1 presents the results of simulating a similar scenario. However, here the task has been modified from before. The vision system is used to guide the object to the goal (insertion site). Hence, the target is the displacement from the object to the goal. The task is successful if the object can be visually guided to the goal within the defined tolerances. Several different task scenarios were considered where the task difficulty and the occlusion density were varied. Task difficulty is measured as the tolerance (as a function of object size) permitted in the task that still allows for success. Occlusion density is measured as the percentage of the angular projection from the target onto a target centered circle that results in an occlusion (where the maximum circle radius is given by the vision system depth of field). The simulation is carried out 300 times for each different combination of task difficulty and occlusion density, and compares the approach developed in this paper with another heuristic

method: biased random placement—vision system position is selected randomly with some heuristic constraints (bounded region of placement, bounded target angular view, etc.). Note that as the task difficulty increases for low occlusion density (5%), the optimal camera placement algorithm still performs well. For the same change, the heuristic algorithm starts to decay

significantly in performance. However, for high occlusion densities, both algorithms yield significant performance decay. This may suggest the need to re-evaluate the field scenario for task execution when occlusion density is high. Finally, the influence of secondary targets is seen as task difficulty or occlusion density increases.

Table 1: Results of changing task difficulty, occlusion density and task execution success

300 tests per scenario		Occlusion Density 1 (5%)		Occlusion Density 2 (20%)		Occlusion density (35%)	
		Without sec. target	With sec. target	Without sec. target	With sec. target	Without sec. target	With sec. target
Task difficulty: easy→20% tolerance	Optimal camera placement success (%)	100	100	76	95	13	25
	Random camera placement success (%)	45	58	16	28	5	9
Task difficulty: medium→10% tolerance	Optimal camera placement success (%)	99	100	63	86	10	18
	Random camera placement success (%)	23	30	8	15	3	4
Task difficulty: hard→1% tolerance	Optimal camera placement success (%)	97	99	30	52	3	7
	Random camera placement success (%)	1	1	<<1	<<1	<<1	<<1

4. Conclusions

In field environments it is often not possible to provide robotic teams with detailed a priori environment and task models. In such unstructured environments, cooperating robots will need to create a dimensionally accurate 3-D geometric model by performing appropriate sensor actions and position their sensors in a task directed optimal way. A new algorithm based on iterative sensor planning and sensor redundancy is proposed to overcome the camera occlusion from fixed poses, to build the 3-D environment model, and to position sensors to observe a task. The environment modeling stage of the algorithm was developed in detail in a previous paper [16]. This paper addresses the process of task directed optimal camera placement, given the developed environment model. This algorithm is based on iterative sensor planning and exploiting the sensor redundancy of cooperative robotic systems. A rating function is developed and optimized to find the most suitable pose to view the target. Simulations show promising results.

References

[1] Asada, M. *Map building for a mobile robot from sensory data.* IEEE Trans. on Sys., Man, and Cyb. Vol. 37, no. 6, Nov/Dec 1990.
[2] Baumgartner, E.T., P. S. Schenker, C. Leger, and T. L. Huntsberger. *Sensor-fused navigation and manipulation from a planetary rover.* Proc. SPIE Sym. on Sensor Fusion and Decentralized Cntrl. in Rob. Sys. Vol. 3523, Boston, Nov 1998.
[3] Betge-Brezetz, S., Hebert, P., Chatila, R., and Devy, M. *Uncertain map making in natural environments.* Proc. of 1996 IEEE ICRA, Minneapolis, Minnesota, April, 1996.
[4] Burschka, D. Eberst, C. and Robl, C. *Vision based model*

generation for indoor environments. Proc. of 1997 IEEE ICRA, Albuquerque, New Mexico, April, 1997
[5] Chu, G-W., Yu, W.P., and Chung, M.J. *A simple method of finding a visible region for the optimal camera position.* Proceedings of the 1997 IEEE ICAR. Monterey, CA. July 1997.
[6] Connolly, C.I. *The determination of the next best views.* Proc. of the IEEE ICRA, pp. 432-435, 1985.
[7] Cowan, G.K. and Kovesi, P.D. *Automatic sensor placement from vision task requirements.* IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 10, no.3, pp. 407-416, May 1988.
[8] Huntsberger, T.L. *Autonomous multi-rover system for complex planetary retrieval operations.* Proc. SPIE Sym. on Sensor Fusion and Decentralized Control in Autonomous Robotic Agents, Pittsburgh, PA, Oct 1997, pp. 220-229.
[9] Kececi, F., et. al. *Improving visually servoed disassembly operations by automatic camera placement.* Proc. of 1998 IEEE ICRA, Leuven Belgium. May 1998.
[10] Kruse, E., Gutsche, R., and Wahl, F.M. *Efficient, iterative, sensor based 3-D map building using rating functions in configuration space.* Proc. of IEEE ICRA, Minneapolis, Minnesota, April, 1996.
[11] Laugier, C. and Triggs, B. *Automatic camera placement for robot vision tasks.* IEEE ICRA, Nagoya, Japan, 1995.
[12] Lumelsky, V., Mukhopadhyay, S. and Sun, K. *Sensor-based terrain acquisition: the "sightseer" strategy.* Proc. of 28th Conf. on Decision and Control. Tampa, Florida. December 1989.
[13] Luo, R.C. and M.G. Kay. *Multisensor integration and fusion in intelligent systems.* IEEE Transactions on Systems, Man, and Cybernetics, Vol. 19, No. 5, September, 1989.
[14] Shaffer, G.; Stentz, A. *A robotic system for underground coal mining.* Proc. of IEEE ICRA, 1992. Pages(s): 633 -638 vol.1
[15] Smith, R.C. and Cheeseman, P. *On the representation and estimation of spatial uncertainty.* International Journal of Robotics Research. 5(4): 56-68
[16] Sujan, V.A. and Dubowsky, S. *Visually Built Task Models for Robot Teams in Unstructured Environments.* Proceedings of the IEEE ICRA, May 11-15, 2002, Washington, D.C., U.S.A.
[17] Tarabanis, K.A., Allen, P.K. and Tsai, R.Y. *A survey of sensor planning in computer vision.* IEEE Transactions on Robotics and Automation, Vol. 11 no. 1, Pp. 86-104. February, 1995.