AI Magazine Volume 9 Number 2 (1988) (© AAAI)

Sensor Fusion in Certainty Grids for Mobile Robots

Hans P. Moravec

A numeric representation of uncertain and incomplete sensor knowledge called certainty grids was used successfully in several recent mobile robot control programs developed at the Carnegie-Mellon University Mobile Robot Laboratory (MRL) Certainty grids have proven to be a powerful and efficient unifying solution for sensor fusion, motion planning, landmark identification, and many other central problems MRL had good early success with ad hoc formulas for updating grid cells with new information. A new Bayesian statistical foundation for the operations promises further improvement MRL proposes to build a software framework running on processors onboard the new Uranus mobile robot that will maintain a probabilistic, geometric map of the robot's surroundings as it moves. The certainty grid representation will allow this map to be incrementally updated in a uniform way based on information coming from various sources, including sonar, stereo vision, proximity, and contact sensors The approach can correctly model the fuzziness of each reading and, at the same time, combine multiple measurements to produce sharper map features; it can also deal correctly with uncertainties in the robot's motion The map will be used by planning programs to choose clear paths, identify locations (by correlating maps), identify well-known and insufficiently sensed terrain, and perhaps identify objects by shape. The certainty grid representation can be extended in the time dimension and used to detect and track moving objects Even the simplest versions of the idea allow us to fairly straightforwardly program the robot for tasks that have hitherto been out of reach MRL looks forward to a program that can explore a region and return to its starting place, using map "snapshots" from its outbound journey to find its way back, even in the presence of disturbances of its motion and occasional changes in the terrain.

Robot motion planning systems have used many space and object representations. Objects have been modeled by polygons and polyhedra or bounded by curved surfaces. Free space has been partitioned into Vornoi regions or, heuristically, free corridors. Traditionally, the models have been hard edged; positional uncertainty, if considered at all, was used in just a few special places in the algorithms, expressed as a Gaussian spread. Partly, this oversimplification of uncertainty information is the result of analytic difficulty in manipulating interacting uncertainties, especially if the distributions are not Gaussian. Incomplete error modeling reduces positional accuracy. Seriously, it can produce entirely faulty conclusions: A false determination of an edge in a certain location, for instance, can derail an entire train of inference about the location or existence of an object. Because they neglect uncertainties and alternative interpretations, such programs are brittle. When they jump to the right conclusions, they do well, but a small error early in the algorithm can be amplified to produce a ridiculous action. Most AI-based robot controllers have suffered from this weakness.

The Mobile Robot Laboratory (MRL) at Carnegie-Mellon University has built their share of brittle controllers. Occasionally, however, MRL stumbled across numeric (as opposed to analytic) representations that seem to escape this fate. One numeric representation is deep inside the program that drove the Stanford Cart in 1979 (Moravec 1981). Each of 36 pairings of nine images from a sliding camera produced a stereo depth measurement of a given feature, identified by a correlator, in the nine images. Some pairings were from short baselines and had large distance uncertainty; others were from widely separated viewpoints with small spread. The probability distributions from the 36 readings were combined numerically in a 1000-cell array, each cell representing a small range interval (see figure 1). Correlator matching errors often produced a multipeaked resultant distribution, but the largest peak almost always gave the correct range. The procedure was the most error-tolerant step in the Cart navigator, but it alone did not protect the whole program from brittleness.

A descendant of the Cart program by Thorpe and Matthies had a path planner (Thorpe 1984) that modeled floor space as a grid of cells containing numbers representing the suitability of each region to be on a path. Regions near obstacles had low suitability, and empty space was high. A relaxation algorithm found locally optimum paths (see figure 2). The program represented uncertainty in the location, even existence, of obstacles by having the suitability numbers for them vary according to extended, overlapping probability distributions The method dealt reliably and completely with uncertainty but suffered from being embedded in an otherwise brittle program.

MRL's earliest thorough use of a numeric model of position uncertainty was in a sonar mapper, map matcher, and path planner developed initially for navigating the Denning Sentry (Moravec and Elfes 1985; Elfes 1987; Kadonoff et al. 1986). Space was represented as a grid of cells, each mapping an area 30 (in some versions 15) centimeters on a side and containing two numbers, one the estimated probability the area was empty, the other that

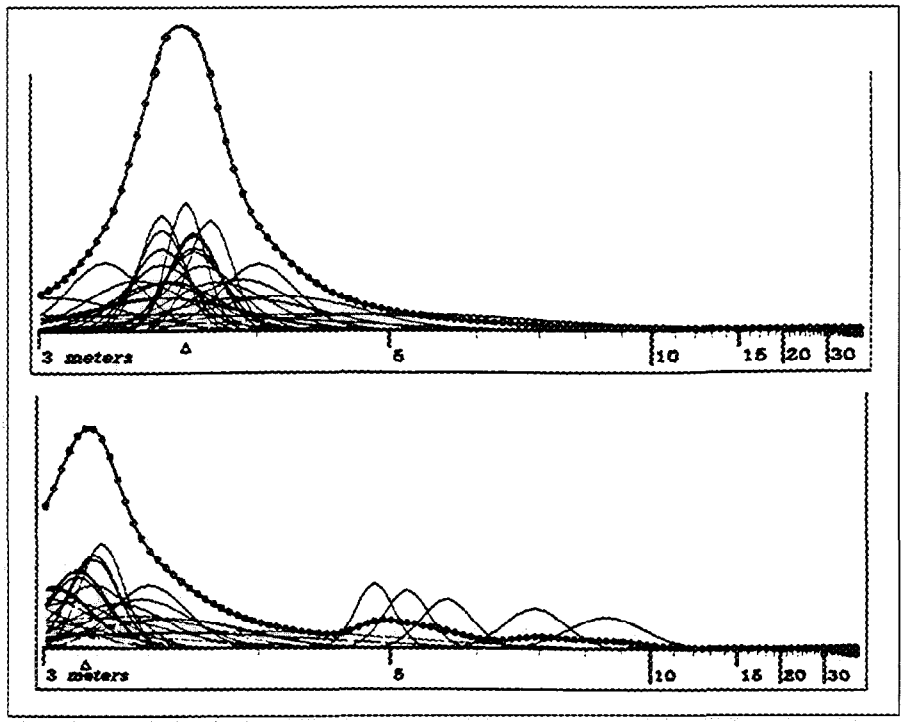

Figure 1 Nine Eyed Stereo.
Identifications of a point on an object are seen in nine different images taken as a camera traversed a track at right angles to its direction of view Each pairing of images gives a stereo baseline, some short, some long Long baselines have less uncertainty in the calculated distance The distributions for all 36 possible pairings are added in a one-dimensional certainty grid, and the peak of the resultant sum is taken as the actual distance to the object The top graph represents a case where all nine identifications of the point in the images are correct The bottom graph represents a case where one image is in error The error produces eight small peaks at incorrect locations, but these small peaks are no match for the large peak resulting from the accumulation of correct values

it was occupied. Cells whose state of occupancy was completely unknown had both probabilities zero, and inconsistent data were indicated if both numbers were high. Many of the algorithms worked with the differences of the numbers. Each wide-angle sonar reading added a 30-degree swath of emptiness and a 30-degree arc of occupancy, by itself a fuzzy image of the world. Together, several hundred readings produced an image with a resolution often better than 15 centimeters, despite many aberrations in individual readings (see figure 3). The resiliency of the method was demonstrated in successful multihour-long runs of Denning robots around and around long trajectories, using 3-second map-building and 3-second map-matching pauses at key intersections to repeatedly correct their position. These runs worked well in clutter and survived disturbances such as people milling around the running robot.

Ken Stewart of MIT and Woods Hole has implemented a three-dimensional version of the sonar mapper for use with small submersible craft. Initially tested in simulation in the presence of large simulated errors, Stewart's program provided extremely good reconstructions, in a 128 x 128 x 64 array, of large-scale terrain, working with about 60,000 readings from a sonar transducer with a 7-degree beam. Running on a Sun computer, his program can process sonar data fast enough to keep up with the approximately 1-second pulse rate of the transducers on the two candidate submersibles at Woods Hole The program was recently tested on real sonar data from a scanning transducer on an underwater robot that swam over the remains of the civil war battleship USS Monitor as part of a National Oceanic and Atmospheric Administration and United States Navy survey (Stewart 1988, 1987). The impressive results are shown in figure 4.

Recently, Serey and Matthies (1986) demonstrated the utility of the grid representation in a stereo vision-based navigator running on the MRL "Neptune" (see figure 5) mobile robot. Edges crossing a particular scan line in the two stereo images are matched by a dynamic programming method to produce a range profile. The wedge-shaped space from the camera to the range profile is marked empty; cells along the profile itself are marked occupied. The resulting map is then used to plan obstacle-avoiding paths, as with the stereo and sonar programs mentioned earlier (see figure 6).

Matthies and Elfes (1987) combined improved versions of the sonar and stereo programs into a single one that builds maps integrating data from both sensors. Their first results, also from a run of Neptune, are shown in figure 7.

In work in progress, In So Kweon (1987) of Carnegie-Mellon University successfully demonstrated the mapping of data from a scanning laser range finder (from ERIM Company, Ann Arbor, Michigan) into a three-dimensional grid.

Despite its effectiveness, in each instance, MRL reluctantly adopted the grid representation of space. This reluctance might reflect habits from a recent time when analytic approaches were more feasible and seemed more elegant because computer memories were too small to easily handle numeric arrays of a few thousand to a million cells. I think the reluctance is no longer appropriate. The straightforwardness, generality, and uniformity of the grid representation have proven themselves in finite-element approaches to problems in physics and raster-based approaches to computer graphics and have the same promise in robotic spatial representations. At first glance, a grid's finite resolution inherently seems to limit positioning accuracy. This impression is false. Cameras, sonar transducers, laser scanners, and other long-range sensors have intrinsic uncertainties and resolution limits that can be matched by grids no larger than a few hundred cells on a side, giving a few thousand cells in two dimensions or a few million in three dimensions. Because the accura-

cy of most transducers drops with range, even greater economy is possible using a hierarchy of scales, covering the near field at high resolution and successively larger ranges with increasingly coarser grids. The implicit accuracy of a certainty grid can be better than the size of its cell. The grid can be thought of as a discrete sampling of a continuous function. Extended features such as lines (perhaps representing walls) can be located to high precision by examining the parameters of surfaces of best fit. The Denning robot navigator convolves two maps to find the displacement and rotation between them. In the final stages of the matching correlation, values are obtained for a number of positions and angles in the vicinity of the best match. A quadratic least squares polynomial is fitted to the correlation values, and its peak is located analytically. Controlled tests of the procedure usually give positions accurate to better than one-fourth of a cell width.

MRL's results to date suggest that many mobile robot tasks can be solved with this unified, sensor-independent approach to space modeling. The key ingredients are a robot-centered, multiresolution map of the robot's surroundings; procedures for efficiently inserting data from sonar, stereo vision, proximity, and other sensors into the map; other procedures for updating the map to reflect the uncertainties introduced by imprecise robot motion; and yet other procedures to extract conclusions from the maps. Procedures that produce local and global navigational fixes and obstacle-avoiding paths from such maps have already been demonstrated. Other tasks, such as tracking corridors, finding vantage points with good views of unseen regions, and identifying larger features such as doors and desks by general shape, seem within reach.

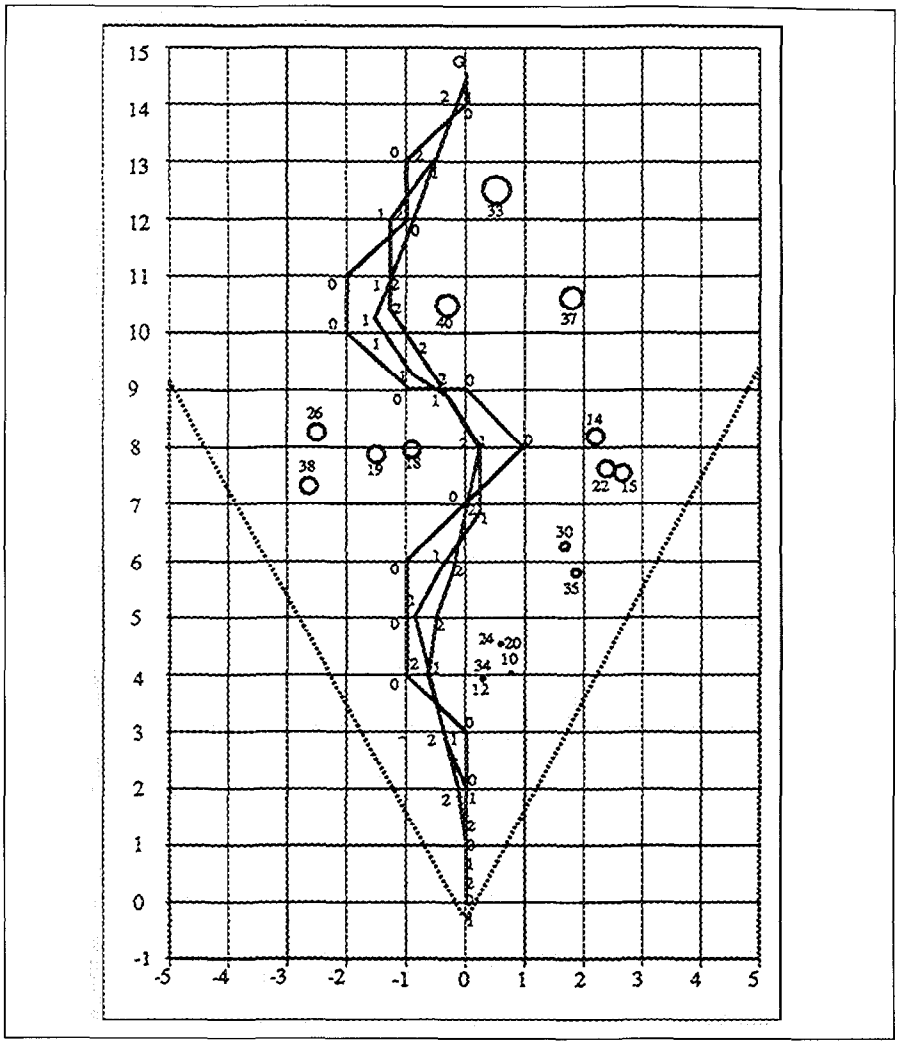

Figure 2 Relaxation Path Planner.
A path is chosen that minimizes a given cost function in a certainty grid Small perturbations are made in the vertexes of the path in directions that reduce the cost.

The Representation

The sonar mappers mentioned earlier is MRL's most thorough use to date of the certainty grid idea. Although original implementations used two grids to represent occupancy knowledge (labeled $P_{occupied}$ and P_{empty}, Stewart's three-dimensional system used only one An analysis of the steps in MRL's code revealed that one grid did indeed suffice, and this simplification clarified several puzzling issues in the original formulation. Here, I present a sample of the methods MRL used in most of its work to date. A discussion of a new, Bayesian approach follows.

Before any measurements were made, the grid was initialized to a background occupancy certainty value, Cb. This number represented the average occupancy certainty expected in a mature map and encoded a little a priori information MRL had about the world. In MRL, a good Cb seemed to be about the number of cells in the perimeter of the grid divided by the total cells (4 x 32/(32 x 32) = 1/8) in the case of the Denning code. If the space was cluttered, Cb should be larger. As the map was used, values near Cb would indicate regions whose occupancy state was essentially unknown, those much nearer zero would represent empty places, and those much nearer unity were likely to be occupied. Most of the planning algorithms that used the grid are better off if they do not make sharp distinctions but instead numerically combine the certainty values from various cells to produce "goodness of fit" numbers for their various hypotheses In this way, the essential uncertainties in the measurements are not masked, and the algorithms do not jump to unnecessary, possibly false, conclusions.

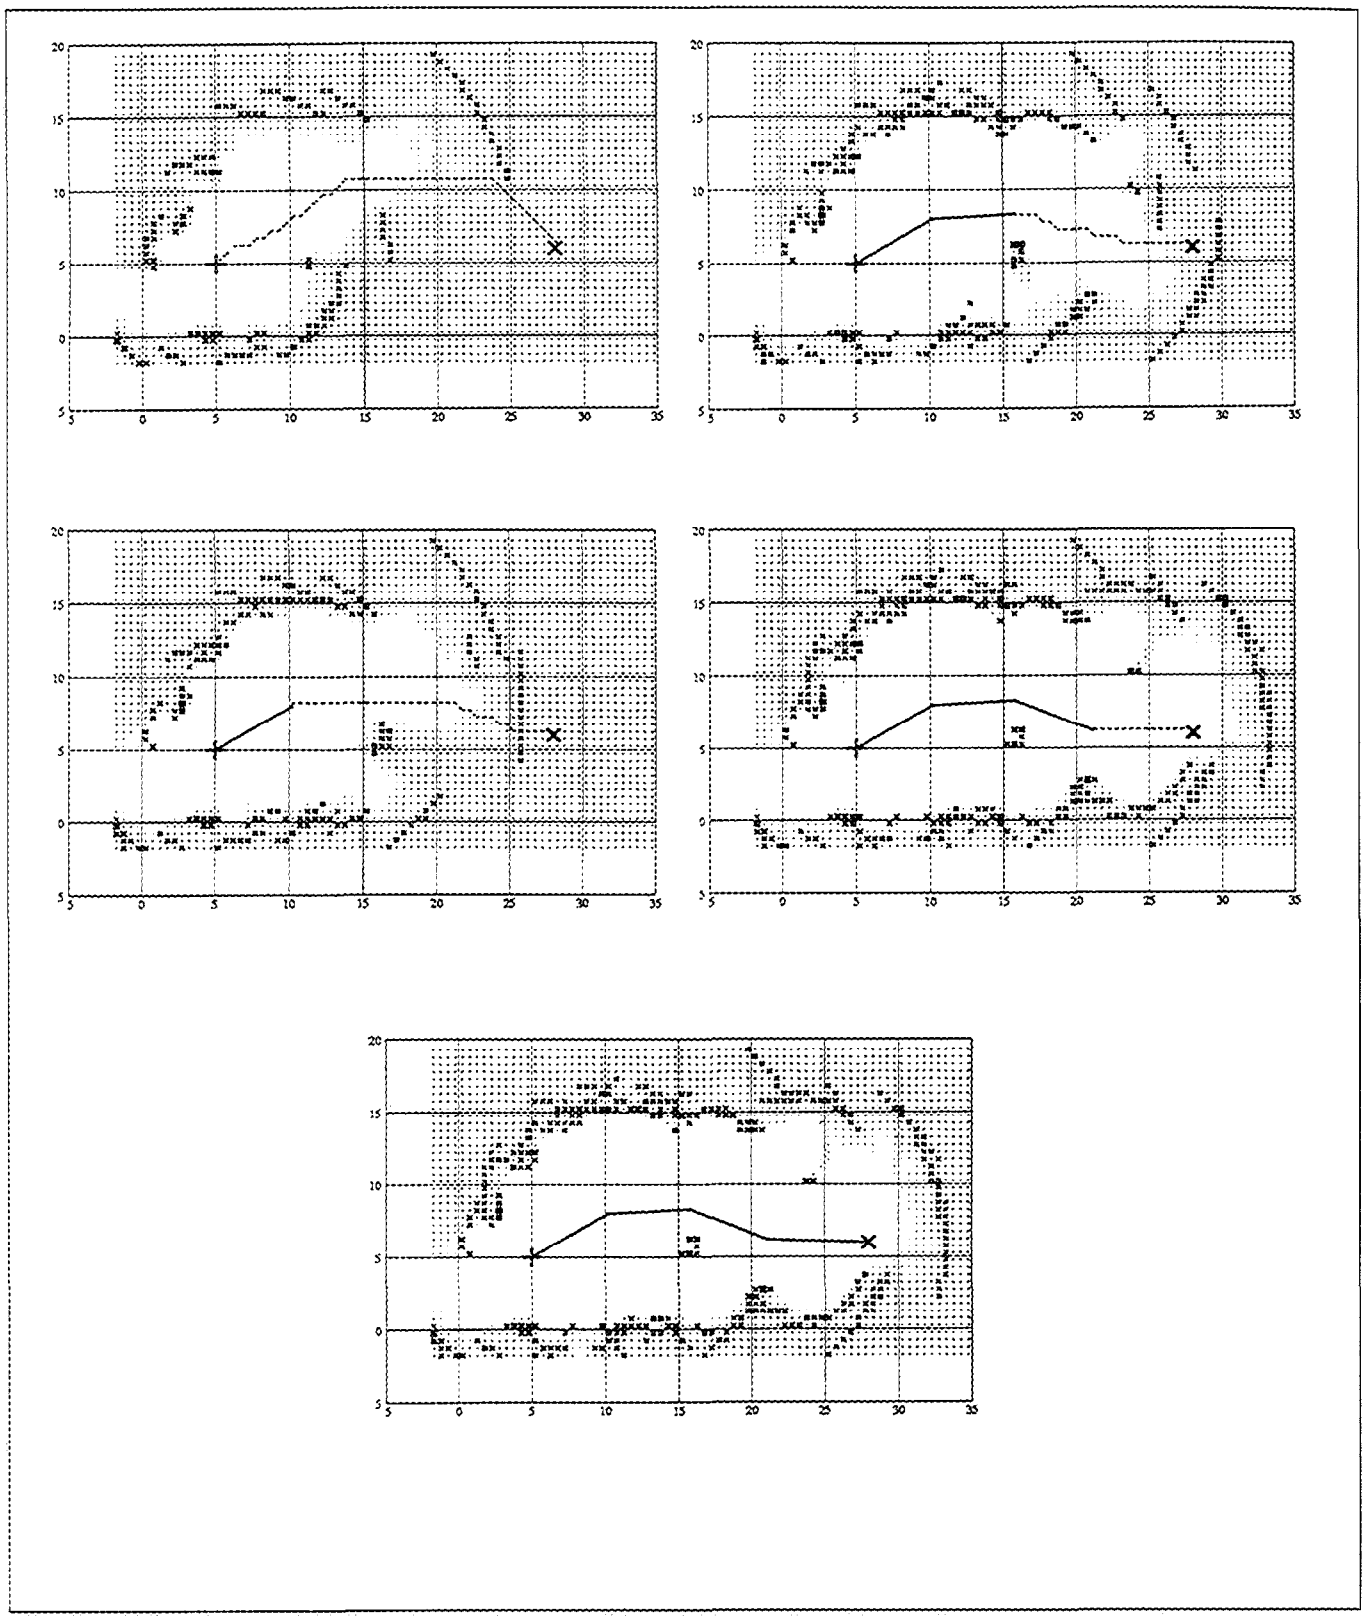

Figure 3. Sonar Mapping and Navigation
The figure illustrates a plan view of the certainty grid built by a sonar-guided robot traversing the MRL laboratory The scale marks are in feet Each point on the dark trajectory is a stop that allowed the on-board sonar ring to collect 24 new readings The grid cells are white if the occupancy probability is low, dots if unknown, and x if high The forward paths were planned by an A algorithm working in the grid as it was incrementally generated*

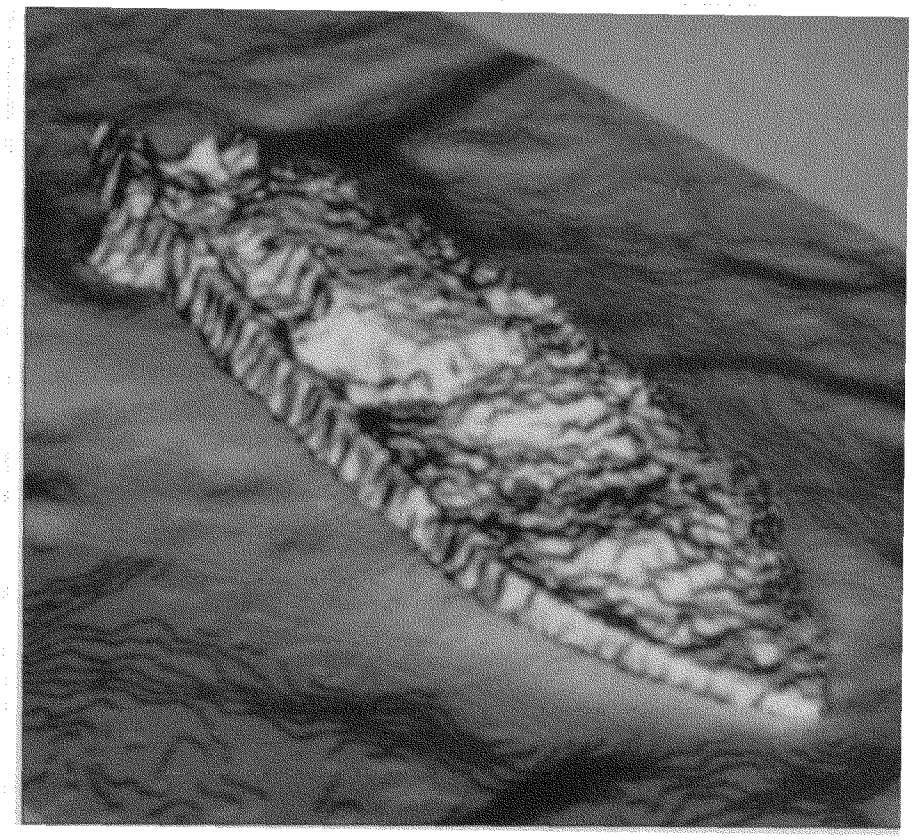

Figure 4.
Three-Dimensional Underwater Sonar Image of USS Monitor.

A surface was extracted from a three-dimensional certainty grid 128 cells on a side The grid was built from about 100,000 readings from a 1 5-degree scanning sonar on a free swimming robot The ship is lying upside down, with many parts of the hull collapsed

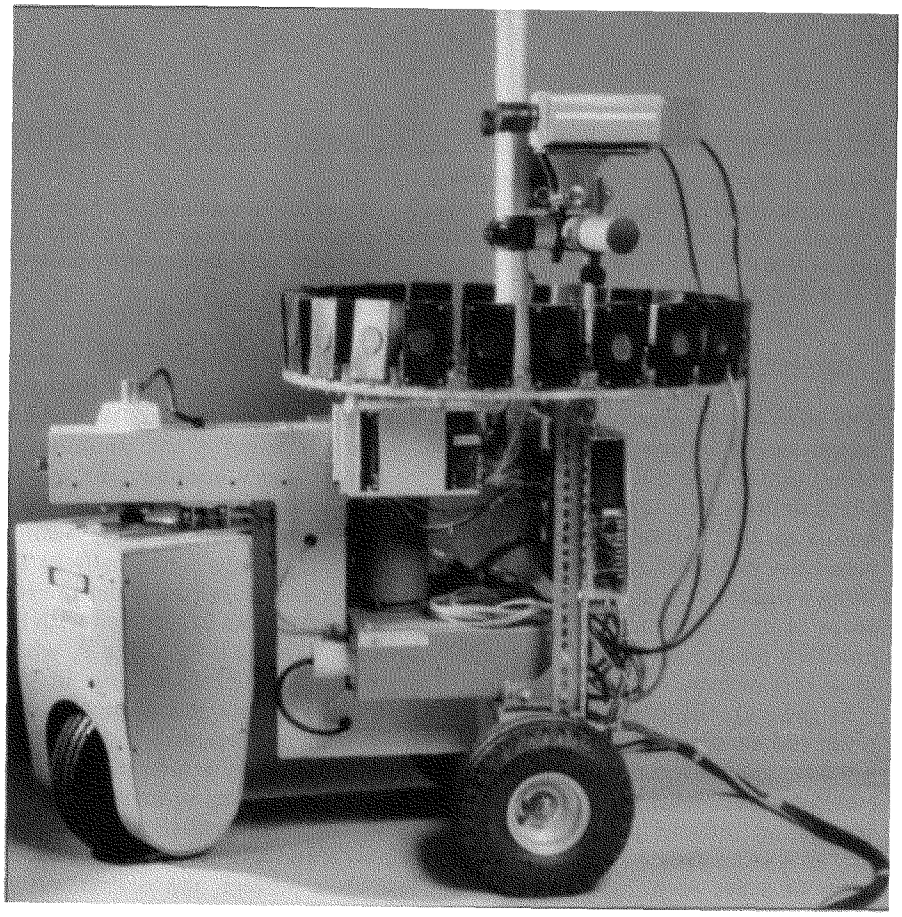

Figure 5.
The Neptune Mobile Robot.

The Host for Many Early Certainty Grid Experiments.

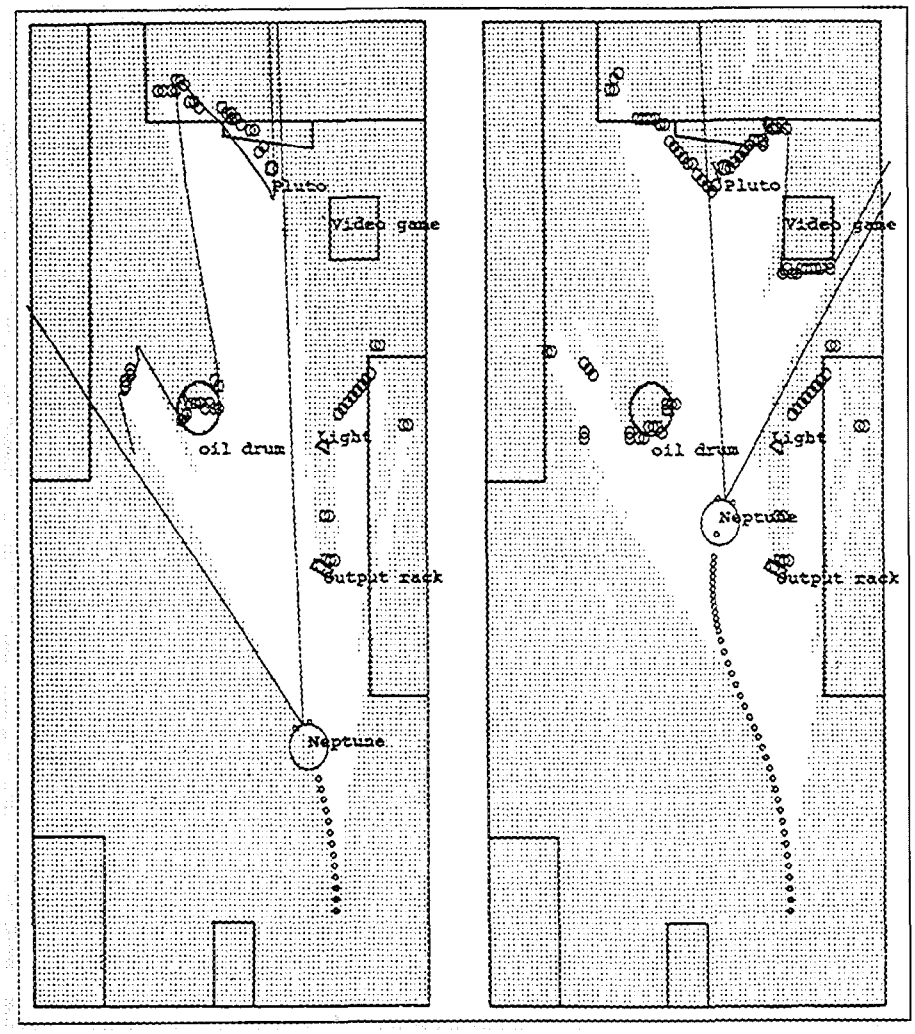

Figure 6 Stereo Mapping and Navigation
The figure illustrates a plan view of the certainty grid built by a stereo-guided robot traversing the MRL laboratory The situation is analogous to the sonar case of figure 3, but the range profiles were gathered from a scan-line stereo method using two television cameras rather than a sonar ring

Inserting Measurements

The readings of almost any kind of sensor can be incorporated into a certainty grid if they can be expressed in geometric terms. The information from a reading can be as minimal as a proximity detector's report that something probably is in a certain region of space or as detailed as a stereo depth profiler's precise numbers on the contours of a surface.

In general, the first step is to express the sensor's measurement as a numeric spatial certainty distribution commensurate with the grid's geometry. For an infrared proximity detector, this measurement can take the form of a set of numbers P_x in an elliptical envelope, with high certainty values in a central axis (meaning detection is likely there) tapering to zero at the edges of the illumination envelope. Let's suppose the sensor returns a binary indication that something is or is not in its field of view. If the sensor reports a hit, cells in the certainty grid C_x falling under the sensor's envelope can be updated with the formula $C_x := C_x + P_x - C_x \times P_x$ which will increase the C values. In this case, the P values should be scaled so their sum is 1 because the measurement describes a situation where there is something somewhere in the field of view, probably not everywhere. If the reliability of the sensor is less than perfect, the normalization can be to a sum less than unity. If, however, the detector registers no hit, the formula might be $C_x := C_x \times (1 - P_x)$ and the Cs will be reduced. In this case, the measurement states there is nothing anywhere in the field of view, and the *P* values should reflect only the chance that an object has been overlooked at each particular position; that is, they should not be normalized. If the sensor returns a continuous value rather than a binary one, perhaps expressing some kind of rough range estimate, a mixed strategy similar to the one described for sonar in the following paragraph is called for.

A Polaroid sonar measurement is a number giving the range of the nearest object within about a 30-degree cone in front of the sonar transducer. Because of the wide angle, the object position is known only to be somewhere on a certain surface. This range surface can be handled in the same manner as the sensitivity distribution of a proximity detector "hit" discussed earlier. The sonar measurement has something else to say, however. The volume of the cone as high as the range reading is probably empty, else a smaller range would have been returned. The empty volume is like the "no hit" proximity detector case and can be handled in the same fashion. Thus, a sonar reading is like a proximity detector hit at some locations, increasing the occupancy probability there, and like a miss at others, decreasing the probability. If we have a large number of sonar readings taken from different vantage points (say, as the robot moves), the gradual accumulation of such certainty numbers will build a respectable map. We can, in fact, do a little better than simply accumulate probabilities.

Imagine two sonar readings whose volumes intersect, and suppose the empty region of the second overlaps part of the range surface of the first. Now, the range surface says, "Somewhere along here there is an object," and the empty volume says, "There is no object here." The second reading can be used to reduce the uncertainty in the position of the object located by the first reading by decreasing the probability in the area of the overlap

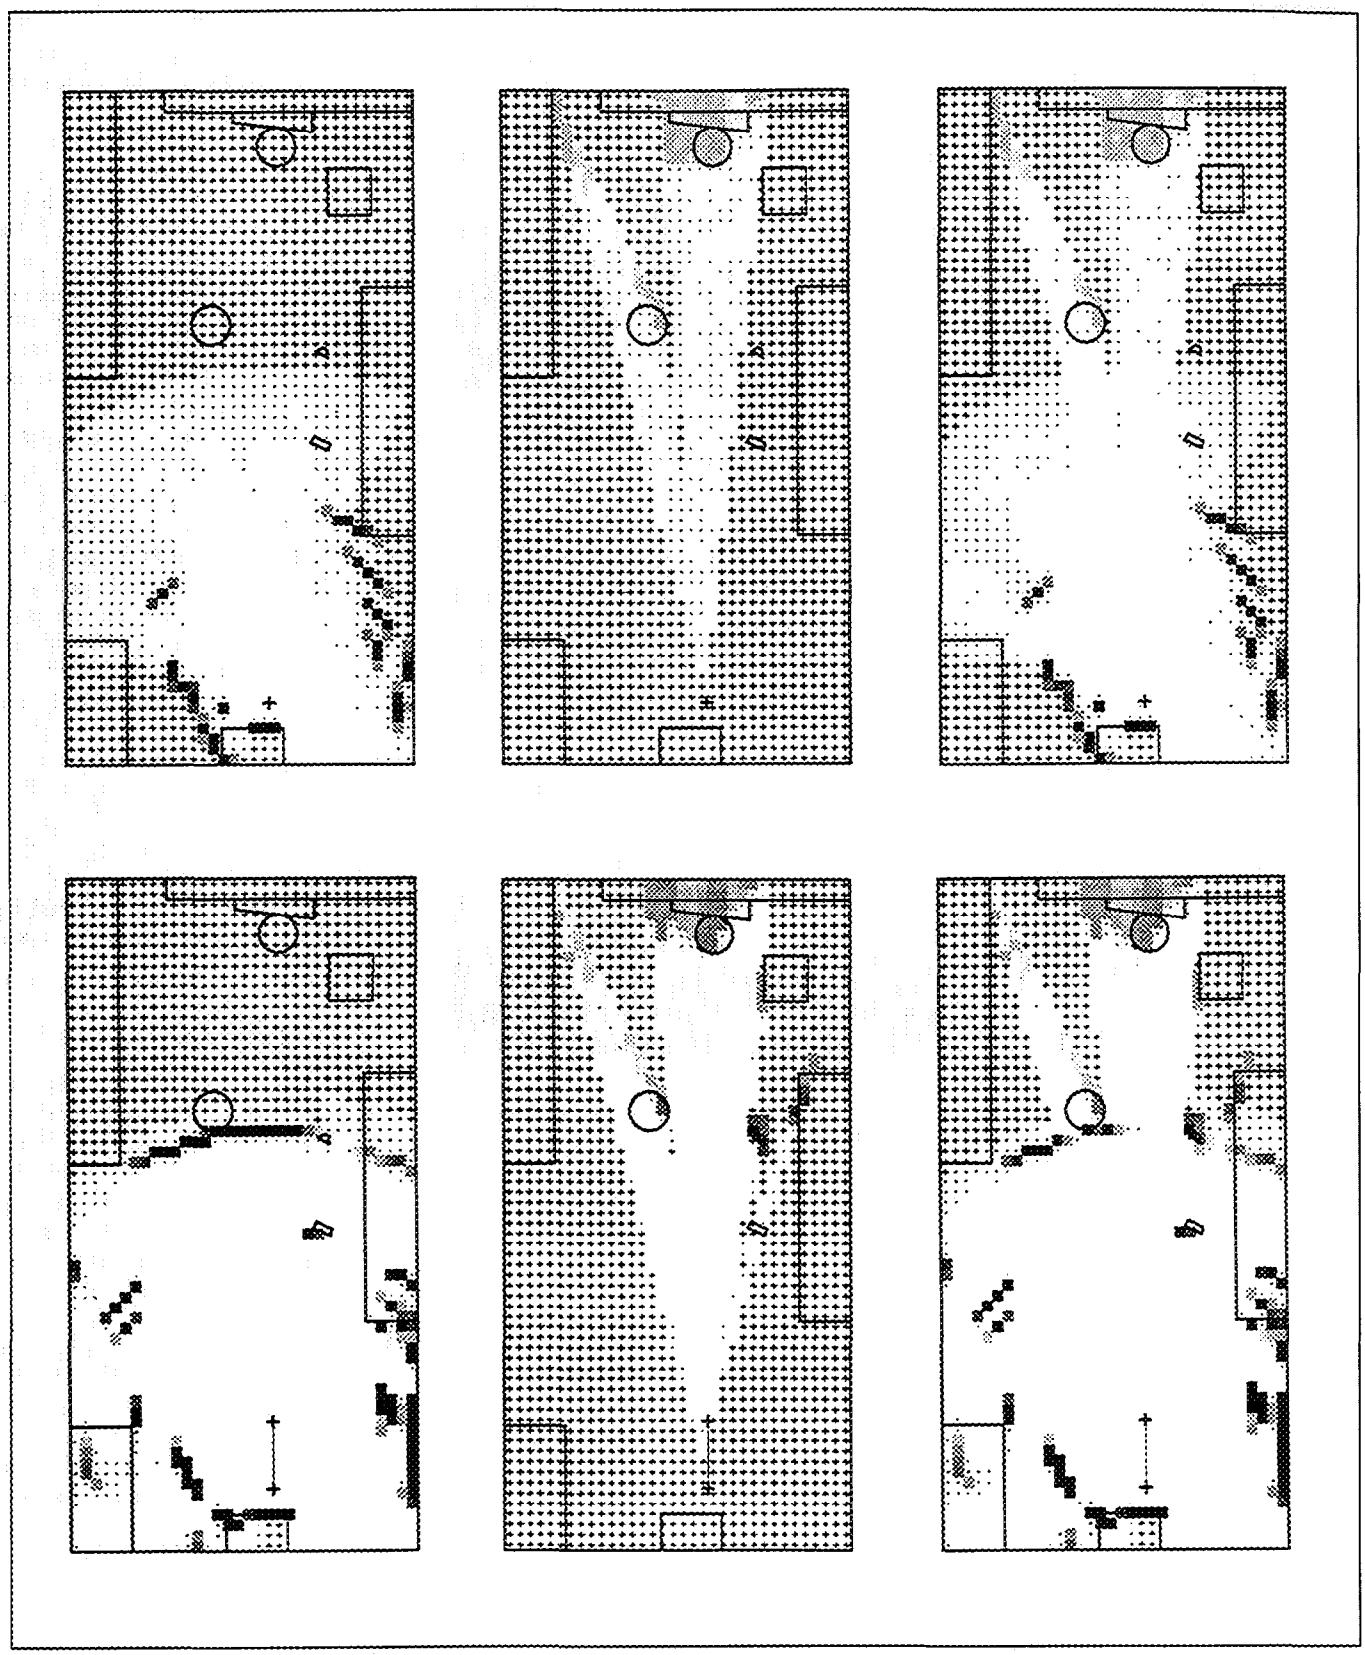

Figure 7. Stereo and Sonar Sensor Integration.
The figure illustrates a plan view of the certainty grids built on the first and tenth steps of a dual sensor run. The leftmost grids contain sonar data only, the center grid has stereo vision only, and the rightmost is the combination of the two Occupied regions are marked by shaded squares, empty areas by dots fading to white, and unknown territory by + signs.

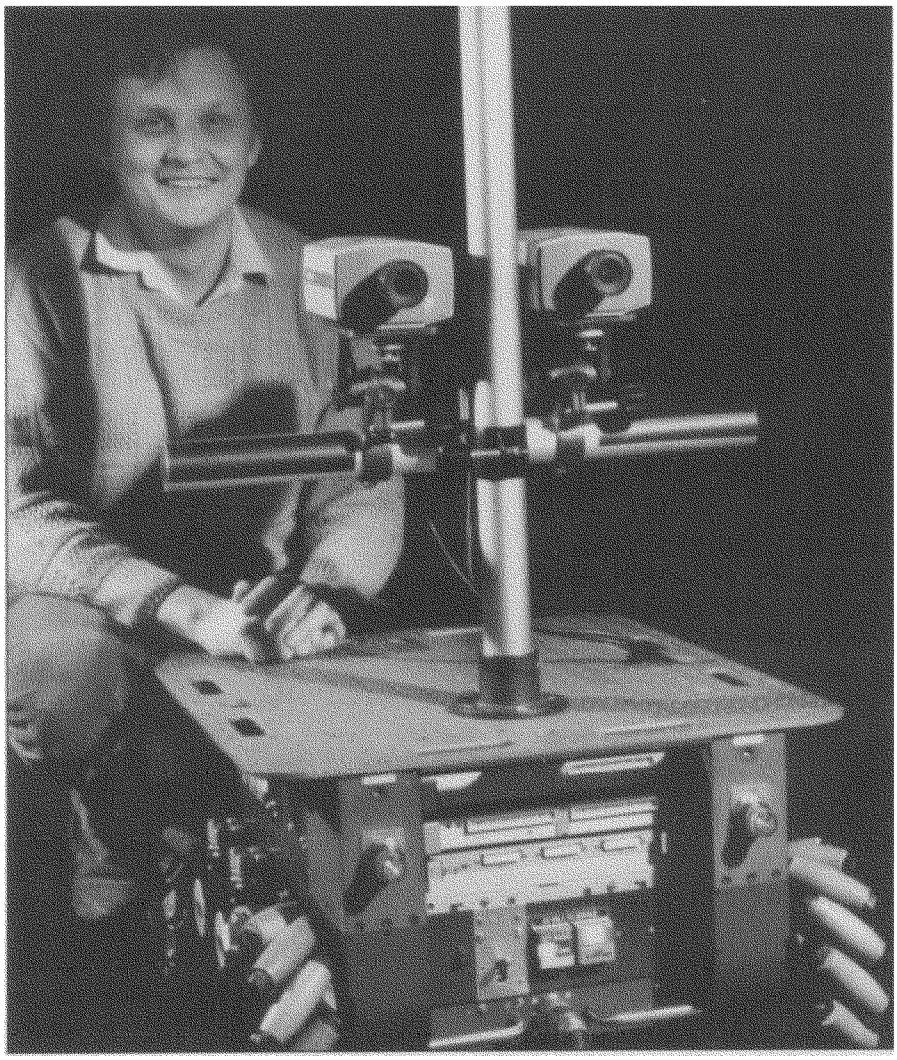

Figure 8 The Uranus Mobile Robot—Great Expectations.

and correspondingly increasing it in the rest of the range surface. This redistribution of range-surface probabilities can be accomplished by reducing the range surface certainties R_x with the formula $R_x := R_x \times (1-E_x)$ where E_x is the empty certainty at each point from the second reading, and then normalizing the *R*s. This method is used to good effect in the existing sonar navigation programs, with the elaboration that the *E*s of many readings are first accumulated and then used to condense the *R*s of the same readings. (It is this two-stage process that led MRL to use two grids in the original programs. In fact, the grid in which the Es are accumulated need merely be temporary working space.)

The stereo method of Serey and Matthies provides a depth profile of visible surfaces. Although, like a sonar reading, it describes a volume of emptiness bounded by a surface whose distance has been measured, it differs by providing a high certainty that there is matter at each point along the range surface. The processing of the empty volume is the same, but the certainty reduction and normalization steps we apply to sonar range surfaces are, thus, not appropriate. The grid cells along a tight distribution around the range surface should simply be increased in value according to the hit formula. The magnitude and spread of the distribution should vary according to the confidence of the stereo match at each point. The method used by Serey and Matthies matches an edge crossing along corresponding scan lines of two images and is likely to be accurate at these points. Elsewhere, it interpolates, and the expected accuracy declines.

If the robot has proximity or contact sensors, its own motion can contribute to a certainty grid. Areas traversed by the robot are almost certainly empty, and their cells can be reduced by the no-hit formula, applied over a confident sharp-edged distribution in the shape of the robot. This approach becomes more interesting if the robot's motion has inherent uncertainties and inaccuracies. If the certainty grid is maintained so it is accurate with respect to the robot's present position (so-called robot coordinates), then the past positions of the robot will be uncertain in this coordinate system. The uncertain relation between past and present coordinate systems can be expressed by blurring the certainty grid accumulated from previous readings in a particular way after each move to reflect the uncertainty in this move. New readings are inserted without blur (essentially the robot is saying, "I know exactly where I am now; I'm just not sure where I was before"). The track in the certainty grid of a moving robot's path in this system will resemble the vapor trail of a high-flying jet—tight and dense in the vicinity of the robot, diffusing eventually to nothing with time and distance.

Extracting Deductions

The purpose of maintaining a certainty grid in the robot is to plan and monitor actions. Thorpe and Elfes showed one way to plan obstacle-avoiding paths. Conceptually, the grid can be considered an array of topographic values—high-occupancy certainties are hills, and low certainties are valleys. A safe path follows valleys such as running water. A relaxation algorithm can perturb portions of a trial path to bring each part to a local minimum. In principle, a decision need never be made about which locations are actually empty and which are occupied, although perhaps the program should stop if the best path climbs beyond some threshold "altitude." If the robot's sensors continue to operate and update the grid as the path is executed, impasses will become obvious as proximity and contact sensors raise the occupancy certainty of locations where they make

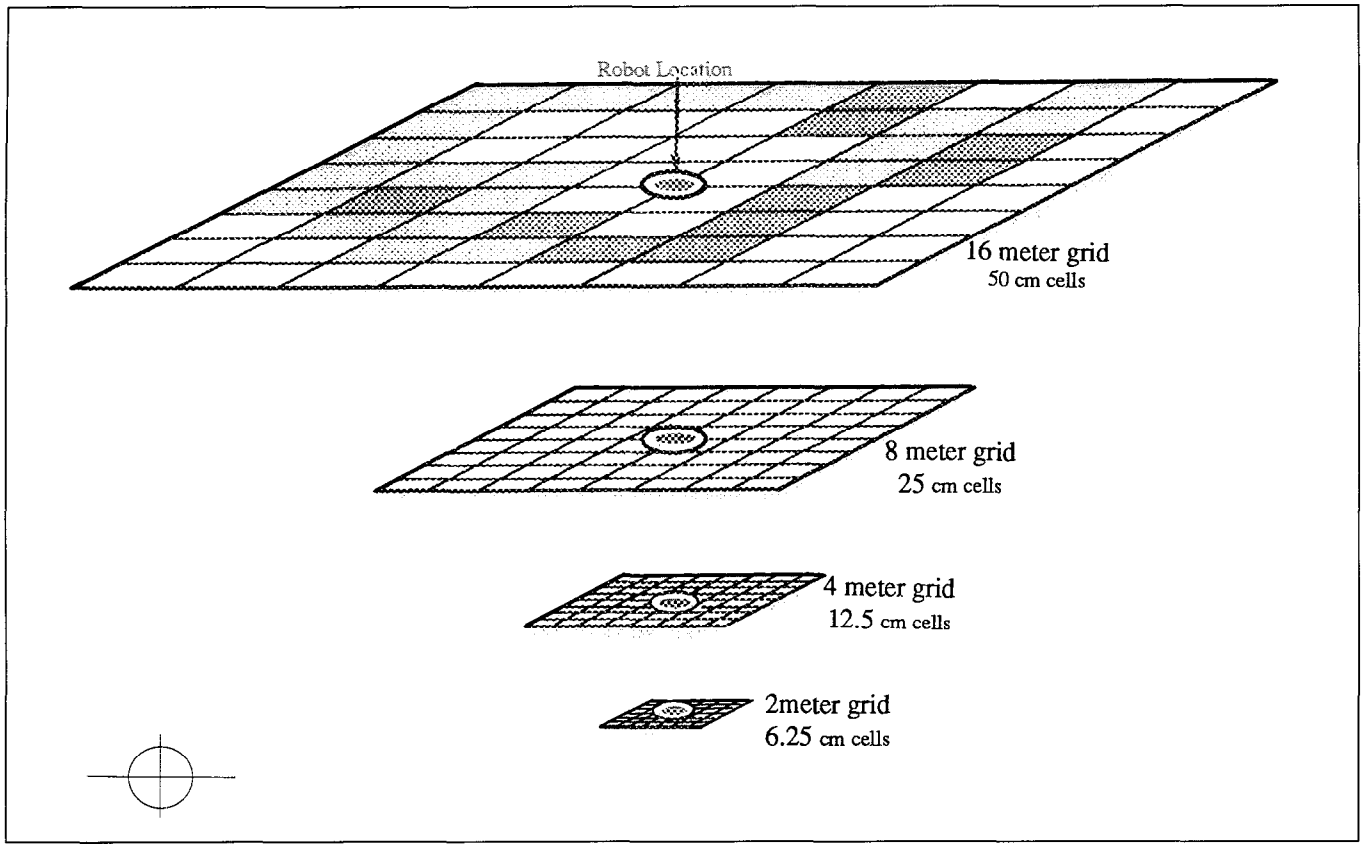

Figure 9. Map Resolution Hierarchy
Coarse maps are for the big picture, fine ones for the fiddly details in the immediate environment All the maps are scrolled to keep the robot in the center cells

contact with solid matter.

As indicated in the beginning of this article, MRL has already demonstrated effective navigation by convolving certainty grids of given locations built at different times, allowing the robot to determine its location with respect to previously constructed maps. This technique can be extended to subparts of maps and can be suitable for recognizing particular landmarks and objects. For instance, MRL is currently developing a wall tracker that fits a least squares line to points which are weighted by the product of the occupancy certainty value and a Gaussian of the distance of the grid points from an a priori guess of the wall location. The parameters of the least squares line are the found wall location and, after being transformed for robot motion, serve as the initial guess for the next iteration of the process.

For tasks that would benefit from an opportunistic exploration of unknown terrain, the certainty grid can be examined to find interesting places to go next. Unknown regions are those whose certainty values are near the background certainty Cb. By applying an operator that computes a function such as $\Sigma (Cx - Cb)^2$ over a weighted window of suitable size, a program can find regions whose contents are relatively unknown and head for them. Other operators similar in spirit can measure other properties of the space and the robot's state of knowledge about it. Hard-edged characterizations of the stuff in the space can be left to the last possible moment by this approach or avoided altogether.

A Plan: Awareness for a Robot

Uranus is MRL's latest and best robot and the third and last one we intend to construct for the foreseeable future. About 60 centimeters square with an omnidirectional drive system intended primarily for indoor work, Uranus carries two racks wired for the industry standard VME computer bus and can be upgraded with off-the-shelf processors, memory, and input output boards (see figure 8). In the last few years, the speed and memory available on single boards have begun to match that available in our mainframe computers. The growth in power of microprocessors removes the main arguments for operating the machine primarily by remote control. With most computing done on board by dedicated processors, enabling high bandwidth and reliable connection of processors to sensors and effectors, real-time control is much easier. Also favoring this change in approach is the realization, growing from experience with robot control programs from the very complex to the relatively simple, that the most complicated programs are probably not the most effective way to learn about programming robots. Very complex programs are slow, limiting the number of experiments possible in any given time, and they involve too many simultaneous variables whose

effects can be hard to separate. A manageable intermediate complexity seems likely to get MRL to long-term goals fastest. The most exciting element in MRL's current plans is a realization that certainty grids are a powerful and efficient unifying solution for sensor fusion, motion planning, landmark identification, and many other central problems.

As the core of the robot and the research, MRL will prepare a kind of operating system based on the certainty grid idea. Software running continuously on processors on board Uranus will maintain a probabilistic, geometric map of the robot's surroundings as it moves. The certainty grid representation will allow this map to be incrementally updated in a uniform way from various sources, including sonar, stereo vision, proximity, and contact sensors. The approach can correctly model the fuzziness of each reading and at the same time combine multiple measurements to produce sharper map features; it can also correctly deal with uncertainties in the robot's motion. The map will be used by planning programs to choose clear paths, identify locations (by correlating maps), identify well-known and insufficiently sensed terrain, and perhaps identify objects by shape. To obtain both adequate resolution of nearby areas and sufficient coverage for longer-range planning without excessive cost, a hierarchy of maps will be kept, the smallest covering a 2-meter area at 6.25-centimeters resolution and the largest 16 meters at 50-centimeters resolution (see figure 9). This map will be "scrolled" to keep the robot centered as it moves, but rotations of the robot will be handled by changing elements of a matrix that represents the robot's orientation in the grid. The map forms a kind of consciousness of the world surrounding the robot; reasoning about the world would actually be done by computations in the map. It might be interesting to take one more step in the hierarchy, to a 1-meter grid that simply covers the robot's own extent. It would be natural to keep this final grid oriented with respect to the robot chassis itself rather than approximately to the compass, as with the other grids. This change of coordinate system would provide a natural distinction between "world" awareness and "body" or "self" awareness. Such encoding of a sense of self might even be useful if the robot were covered with many sensors or, perhaps, were equipped with manipulators. MRL has no immediate plans in this direction and so will pass by this interesting idea for now.

The initial version will contain a pair of two-dimensional grid sets, one mapping the presence of objects at the robots operating height of a few feet above ground level. The other will map the less complex idea of the presence of passable floor at various locations. The object map will be updated from all sensors and the floor map primarily from downward-looking proximity detectors, although possibly also from long-range data from vision and sonar. The robot will navigate by dead reckoning, integrating the motion of its wheels. This method accumulates error rapidly, and this uncertainty will be reflected in the maps by a repeated blurring operation. Old readings, whose location relative to the robot's present position and orientation are known with decreasing precision, will have their effect gradually diffused by this operation until they eventually evaporate to the background certainty value.

It would be natural to extend the two-grid system to many grids, each mapping a particular vertical slice, until we have a true three-dimensional grid. MRL will do this as research results and processing power permits. The availability of single-board array processors that can be installed on the robot would help handle the increased computational demands because the certainty grid operations are amenable to vectorizing. The certainty grid representation can also be extended in the time dimension, with past certainty grids being saved at regular intervals like frames in a movie film, and registered to the robot's current coordinates (and blurred for motion uncertainties). Line operators applied across the time dimension could detect and track moving objects and give the robot a sense of time as well as space. This possibility has some thrilling conceptual (and perceptual) consequences, but MRL might not get to it for a while.

Even the simplest versions of the idea allows to fairly straightforwardly program the robot for tasks that have hitherto been out of reach. MRL looks forward to a program that can explore a region and return to its starting place, using map snapshots from its outbound journey to find its way back, even in the presence of motion disturbances and occasional changes in the terrain. By funneling the sensor readings through a certainty grid, which collects and preserves all the essential data and indications of uncertainties and makes it available in a uniform way, the problem—that for each combination of sensor and task, a different program is required—is avoided. Now, the task execution is decoupled from the sensing and, thus, becomes simpler.

Bayesian Reasoning

In most of MRL's work to date, ad hoc formulas were used to update the certainty grid estimates. Recently, a less arbitrary statistical approach derived from Bayes's theorem (Berger 1985) has captured our attention. Preliminary results using this approach are at least as good as those from the old formulas. Many puzzling aspects of the old scheme have been clarified in the process.

Let p(A|B) represent our best estimate of the likelihood of situation A if we have received information B. By definition,

$$p(A|B) = \frac{p(A \wedge B)}{p(B)} \quad (1)$$

Plain p(A) represents our estimate of A given no special information. The alternative to event A is referred to as ¬A (not A). For any B,

$$p(A|B) + p(\neg A|B) = 1 \quad (2)$$

Certainty grid is a regular finite-element model of space. Each cell of the grid contains a likelihood estimate (our certainty) of a given property of the corresponding region of space. Primarily, we are concerned with simple occupancy of the region, represented by p(o[x]), the probability that region x is occupied. When a discussion contains only one particular x, we drop the subscript and refer simply to p(o).

MRL is considering data derived from wide-angle sonar range measurements. A given measurement is designated M[i], with i the sequential number of the reading. The intersection of a set of readings can be designated by a range in subscript, as in

$$M[<n] = \bigcap_{i=1}^{i<n} M[i] \quad (3)$$

or by a list as in

$$M[i,j,l] = M[i] \wedge M[j] \wedge M[l] \quad (4)$$

When only one reading M[i] enters into a discussion, we abbreviate its name to M. Each measurement has a value, a sonar range R(M). The sonar sensor is quantized, and R is an integer.

P(o) is the probability that any particular cell is occupied, that is, the average occupation density of our space. Our measurements don't give P(o) directly, but it is approximately the overall average of the p(o)'s of all the cells of a typical map of the space By definition, $P(\neg o) = \neg P(o) = 1 - P(o)$

Fundamental Formulas

For two occupancy possibilities o and ¬ o of a cell, new information B and old information A, one form of Bayes's theorem gives (5).

$$p(o|B \wedge A) = \frac{p(B|o \wedge A) \times p(o|A)}{p(B|o \wedge A) \times p(o|A) + p(B|\neg o \wedge A) \times p(\neg o|A)}$$

and (6):

$$p(\neg o|B \wedge A) = \frac{p(B|\neg o \wedge A) \times p(\neg o|A)}{p(B|o \wedge A) \times p(o|A) + p(B|\neg o \wedge A) \times p(\neg o|A)}$$

The odds formulation of equations 5 and 6 is compact and convenient for computation and will be important later:

$$\frac{p(o|B \wedge A)}{p(\neg o|B \wedge A)} = \frac{p(B|o \wedge A)}{p(B|\neg o \wedge A)} \times \frac{p(o|A)}{p(\neg o|A)} \quad (7)$$

Equations 5 and 6 are somewhat complicated for repeated use Equation 7 is better; it is formulated in terms of a product of odds. When the odds ratio involves a probability and its complement, odds and probabilities can be interconverted by the relationship that is shown in the following:

$$Odds(A) = \frac{p(A)}{p(\neg A)} = \frac{p(A)}{1 - p(A)} \quad (8)$$

and

$$p(A) = \frac{Odds(A)}{1 + Odds(A)} \quad (9)$$

The p(B|...) ratio in equation 7 is not of this form. To compute its value, both numerator and denominator must be evaluated separately To make this difference apparent, the ratio $p(B|o \wedge A)/p(B|\neg o \wedge A)$ is referred to as Odds(B||o ^ A). Once the ratio is obtained, however, it can be treated as any other odds number.

If we deal exclusively with odds, all the combining operations become multiplications. Equation 7 is expressed as

$$Odds(o|A \wedge B) = Odds(B||o \wedge A) \times Odds(o|A) \quad (10)$$

An additional transformation further streamlines the computation. Let L(A) represent the logarithm, to some suitable base, of Odds(A) The formula then becomes a simple addition

$$L(o|A \wedge B) = L(B||o \wedge A) + L(o|A) \quad (11)$$

The terms can be integers if the base of the logarithms is chosen well. Perfect certainty (p(o) = 0 and p(o) = 1) can no longer be represented, but such values are probably a mistake in any representation because they are unalterable by any input

A Combining Formula

Bayes's theorem is a formula that combines independent sources of information A and B into an estimate of a single quantity p(o|A ^ B). The new information B occurs in terms of the probability of B in the situation that (in our case of interest) a particular cell is or is not occupied, p(B|o) and p(B|¬o). This inversion is central to the usefulness of Bayes' theorem. However, consider the problem of generating a map from information A and B when each has already been individually processed into a map; that is, find p(o|A ^ B) given p(o|A) and p(o|B).

Bayes's formula (equation 7) applied to information B and null A (that is, only global information from the A side) makes the relationship between p(o|B) and p(B|o) clear .

$$\frac{p(o|B)}{p(\neg o|B)} = \frac{p(B|o)}{p(B|\neg o)} \times \frac{P(o)}{P(\neg o)} \quad (12)$$

thus,

$$\frac{p(B|o)}{p(B|\neg o)} = \frac{p(o|B)}{p(\neg o|B)} \times \frac{P(\neg o)}{P(o)} \quad (13)$$

It would be nice to substitute this ratio back into the original version of Bayes's formula (equation 7) to produce a formula giving p(o|A ^ B) in terms of p(o|A) and p(o|B), thus allowing measurements A and B to be incorporated into maps independently and combined afterward. Such independent insertion is not possible in general because the two measurements might interact in some way; for example, either A or B alone might indicate a high probability of o, but together, they might confirm some other hypothesis and reduce the probability of o. If, however, we make a strong assumption of independence of B from A:

$$\frac{p(B|o \wedge A)}{p(B|\neg o \wedge A)} = \frac{p(B|o)}{p(B|\neg o)} \quad (14)$$

we can use equations 7, 12, and 14 to produce a map-combining formula (15):

$$\frac{p(o|A \wedge B)}{p(\neg o|A \wedge B)} = \frac{p(o|B)}{p(\neg o|B)} \times \frac{p(o|A)}{p(\neg o|A)} \times \frac{P(\neg o)}{P(o)}$$

Although the noninformative reading in equation 7, that which leaves the original p(o) estimate unchanged, is found when p(B|o)/p(B|¬o) = 1, the noninformative case in equation 15 happens when p(o|B)/p(¬o|B) = P(o)/P(¬o), that is, when the cell density estimated from the reading is the same as the average cell density of the whole map

Formula 15 is most important because it provides a means for combining maps of the same area obtained by different means, for example, by independent scans of different sensors. It can also be used to incorporate individual sensor readings as we build a map, but in general, it is inferior to Bayes's formula (equation 7) for this purpose because it precludes the use of any knowledge we might have of sensor interactions. In odds and log odds form, it becomes (16).

$$Odds(o|A \wedge B) = \frac{Odds(o|A) \times Odds(o|B)}{Odds(o)}$$

and

$$L(o|A \wedge B) = L(o|A) + L(o|B) - L(o) \quad (17)$$

Sonar Wedges

Simple reasoning from first principles can produce estimates for p(M|o) and p(M|¬o) for a sonar reading These values can be used to update maps by direct application of Bayes's formula (equation 7). In Measurement Errors, how to systematically incorporate in the possibility of errors in the readings is shown. For now, assume that the sensor always perfectly returns the range to the nearest occupied cell within its angle of sensitivity. Define the sonar regions according to the diagram shown in figure 10.

When incorporating a new reading B, the map built from prior readings A can be used to help provide an estimate for the p(B|o ^ A) and p(B|¬o ^A) of equation 7

Let P(R) be the probability prior to the reading that the next sonar measurement will result in a given range R P(R) can be approximated by stepping through the possible ranges R_1, R_2, ., R_n, starting with the shortest, R_1, and multiplying the occupancy probability p(o|A) of each cell on the range surface for the range by the probability that the sonar would detect an occupied cell at the location. The sum of these products at R_1, call it S(R_1), is P(R_1). Now P(R_2) = S(R_2) x (1 - P(R_1)) because an echo from R_2can happen only if the sonar pulse has not already been intercepted at the shorter range R_2 In general,

$$P(R_i) = S(R_i) \times (1 - \sum_{j<i} P(R_j)) \qquad (18)$$

By definition, Σ P(R) over all R is unity. I will suggest later exactly how to compute S and how the detection probabilities required in the calculation —equation 18—can be determined empirically by collecting statistics from many maps of the correlation of individual sensor readings with the composite maps they helped create.

External Region

Consider a cell in the external region. p(M|o) is our estimate that a measurement range of R(M), as opposed to some other range, will occur if we happen to know only that the cell is occupied. Because the cell is outside the sonar cone, its state of occupancy has no effect on the reading, and we can refer only to the uniform range distribution to conclude that

$$p(M|o) = p(M|\neg o) = P(R(M)) \qquad (19)$$

and

$$p(M|o) / p(M|\neg o) = 1 \qquad (20)$$

Inserting equations 19 and 20 into equation 7 gives (21)

$$\frac{p(o|M[1:i])}{p(o|M[1:i])} = \frac{P(R(M))}{P(R(M))} \times \frac{p(o|M[<i])}{p(\neg o|M[<i])}$$

$$= \frac{p(o|M[<i])}{p(\neg o|M[<i])}$$

so, the occupancy certainty is unchanged, as it should be, because the sonar reading contains no information about the external cell

Range Surface

Now, consider a cell on the range surface, and suppose this surface covers n cells in all. By definition, the range at this surface is R(M).

If the cell is occupied, then a perfect sensor would detect it and, thus, could not return a greater range than R(M). All ranges beyond the occupied cell would thus be short-circuited by the cell. It would, however, be possible for it to return shorter ranges if closer cells within the sonar beam angle happened to be occupied. The probabilities of the shorter ranges should not be changed by the presence of the farther occupied cell The probability of getting just the range reading would be unity minus the a priori probability of getting a lesser reading or, equivalently, the a priori probability of getting a reading greater than or equal to R(M):

$$p(M|o) = 1 - \sum_{R<R(M)} P(R) = \sum_{R \geq R(M)} P(R) \qquad (22)$$

If, however, the cell were unoccupied, the original distribution is hardly altered except that the particular R = R(M) is slightly less likely to occur because one possible way to achieve it is eliminated. The cell in question is one of "n" cells we assumed on the range surface; so, we can say the chance of R(M) happening is reduced by the factor (n-1)/n The probabilities of other ranges are not affected directly; so, the result is a slight "notching" at R = R(M) of the a priori P(R) distribution Simply notching it, however, reduces its total area, which must then be renormalized by the proper factor to restore the area of the distribution to unity, increasing all the other probabilities slightly (23):

$$p(M|\neg o) = \frac{P(R(M)) \times (n-1)/n}{1 - P(R(M))/n}$$

$$= P(R(M)) \times \frac{n-1}{n - P(R(M))}$$

Interior Region

If we assume a perfect sensor, an occupied cell in the wedge at a range less than R(M) would always be detected and, thus, would prevent the reading M. Thus, in this case,

$$p(M|o) = 0 \qquad (24)$$

If the cell in the interior is unoccupied, then the range Rc at which the unoccupied cell occurs is less likely than other ranges. Suppose there would be k cells occupied by a range arc at distance Rc. With reasoning such as in equation 23, the probability of Rc would be reduced by a factor of (k-1)/k, but the overall distribution would have to be normalized by divid-

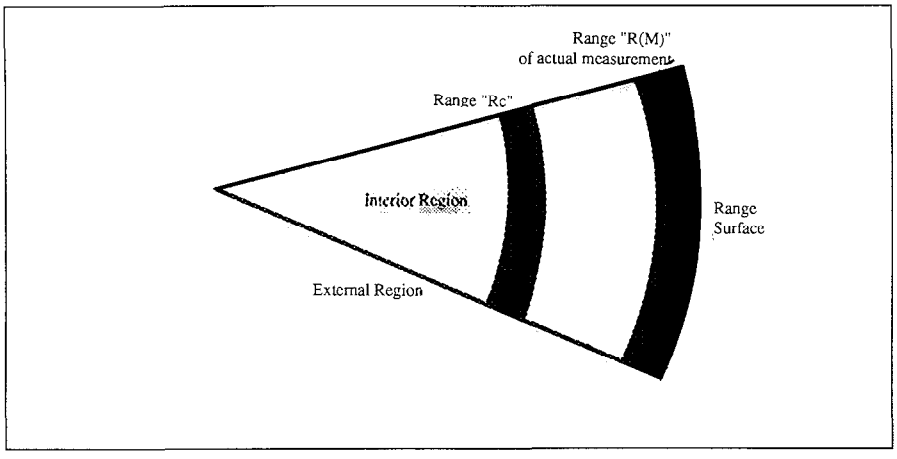

Figure 10. Example of a Sonar Wedge

ing by (1 - P(Rc) / k). This normalization raises the probability of R(M) from its a priori value:

$$p(M|\neg o) = \frac{P(R(M))}{1 - P(Rc) / k} \quad (25)$$

Measurement Errors

Of course, MRL's sensors are not perfect in the sense of those described in the previous subsection. Sometimes, they fail to respond to an occupied location, and at other times, they give a spurious indication of occupancy. It is easy to modify the perfect case formulas for these imperfections. Suppose a small chance e_{hit} exists that an empty cell will act as if it were occupied. Suppose also there is a small possibility e_{miss} that an occupied cell will fail to be detected. We can then construct the formulas with error from the error-free cases using the following general formulas:

$$p(M|o)_{error} = p(M|o) \times (1-e_{miss}) + p(M|\neg o) \times e_{miss} \quad (26)$$

and

$$p(M|\neg o)_{error} = p(M|\neg o) \times (1-e_{hit}) + p(M|\neg o) \times e_{hit} \quad (27)$$

If $e_{hit} = e_{miss} = 1/2$, then the sensors are returning random readings. Random readings are captured in equations 26 and 27 as $p(M|o)_{error} = p(M|\neg o)_{error}$, the noninformative case, which is similar to equation 20. As e_{hit} and e_{miss} grow closer to 1/2, the reading becomes less and less informative. If $e_{hit} = e_{miss} = 1$, the sensor is perfect, but the hit-and-miss indications are mistakenly swapped. In this case, equations 26 and 27 restore the correct pairing.

A real sensor can be modeled, somewhat redundantly, by varying e_{hit} and e_{miss} over the sensed area. Toward the edges of the beam, they creep toward 1/2, achieving this value perfectly outside the beam and perhaps at the boundary of the interior and range surface regions. If $0 \leq e_{hit}, e_{miss} \leq 1$, the error formulas can only reduce or leave unchanged the amount of information provided by the perfect estimates of p(M|o) and p(M|¬o).

The quantities e_{hit} and e_{miss} must be initialized to some reasonable values when a system begins operation. They can then be adjusted by an iterative learning process. Each time a map is constructed from many sensor readings, the correlation of each cell in each individual sensor profile with the area of the total map it overlays is recorded. This correlation is averaged over many maps. If the map often reports a high-occupancy probability in a location where the sensor profile indicates occupancy, e_{hit} will be low for this position in the sensor profile; otherwise, it will be high. Similarly, if the map usually has low probabilities at a site indicated empty in the sensor profile, e_{miss} will be low; otherwise, it will be high.

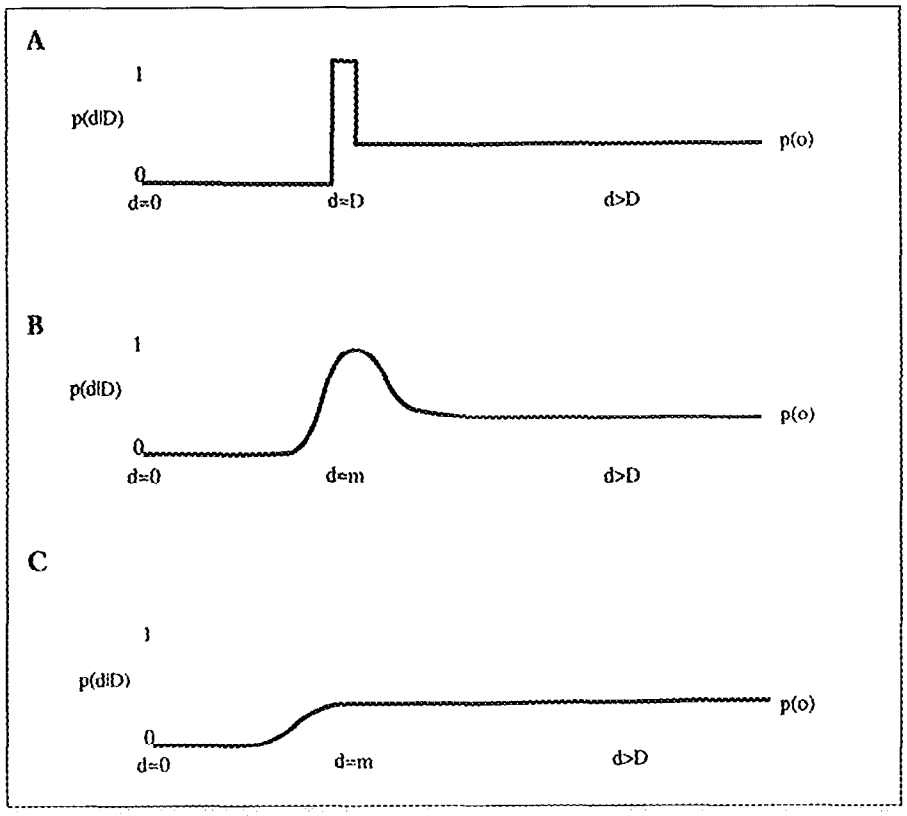

Figure 11

The detection probability required to compute the function S in equation 18 is given by the expression $1-e_{miss}$. A more accurate computation of S would also take into account the chance that an empty cell would falsely register as occupied. Thus, a good estimate for S is

$$S(R) = \Sigma_R p(o|A) \times (1 - e_{miss}) = p(\neg o|A) \times e_{hit} \quad (28)$$

Stereo Range Measurements

Another approach to incorporating individual sensor readings uses the independent combining formula, equation 15. In general, this approach results in inferior maps, but it requires less computation. The following paragraph presents an example from a consideration of a stereo vision measurement

One of MRL's stereo approaches matches edges crossing a given scan line in the left and right images. A range is deduced from the relative horizontal shift of each edge between the two images. The edges cannot be located with infinite precision. Thus, there is some uncertainty in depth, given by a distribution P(D|m); the probability the object ranged by stereo measurement m is actually at distance D. Note that P(D|m) is not the probability that location d is occupied because d can contain an object other than the edge being ranged. Let p(d|m) represent the probability, given measurement m, that d is occupied by any object. We can directly estimate p(d|m) by the following reasoning.

Consider the possible case that the ranged edge is at a particular distance D (this possibility has probability P(D|m)). In this case, assuming a perfect sensor as we did in the sonar

example, we can conclude that

$$p(d|D) = 0 \text{ when } d < D \quad (29)$$

because a perfect sensor would detect an intervening object

$$p(d|D) = 1 \text{ when } d = D \quad (30)$$

because the object is at D:

$$p(d|D) = p(o) \text{ when } d > D \quad (31)$$

Because the edge blocks the view behind it, the measurement gives us no special information beyond d = D.

This is shown in graphic form in figure 11a.

The relationship represented in equations 29 to 31 holds only if the edge is exactly at distance D. Measurement m tells us only there is a probability P(D|m) that such is the case. To get the overall probability, we must sum the p(d|D)s over all possible Ds, each weighted by its probability of actually being the case. Thus

$$p(d|m) = \sum_D (p(d|D) \times p(D|m)) \quad (32)$$

Equation 32 will generally have the form shown in figure 11b.

Although if the spread of p(D|m) is large, the hump at d = m will be attenuated, and the curve will have the approximate shape of figure 11c.

The values of p(d|m) can be used directly to update maps using the combining formula (equation 15).

Addendum

MRL recently discovered a general formulation that captures and improves on the ideas of equations 21, 23, 25, 28, and 32. It results in a single procedure that can process a broad class of range readings, of any dimensionality, at a cost linear in the volume of the sensor footprint and quadratic in its volume of range uncertainty. See Moravec and Cho (1988).

Acknowledgment

This work has been supported since 1981 by the Office of Naval Research, under contract N00014-81-K-503 Alberto Elfes and Larry Matthies made key contributions to the development of these ideas, and I thank them warmly I also wish to thank Peter Cheeseman for setting my nose firmly in the Bayesian groove

References

Berger, J O 1985 *Statistical Decision Theory and Bayesian Analysis* New York: Springer-Verlag

Elfes, A E, 1987 A Sonar-Based Mapping and Navigation System *IEEE Journal of Robotics and Automation* RA-3(3): 249-266.

Kadonoff, M ; Benayad-Cherif, F ; Franklin, A ; Maddox, J ; Muller, L ; Sert B ; and Moravec, H. 1986 Arbitration of Multiple Control Strategies for Mobile Robots In *Mobile Robots Proceedings of the SPIE Conference on Advances in Intelligent Robotics Systems*, W. J Wolfe and N Marquina vol. 727 Bellingham, Wash : International Society for Optical Engineering

Kweon, I S 1987 Personal communication, Robotics Institute, Carnegie-Mellon University, Pittsburgh, Pa , October 1987

Matthies, L H., and Elfes, A 1987. Sensor Integration for Robot Navigation: Combining Sonar and Stereo Range Data in a Grid-Based Representation In Proceedings of the Twenty-sixth IEEE Decision and Control Conference, Los Angeles, 9-11 December 1987

Moravec, H P 1981 *Robot Rover Visual Navigation* Ann Arbor, Mich : UMI Research Press

Moravec, H. P, and Cho, D. W 1988 A Bayesian Method for Certainty Grids Robotics Institute, Mobile Robot Laboratory, Carnegie-Mellon Univ Forthcoming

Moravec, H P, and Elfes, A E 1985 High Resolution Maps from Wide Angle Sonar In Proceedings of the 1985 IEEE International Conference on Robotics and Automation, 116-121 Washington D C.: IEEE Computer Society

Serey, B , and Matthies, L H. 1986 Obstacle Avoidance Using 1-D Stereo Vision, Technical Report, Robotics Institute, Mobile Robot Laboratory, Carnegie-Mellon Univ.

Stewart, W K 1988 A Model-Based Approach to 3-D Imaging and Mapping Underwater In Proceedings of the Seventh International Conference and Exhibit on Offshore Mechanics and Arctic Engineering (OMAE), Houston, Tex., 7-12 February 1988

Stewart, W K 1987 A Non-Deterministic Approach to 3-D Modeling Underwater Presented at the Fifth Symposium on Unmanned Untethered Submersible Technology, University of New Hampshire, Durham, N H , June 1987

Thorpe, C E. 1984 Path Relaxation: Path Planning for a Mobile Robot In Proceedings of the Third National Conference on Artificial Intelligence, 318-321 Los Altos, Calif : Morgan Kaufmann