SENSOR-BASED INTERIOR MODELING

Martial Hebert
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268-2585
E-mail: hebert@ri.cmu.edu

Regis Hoffman
Field Robotics Center
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268-6987
E-mail: rmh@cs.cmu.edu

Andrew Johnson
Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268-6880
E-mail: aej@ri.cmu.edu

James Osborn
Field Robotics Center
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268-6553
E-mail: oz@ri.cmu.edu

I. ABSTRACT

Robots and remote systems will play crucial roles in future decontamination and decommissioning (D&D) of nuclear facilities. Many of these facilities, such as uranium enrichment plants, weapons assembly plants, research and production reactors, and fuel recycling facilities, are dormant; there is also an increasing number of commercial reactors whose useful lifetime is nearly over. To reduce worker exposure to radiation, occupational and other hazards associated with D&D tasks, robots will execute much of the work agenda. Traditional teleoperated systems rely on human understanding (based on information gathered by remote viewing cameras) of the work environment to safely control the remote equipment. However, removing the operator from the work site substantially reduces his efficiency and effectiveness. To approach the productivity of a human worker, tasks will be performed telerobotically, in which many aspects of task execution are delegated to robot controllers and other software.

This paper describes a system that semi-automatically builds a virtual world for remote D&D operations by constructing 3-D models of a robot's work environment. Planar and quadric surface representations of objects typically found in nuclear facilities are generated from laser rangefinder data with a minimum of human interaction. The surface representations are then incorporated into a task space model that can be viewed and analyzed by the operator, accessed by motion planning and robot safeguarding algorithms, and ultimately used by the operator to instruct the robot at a level much higher than teleoperation.

II. INTRODUCTION

A. Scenario

The DOE Robotics Technology Development Program has identified the Integrated Process Demonstration Facility (IPDF) at Oak Ridge National Laboratory as a suitable location for testing and demonstrating robotic equipment currently being developed for decontamination and dismantlement. Within the IPDF is a large apparatus known as the IODOX (Figure 1) that was built to remove iodine from effluent gases generated in other parts of the reprocessing plant. Because the IODOX contains a myriad of process equipment and their interconnections, it is an excellent candidate for testing and demonstrating selective equipment removal scenarios using mobile worksystems. The majority of the process equipment is housed within a framework of painted carbon steel I-beams measuring approximately 14 feet wide x 14 feet deep x 24 feet high in three 8-foot stories with open lattice floor grates. Equipment within this

Figure 1. The IODOX facility, a DOE test site for selective equipment removal.

frame includes several process tanks, un-insulated pipes and conduits ranging in size from 1/4” instrumentation lines to 4” supply lines, insulated pipes with overall diameters up to 8”, pumps, valves, and various types of control components.

Remote D&D tasks require manipulators to move objects and position end effectors; the size and extent of most facilities dictates the need for transporters, such as mobile robots and overhead gantries, to move the manipulators. In near term DOE facility D&D, such mobile worksystems will be used for selective equipment removal, in which some part of an apparatus is extricated while minimally disturbing the surrounding objects. The motivation for such a mission is to remove highly contaminated components from an area that is otherwise relatively clean thus lowering the overall average level of contamination. Such surgical dismantlement is necessary in cases where the contamination is inaccessible; washdown, scarification and other surface treatment procedures are ineffective if the equipment is internally contaminated. Selective equipment removal will be particularly relevant during the deactivation and deinventory stages of facility decommissioning as a means to reduce the costs and risks associated with subsequent surveillance and monitoring.

B. Sensor-Based Modeling

Before a truly telerobotic selective equipment removal system can become a reality, issues of sensing, navigation, manipulation and planning must be addressed. A critical enabling technology is the task based scene analysis required to make the robot and the human operator aware of the operating environment surrounding the robot. Task space scene analysis is a paradigm in which geometric data from the site is collected, processed, and displayed by a computer to an operator in a useful and convenient form.

Facilities such as the IODOX pose unique challenges to the design of a task space scene analysis system. The facility to be modeled is typically much larger than the robot; this implies that data from multiple sensor viewpoints must be merged to construct a complete model. Complexity of the environment is the second challenge. The IODOX is a maze of process equipment with irregular sizes and locations. The number of objects and occlusions is large and there is a tremendous variation in scale. Fortunately, one can safely assume that all objects in the facility are man-made. This limits the types of objects that will be encountered in the model building phase. Further, most can be modeled by simple geometric descriptions and can be readily identified by a human observer. A third challenge then is to devise a system that takes maximum advantage of the human’s object recognition ability while minimally burdening him.

Previous work uses a variety of sensors and representations to construct models of interior workspaces. Christensen [1] describes a supervised teleoperated system with a world model that includes *a priori* knowledge, the robot configuration, and information gathered by the robot’s sensors. The world model allows the operator to preview the operation, have the computer automatically plan an end-effector trajectory, and view the task from many different viewing positions. Trivedi [2] reports on another model-based system using range sensors that allows testing of robot plans in simulation. Azarbayejani [3] has a system to semi-automatically construct CAD models from uncalibrated video images. Thayer [4] computes an object’s location and orientation using stereo vision. The operator performs the stereo matching of some points on the object and the system computes the pose of the object.

These and other modeling systems share common characteristics that justify their development:

1. They are a graphical aid to the operator by allowing viewing the robot and its workspace from many more angles than are possible with on-board cameras.
2. They store knowledge of the world in the form of a model; this capability is essential for automatic planning agents.
3. They augment uncertain or unrecorded *a priori* information with current *in situ* data.

Our system retains these characteristics but differs from the cited work in that there need not be any *a priori* world model, the input data consists of 3-D range images, the internal representation is surface-based, and the system requires limited human interaction.

III. OBJECT REPRESENTATION

Many representations of the robot’s task space are possible; the appropriateness of the representation depends on the function the robot is called on to perform. For example, a volumetric representation that segregates space into regions that are either occupied or empty is sufficient for planning collision-free paths. More sophisticated methods connect associated regions of space; for example a surface-based representation that groups regions with similar curvature. Such a representation is needed to plan actions like grasping an object, applying a surface coating, or obtaining a sample. To plan actions at the task level (*e.g.*, “close that valve”, or “cut that pipe”), a representation that classifies and assigns semantic meaning to objects in the robot’s task space is required.

In the context of selective equipment removal, the robot workspace contains a large number of man-made objects - pipes of various dimensions, waste barrels, control boxes, valves, holding tanks, and pressure vessels (Figure 1). Such objects can be concisely modeled with planar and/or quadric surfaces patches: groupings of points in space belonging to surfaces described by three dimensional functions.

A planar surface is described by its surface normal (n_1, n_2, n_3) and cartesian distance d from the origin of the world coordinate frame as follows:

$$n_1 x + n_2 y + n_3 z + d = 0 \tag{1}$$

Flat surfaces such as walls, floors, and beams, are readily modeled with planar patches. Curved objects such as cylinders, spheres, and cones which are prevalent in man-made environments can be represented with quadric surfaces.

A quadric surface is a second order three-dimensional surface with ten parameters given by:

$$ax^2 + by^2 + cz^2 + dxy + exz + fyz + gx + hy + iz + j = 0 \tag{2}$$

Combinations of planar and quadric surfaces can be used to describe most objects that mobile worksystems will encounter. Planar and quadric surface patches were chosen to represent objects because their parametric representations do not vary when surfaces are partially occluded or the sensing viewpoint changes. Furthermore, the position and orientation of surface patches in the scene can be easily computed from their corresponding parameters.

IV. MODELING SYSTEM

Our system (which we call “Artisan”) is a combination of sensors, modeling and analysis software, and an operator interface that creates 3-D models of indoor man-made environments as they are discovered. The interactive system described performs the following functions:

1. Acquire 3-D sensor data of the facility in the form of a range image.
2. Select a region of interest in the range image.
3. Generate a mesh connecting range pixels within the region of interest.
4. Partition mesh into planar surfaces patches using region growing.
5. Group planar surfaces patches into quadric surface patches.
6. Incorporate surface patches into a world model of the robot and workspace.

A block diagram of the above operations and data flow is shown in Figure 2.

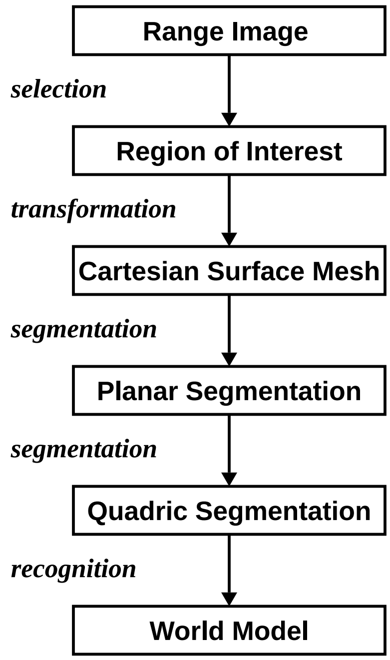

Figure 2. Modeling system block diagram.

A. 3-D Sensor Data

The sensor is a 3-D scanning laser rangefinder manufactured by Perceptron, Inc. It acquires 256 x 256 pixel range and intensity images over a vertical and horizontal field of view of 60 degrees at a frame rate of 2 Hz. The scanner’s range is 2 to 40 meters and its range precision is 5-7 cm [5]. To map a facility, the scanner is remotely positioned by a mobile worksystem and commanded to acquire images by the human operator. The range and intensity images acquired are displayed on the operator’s console. Figure 3 shows a sample image and reflectance image pair taken at our experimental test site. The scene consists of a series of cylindrical pipes of various dimensions against a flat background.

B. Region of Interest

To focus the system on a specific object in the scene to be analyzed, the operator selects a rectangular region of interest in the image. This limits the amount of data that is processed in the surface generation phase to that which is

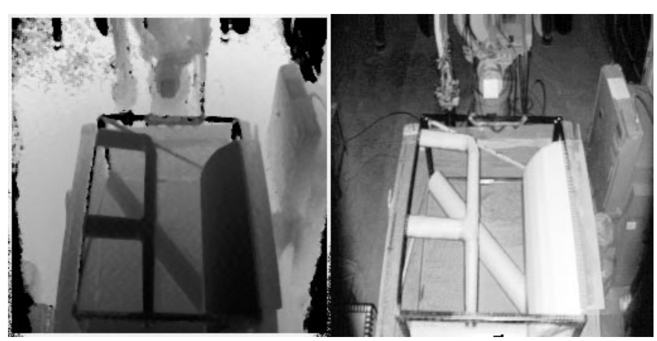

Figure 3. Range image (left) and reflectance image (right) of experimental testbed. Contrast enhanced for viewing.

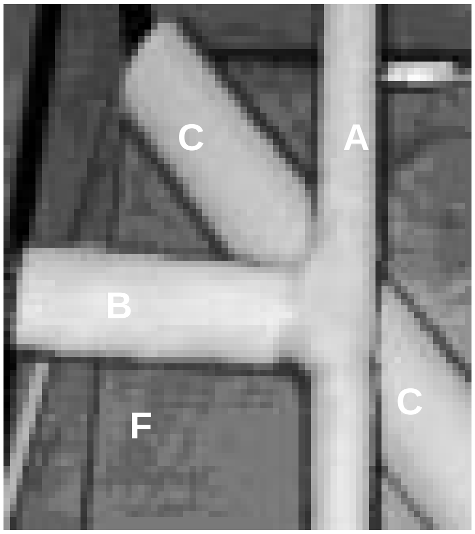

Figure 4. Region of interest in reflectance image from Figure 3. It contains a large pipe partially occluded by two pipes forming a T-joint. Labels do not appear in the actual image.

important to recognizing the specific object. Figure 4 shows a region of interest selected from the reflectance image of Figure 3. For discussion purposes, the two pipes that form a T-joint are labelled “A” and “B”, the pipe beneath it (and occluded by it) is labelled “C”, and the flat region is labelled “F”.

C. Generating a Mesh

Once the region of interest is determined all of the points in it are filtered by a temporal averaging of an image sequence to reduce noise. This is followed by the application of a spatial smoothing filter to eliminate outliers while preserving range discontinuities. The resulting range image is then converted from spherical sensor coordinates (ρ,θ,ϕ) to cartesian world coordinates (x,y,z).

To group points into surface patches a surface mesh must be created that establishes the local connectivity of the points in space The arrangement of pixels in the range image is used to determine the nearest neighbors for each corresponding cartesian point. Nearest neighbors are connected if the distance between them is less than a specified threshold. The purpose of this distance threshold is to preserves range discontinuities when converting from the sensor to world coordinates. The resulting set of connected points constitutes a surface mesh in cartesian space which is used in all subsequent processing stages. Though more sophisticated techniques exist, this approach is sufficient for generating meshes from individual range images. The benefit of converting to cartesian coordinates is the ability to fuse range data collected from images taken from different viewpoints

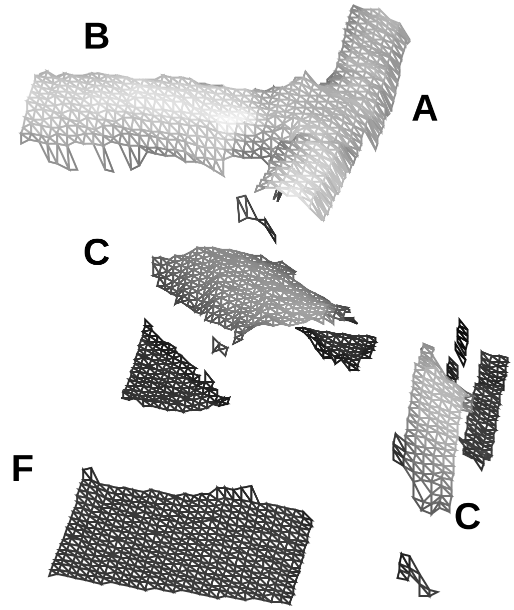

Figure 5. Result of constructing a surface mesh from the region of interest. Meshes are identified with the corresponding labels from Figure 4.

into a single representation. Figure 5 shows the surface mesh generated from the range image data from Figure 4. To illustrate the correspondence between the mesh and the range image data, the meshes are identified in the figure with the corresponding labels from Figure 4.

D. Finding Planar Patches

The next step after the construction of the surface mesh is to segment the mesh into planar surface patches. Planar surface patches are groupings of points that are locally connected by the mesh and are close to some planar surface. In man-made environments planes are very prevalent, so planar patches can be used to describe many surfaces that appear in interior workspaces. Planar patches are also used as an intermediate representation of curved surfaces to provide a grouping of points that can later be merged into quadric surfaces patches.

The algorithm for segmenting the mesh into planar and quadric surface patches follows closely that presented by Faugeras and Hebert [6]. This algorithm segments the mesh through region-growing. Initially each point in the mesh is considered to be a planar surface patch whose parameters are determined by fitting a plane to the point and all of its nearest neighbors as determined by the mesh. Then each initial planar surface patch is inserted into a region adjacency graph and is connected to its nearest neighbors through weighted links. The links between patches are weighted by the least squares fit error of merging the points in the two adjacent planar patches into one planar patch. All of the links are stored in a priority queue sorted by the fit error with the link with the smallest fit error at the top of the queue.

The two regions connected by the link at the top of the queue with the minimum fit error are then merged into one region. The region adjacency graph is updated by adding and deleting links to reflect the creation of the new larger region and the elimination of the two old regions. Similarly, the priority queue is updated by deleting links that no longer exist in the region adjacency graph and inserting links in their appropriate position based on the fit errors between the new region and its neighbors in the graph. This merging procedure continues until the total fit error given by

$$E_T = E_{ij} - E_i - E_j \tag{3}$$

where

E_{ij} is the fit error from merging regions *i* and *j*,
E_i is the fit error on region *i*,
E_j is the fit error on region *j*.

exceeds some predetermined threshold. Because the merge with the smallest fit error between all of the planar patches in the region adjacency graph is executed at each iteration in the region growing process, the best segmentation in a global sense is ensured. This global criterion prevents merging across edges in the mesh. In contrast, local algorithms that only consider merges local to the current region may cause undesirable merging of two regions that are connected across an edge.

The end result of the planar segmentation is a set of planar patches on the surface mesh. Figure 6 shows the result of the planar segmentation of the mesh in Figure 5.

The planar surfaces may constitute actual planar surfaces or may be part of a larger quadric surface that appears in the scene. The next step in the surface segmentation will make this distinction.

E. Finding Quadric Patches

Because objects like pipes, tanks and valves exist in the environment that is to be modeled, planar surface patches are not sufficient to model all of the objects in the scene. Ideally, every type of surface that appears in the scene should be representable by the surface modeling system. However, arbitrarily complex surfaces are not feasible when segmenting range images of finite precision and accuracy because many different surfaces could conceivably fit each grouping of points. Quadric surfaces, which are three-dimensional second-order surfaces, are a compromise between the representation and the ease of computation of the parameters of complex surfaces.

Merging planar surface patches into quadric surface patches enables the modeling system to model cylinders, spheres and cones. From the parameters of the quadric surface the modeling system extracts radii, positions and shape of objects in the scene which are incorporated into a world model of the robot's environment for use in path planning and manipulation.

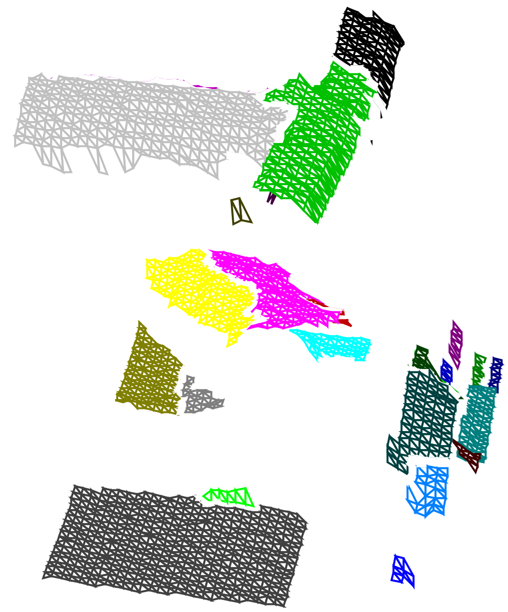

Figure 6. Identification of planar surfaces segmented from the surface mesh.

The algorithm for creating quadric surface patches is very similar to that used to generate planar surface patches through global region growing. A region adjacency graph connecting all of the adjacent planar patches through weighted links in the scene is created. However, planar patches that are flat enough to be considered actual planar surfaces are not included in the region adjacency graph. The links are weighted based on the fit error for merging the two adjacent planar patches into a quadric patch. The difference between the algorithms lies in the fact that links are inserted into a priority queue (sorted on fit error) only if they pass additional tests at the boundaries of the two regions. These tests include a test on the continuity of the surface normals of the points along the boundary, thus preventing regions from being merged across an edge in the scene, and a test on the size of the boundary connecting the two regions which will keep regions with short boundaries from being merged. The purpose of these boundary tests is to prevent regions that should not be merged from being merged. The algorithm proceeds as in the planar case with the region adjacency graph and priority queue being updated at each iteration until the total fit error exceeds a threshold.

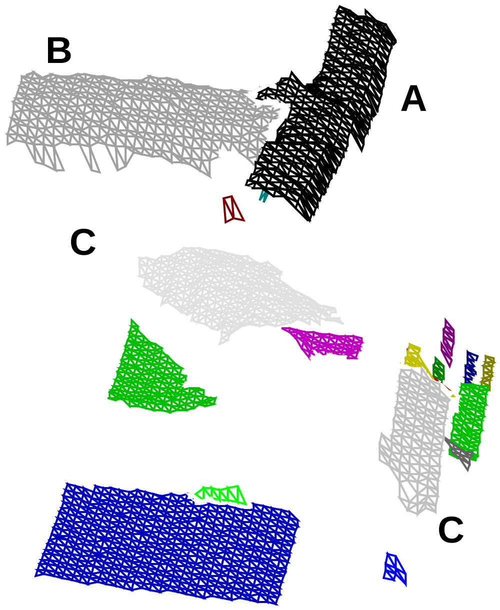

Figure 7. Identification of final planar and quadric surfaces.

The end result of the quadric surface segmentation is a set of quadric and planar surface patches. Figure 7 shows the result of applying the quadric surface algorithm on the planar patches of Figure 6.

F. World Model

The final step in the modeling process is to present useful geometric information to the human operator. The initial version of the system finds the quadric surface patches that are approximately cylindrical. After picking the cylindrical quadrics, the system calculates their axis, position and radii from their quadric parameters. The cylinders are then displayed to the user at their estimated position and orientation in the world model. The view presented to the operator for the region of interest shown in Figure 4 is shown in Figure 8. The cylinders are displayed with the radii that were calculated from their quadric parameters. To effectively convey a view of the robot in its workspace, the commercial robot simulation package IGRIP from Deneb Inc. is used. This package allows viewing of the scene from any angle, manipulator path- planning, and off-line simulation of robot actions. The pre-stored model of the robot vehicle, manipulators and tooling is updated with the sensor-based models as they are constructed. A virtual 3-D world provides richer understanding of the environment than 2-D camera data. The accuracy of the world model dictates the possible functions the operator can perform with the virtual world. Previewing the world from an arbitrary viewing angle does not require high accuracy models. As task are previewed and automatic plans are generated in the virtual world, model accuracy becomes more important.

Table 1 shows a comparison of the calculated radii versus the radii measured from the actual pipes that appear in the scene. Because only a partial view of the cylinders is available to the quadric segmentation routine, the quadric parameters for each cylinder are not be estimated perfectly. As a result, the radii which are calculated from the parameters are consistently less than their measured value. The relatively large range error on the sensor corrupts the parameter estimation even further. In future versions of the system multiple views will be merged to form a more complete model of the scene which will improve the estimation of the parameters of all surfaces imaged.

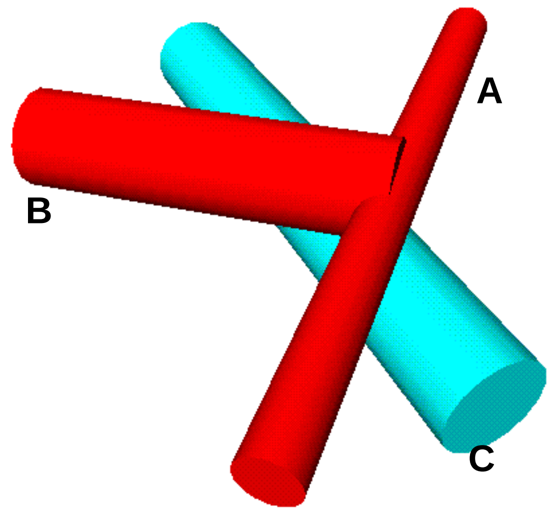

Figure 8. Cylinders calculated from quadric surface patches.

The need for accuracy must be stressed if the models are to be used for purposes beyond synthetic displays of the task space, such as for motion planning. There are many sources of error (sensor noise, sensor calibration errors and inaccuracy in sensor positioning control) which compound to diminish the overall accuracy of each range point. Errors in the range points will propagate to errors in the parameters of the surface patches. An error analysis is in progress that will be used to predict the accuracy of the modeling system.

Cylinder	**Measured Radius**	**Calculated Radius**
A	4.8 cm	3.5 cm
B	8.0 cm	6.8 cm
C	13.0 cm	11.7 cm

Table 1. Comparison of measured and calculated radii for recognized cylinders.

V. DISCUSSION

Many internal details of the system operation have been described and illustrated, from the collection of the range image, through segmentation, to final model display. The system as used by a human operator hides most of this detail. The human operator initiates range image acquisition and selects a region of interest; other steps are performed automatically Once an object is recognized an displayed, the operator can reject a candidate model before it is entered into the world model database.

Future work will concentrate on three primary improvements. First, the worksystem will move the rangefinder to collect images from several viewpoints. The captured data will thus have a larger field of view and fewer occlusions than a single image. A second improvement will be to increase the number of surface primitives from the current set of plane and cylinder. Logical extensions are sphere and cone primitives which can be extracted from the quadric surface representation. Finally, to model more complicated objects such as valves, surfaces must be linked into more complex primitives.

VI. ACKNOWLEDGMENTS

This work is funded under contract DE-AC21-92MC29104 from the United States Department of Energy Morgantown Energy Technology Center. The authors wish to acknowledge the technical support and guidance of the DOE Robotics Technology Development Program, particularly Dr. Linton Yarbourgh and Dr. William Hamel as well as the Telerobotics System Section of Oak Ridge National Laboratories. We also acknowledge the programmatic support of Mr. Vijay Kothari of the Morgantown Energy Technology Center. Deborah Henderson and Patrick Leger contributed technical support at Carnegie Mellon University.

VII. REFERENCES

[1] B.K. CHRISTENSEN et al. “Graphic Model Based Control of Robotics Systems for Waste Remediation,” *Proceedings of the ANS Fifth Topical Meeting on Robotics and Remote Systems*, American Nuclear Society, Chicago, Il (1993), pp. 89-96.

[2] M. TRIVEDI and C.X. CHEN “Developing Sensor-based Robotics Systems Using Virtual Reality Concepts,” *Proceedings of the ANS Fifth Topical Meeting on Robotics and Remote Systems,* American Nuclear Society, Chicago, Il (1993), pp. 165-172.

[3] A.J. AZARBAYEJANI et al. “Recursive Estimation for CAD Model Recovery,” *Proceedings of the Second CAD-Based Vision Workshop*, IEEE Computer Society, Champion, PA (1994), pp. 90-97.

[4] S. THAYER et al. “On-line Stereo Vision and Graphical Interface for Decontamination and Decommissioning Applications Using the Advanced Servo Manipulator,” *Proceedings of the ANS Fifth Topical Meeting on Robotics and Remote Systems,* American Nuclear Society, Chicago, Il (1993), pp. 287-294.

[5] M. HEBERT and E. KROTKOV “3-D Measurements from Imaging Laser Radars: How Good Are They?” *International Journal of Image and Vision Computing*, **10** (3), pp. 170-178 (April 1992).

[6] O.D. FAUGERAS and M. HEBERT “The Representation, Recognition and Locating of 3-D Objects,” *International Journal of Robotics Research*, **5** (3), pp. 27-52 (Fall 1986).