Using Occupancy Grids for Mobile Robot Perception and Navigation

Alberto Elfes

Carnegie Mellon University

The occupancy grid framework provides a robust and unified approach to a variety of problems in spatial robot perception and navigation.

To widen the range of application and deployment of robots, both in research and in industrial contexts, we need to develop more powerful and flexible robotic systems exhibiting higher degrees of autonomy and able to sense, plan, and operate in unstructured environments. For that, the robot must be able to interact coherently with its world, both by being able to recover robust and useful spatial descriptions of its surroundings using sensory information and by efficiently utilizing these descriptions in appropriate short-term and long-term planning and decision-making activities.

This article reviews a new approach to robot perception and world modeling that uses a probabilistic tesselated representation of spatial information called the *occupancy grid*.[1] The occupancy grid is a multidimensional random field that maintains stochastic estimates of the occupancy state of the cells in a spatial lattice. To construct a sensor-derived map of the robot's world, the cell state estimates are obtained by interpreting the incoming range readings using probabilistic sensor models. Bayesian estimation procedures allow the incremental updating of the occupancy grid using readings taken from several sensors over multiple points of view.

The occupancy grid framework represents a fundamental departure from traditional approaches to robot perception and spatial reasoning. By utilizing probabilistic sensor models and representation schemes, this approach supports the development of agile and robust sensor interpretation mechanisms, incremental discovery procedures, explicit handling of uncertainty, multisensor composition of information, and spatial reasoning tasks within an integrated framework.

The following sections give an overview of the occupancy grid framework and illustrate its application to a number of problems in the mobile robot domain, including range-based mapping, multiple-sensor integration, path planning and navigation, handling of sensor position uncertainty due to robot motion, and related tasks. I contrast the occupancy grid framework to geometric approaches to sensor interpretation and suggest that a number of robotic tasks can be performed directly on the occupancy grid representation. I conclude with an overview of further research.

Spatial sensing and modeling for robot perception

One of the long-term goals of the research discussed in this article has been the development of robust mapping and navigation systems for mobile robots operating in and exploring unstructured and unknown environments.

Such scenarios occur in a variety of contexts. Robot rovers being developed for planetary and space exploration, or autonomous submersibles devoted to submarine prospecting and surveying, have to deal with unexpected circumstances and require the ability to handle complex and rough environments with little or no prior knowledge of the terrain. While planetary

rovers may take advantage of terrain maps obtained from orbiting surveyors for global planning strategies, these will be of limited resolution and not useful for detailed path planning and navigation.

On the other hand, mobile robots developed for factory automation purposes or for operation in hazardous mining environments or nuclear facilities generally can be expected to operate in more constrained situations and to have access to precompiled maps derived from plant blueprints. However, such maps may become outdated. Additionally, over the long distances traversed by autonomous vehicles, inertial or dead-reckoning navigation schemes may accumulate substantial positional errors. This makes it difficult for the robot to position itself in precompiled world models, to register sensor information to an absolute frame of reference, or to construct global maps that are precise in Cartesian coordinates.

These considerations lead to some fundamental requirements for mobile robots. Autonomous vehicles must rely heavily on information recovered from sensor data and must be able to operate without precompiled maps. Sensor views obtained from multiple sensors and different locations have to be integrated into a unified and consistent world model, and sensor uncertainty and errors have to be handled. Precompiled maps, when available, should be used to complement sensor-derived maps. Finally, the positional drift of the sensors due to the robot motion has to be taken into account in the mapping and navigation procedures.

Traditional approaches to sensor interpretation for robot perception have largely relied on the recovery and manipulation of geometric world models.[1] Low-level sensing processes extract geometric features such as line segments or surface patches from the sensor data, while high-level sensing processes use symbolic models, geometric templates, and prior heuristic assumptions about the robot's environment to constrain the sensor interpretation process. The resulting geometric world models serve as the underlying representation for other robotic tasks, such as obstacle avoidance, path planning and navigation, or planning of grasping and assembly operations.

These approaches, which as an ensemble characterize what we refer to as the *geometric paradigm* in robot perception, have several shortcomings.[1] Generally speaking, the geometric paradigm leads to sparse and brittle world models; it requires early decisions in the interpretation of the sensor data for the instantiation of specific model primitives; it does not provide adequate mechanisms for handling sensor uncertainty and errors; and it relies heavily on the adequacy of the precompiled world models and the heuristic assumptions used, introducing strong domain-specific dependencies. Better descriptions of the robot's environment are derived primarily from the application of finer tuned prior models and additional constraints to the available sensor data, rather than from strategies based on additional sensing.

Because of these shortcomings, the geometric paradigm implicitly creates a wide gap between two informational layers: the layer that corresponds to the imprecise and limited information actually provided by the sensor data, and the layer of abstract geometric and symbolic world models operated on by the sensing and world modeling processes. Consequently, geometric approaches to robot perception may be useful in highly structured domains, but have limited applicability in more complex scenarios, such as those posed by mobile robots.

Occupancy grids

The occupancy grid framework addresses the requirements and concerns outlined above through the development of spatial robot perception and reasoning mechanisms that employ probabilistic sensor interpretation models and random field representation schemes. In so doing, it supports robust mapping and navigation strategies and allows a variety of robotic tasks to be addressed through operations performed directly on the occupancy grid representation.

This section provides a brief overview of the occupancy grid formulation, while the following sections illustrate the application of occupancy grids to the mobile robot mapping and navigation domain. The actual derivation of the probabilistic estimation models used is beyond the scope of this article and can be found elsewhere,[1,2] as can more detailed discussions of the experimental work.[1-4]

Occupancy grid representation. The occupancy grid representation employs a multidimensional (typically 2D or 3D) tesselation of space into cells, where each cell stores a probabilistic estimate of its state. Formally, an *occupancy field O*(**x**) is a discrete-state stochastic process defined over a set of continuous spatial coordinates $\mathbf{x} = (x_1, x_2, ..., x_n)$, while the occupancy grid is a lattice process, defined over a discrete spatial lattice. The state variable $s(C)$ associated with a cell C of the occupancy grid is defined as a discrete random variable with two states, *occupied* and *empty*, denoted OCC and EMP. Consequently, the occupancy grid corresponds to a discrete-state binary random field.[5] Since the cell states are exclusive and exhaustive, $P[s(C) = \text{OCC}] + P[s(C) = \text{EMP}] = 1$. More general models are possible by using a random vector and encoding multiple properties in the cell state. I refer to these representations as *inference grids*.[1] This article discusses the estimation of a single property, the occupancy state of each cell.

Estimating the occupancy grid. Since a robot can only obtain information about its environment indirectly, through its sensors, the recovery of a spatial world model from sensor data is best modeled as an estimation theory problem. The specific steps involved in estimating the occupancy grid from sensor readings are sketched out in Figure 1.

To interpret the range data obtained from a given sensing device, we use a stochastic sensor model defined by a probability density function of the form $p(r \mid z)$, which relates the reading r to the true parameter space range value z. This density function is subsequently used in a Bayesian estimation procedure to determine the occupancy grid cell state probabilities. Finally, we can obtain a deterministic world model, using optimal estimators such as the maximum a posteriori (MAP) decision rule to assign discrete states to the cells, labeling them occupied, empty, or unknown. We emphasize, however, that many robotic tasks can operate directly on the occupancy grid representation.

In the discussion below, the occupancy grid is modeled as a Markov random field (MRF)[5] of order 0, so the individual cell states can be estimated as independent random variables. We can employ computationally more expensive estimation procedures for higher order MRFs.

To allow the incremental composition of sensory information, we use the sequential updating formulation of Bayes' theorem to determine the cell occupancy probabilities.[1] Given a current estimate of the state of a cell C_i, $P[s(C_i) = \text{OCC} \mid \{r\}_t]$, based on observations $\{r\}_t = \{r_1, ..., r_t\}$ and given a new observation r_{t+1}, the improved estimate is given by
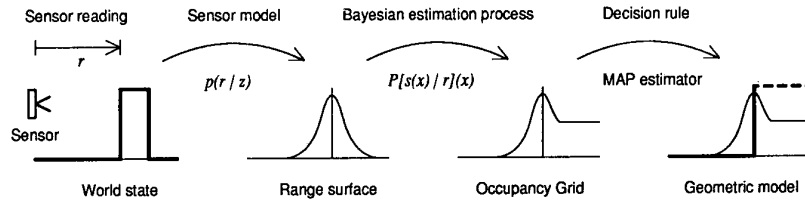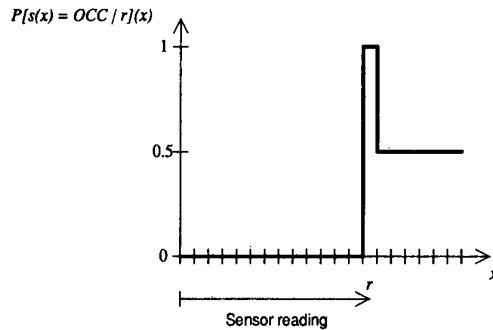
Figure 1. Estimating the occupancy grid from sensor data.

Figure 2. Occupancy probability profile for an ideal sensor.

$$P[s(C_i) = \text{OCC} \mid \{r\}_{t+1}] = \qquad (1)$$

$$\frac{p[r_{t+1} \mid s(C_i) = \text{OCC}] \, P[s(C_i) = \text{OCC} \mid \{r\}_t]}{\sum_{s(C_i)} p[r_{t+1} \mid s(C_i)] \, P[s(C_i) \mid \{r\}_t]}$$

In this recursive formulation, the previous estimate of the cell state, $P[s(C_i) = \text{OCC} \mid \{r\}_t]$, serves as the prior and is obtained directly from the occupancy grid. The new cell state estimate $P[s(C_i) = \text{OCC} \mid \{r\}_{t+1}]$ is subsequently stored again in the map. For the initial prior cell state probability estimates, we use maximum entropy priors.[6] Obtaining the $p[r \mid s(C_i)]$ distributions from the sensor model $p(r \mid z)$ is done using Kolmogoroff's theorem.[7] We can

derive closed form solutions of these equations for certain sensor models and compute numerical solutions in other cases.

To illustrate the approach, Figure 2 shows the occupancy profile derived for the case of a one-dimensional ideal range sensor, characterized by $p(r \mid z) = \delta(r-z)$. Given a range reading r, the corresponding cell has occupancy probability 1. The preceding cells are empty and have occupancy probability 0. The succeeding cells have not been observed and are therefore unknown, so the occupancy probability is 0.5.

A sequence of occupancy profiles obtained from a one-dimensional Gaussian range sensor appears in Figure 3. The sensor model

$$p(r \mid z) = \frac{1}{\sqrt{2\pi}\,\sigma} \exp\left(\frac{-(r-z)^2}{2\sigma^2} \right) \qquad (2)$$

is shown superimposed (dashed line). Several successive updates of the occupancy probabilities are plotted, with the sensor positioned at $x=0.0$ and $r=2.0$. The grid was initialized to $P[s(x) = \text{OCC}](x) = 0.5$. The sequence of occupancy profiles shows that the occupancy grid converges towards the behavior of the ideal sensor.

Finally, a two-dimensional occupancy grid generated from a single sonar range reading is shown in Figure 4. The sonar sensor is modeled with Gaussian uncertainty in both range and angle. The sensor probability density function is given by

$$p(r \mid z,\theta) = \qquad (3)$$

$$\frac{1}{2\pi\sigma_r\sigma_\theta} \exp\left[-\frac{1}{2} \left(\frac{(r-z)^2}{\sigma_r^2} + \frac{\theta^2}{\sigma_\theta^2} \right) \right]$$

The occupancy profile shown corresponds to a range measurement taken by a sonar sensor positioned at the upper left and pointing to the lower right. The horizontal surface corresponds to the unknown level.

Sensor integration. To increase the capabilities and performance of robotic systems in general requires a variety of sensing devices to support the various tasks to be performed. Since different sensor types have different operational characteristics and failure modes, they can in principle complement each other. This is particularly important for mobile robots, where multiple sensor systems can be used to generate improved world models and provide higher levels of safety and fault tolerance.

Within the occupancy grid framework, sensor integration can be performed using

a formula similar to Equation 1 to combine the estimates provided by different sensors.[1,2] For two sensors S_1 and S_2, this requires using the corresponding sensor models $p_1(r \mid z)$ and $p_2(r \mid z)$. As a result, the same occupancy grid can be updated by multiple sensors operating independently.

A different estimation problem occurs when separate occupancy grids are maintained for each sensor system and the integration of these sensor maps is performed at a later stage by composing the corresponding cell probabilities. This scenario requires the combination of probabilistic evidence from multiple sources, which can be addressed using an estimation method known as the independent opinion pool.[6] This method involves summing the evidence for each cell state and performing the appropriate normalization.

Incorporation of user-provided maps. Throughout this article we are mainly concerned with scenarios where the robot operates in unknown environments, so no prior maps can be used. As already mentioned, however, in some situations such knowledge is available and can be represented using geometric and symbolic models.[8]

The occupancy grid framework incorporates information from such high-level precompiled maps using the same methodology outlined in the previous sections. To provide a common representation, the geometric models are scan-converted into an occupancy grid, with occupied and empty areas assigned appropriate probabilities. These precompiled maps can subsequently be used as priors or can simply be treated as another source of information to be integrated with sensor-derived maps.

Decision making. For certain applications, it may be necessary to assign specific states to the cells of the occupancy grid. An optimal estimate of the state of a cell is given by the maximum a posteriori (MAP) decision rule[7]: a cell C is occupied if $P[s(C) = \text{OCC}] > P[s(C) = \text{EMP}]$, empty if $P[s(C) = \text{OCC}] < P[s(C) = \text{EMP}]$, and unknown if $P[s(C) = \text{OCC}] = P[s(C) = \text{EMP}]$.

We could use other decision criteria, such as minimum-cost estimates. Depending on the specific context, it may also be useful to define an unknown band, as opposed to a single thresholding value. However, many robotic tasks can be performed directly on the occupancy grid, precluding the need to make discrete choices concerning the state of individual
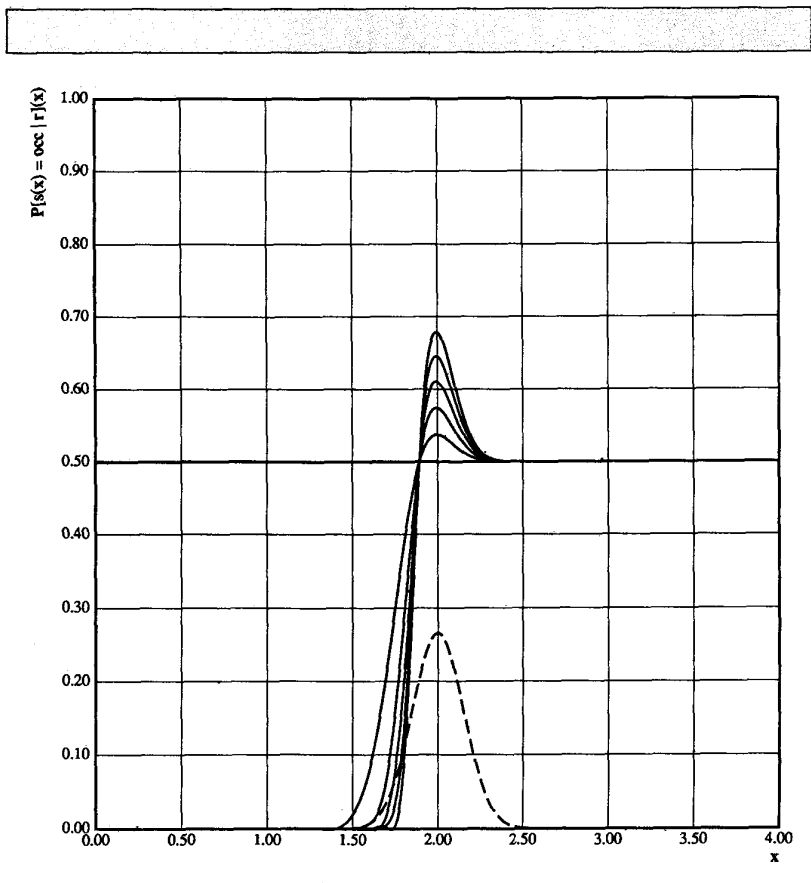

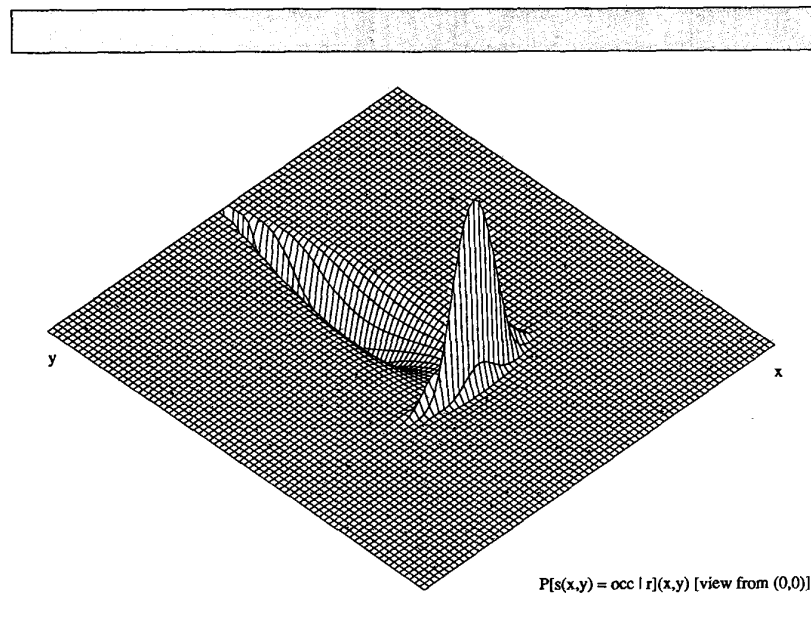**Figure 3. Occupancy probability profiles for a Gaussian sensor.**

$P[s(x,y) = \text{occ} \mid r](x,y)$ [view from (0,0)]

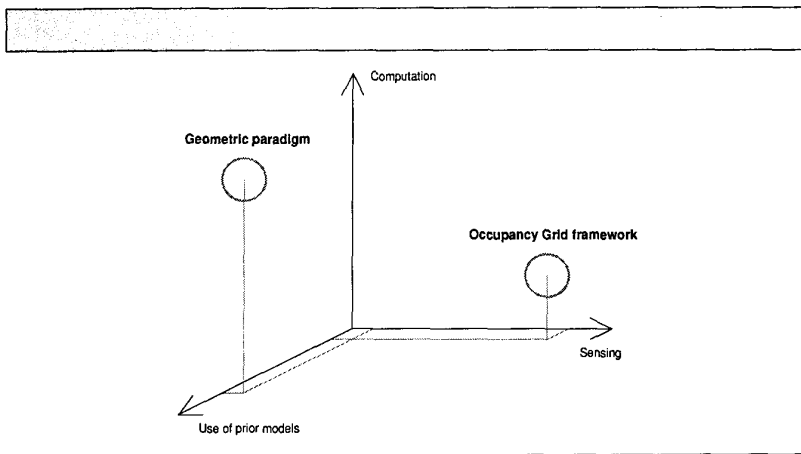**Figure 4. Occupancy grid for a two-dimensional Gaussian sensor.**

Figure 5. A comparison of emphases in the geometric paradigm versus the occupancy grid framework.
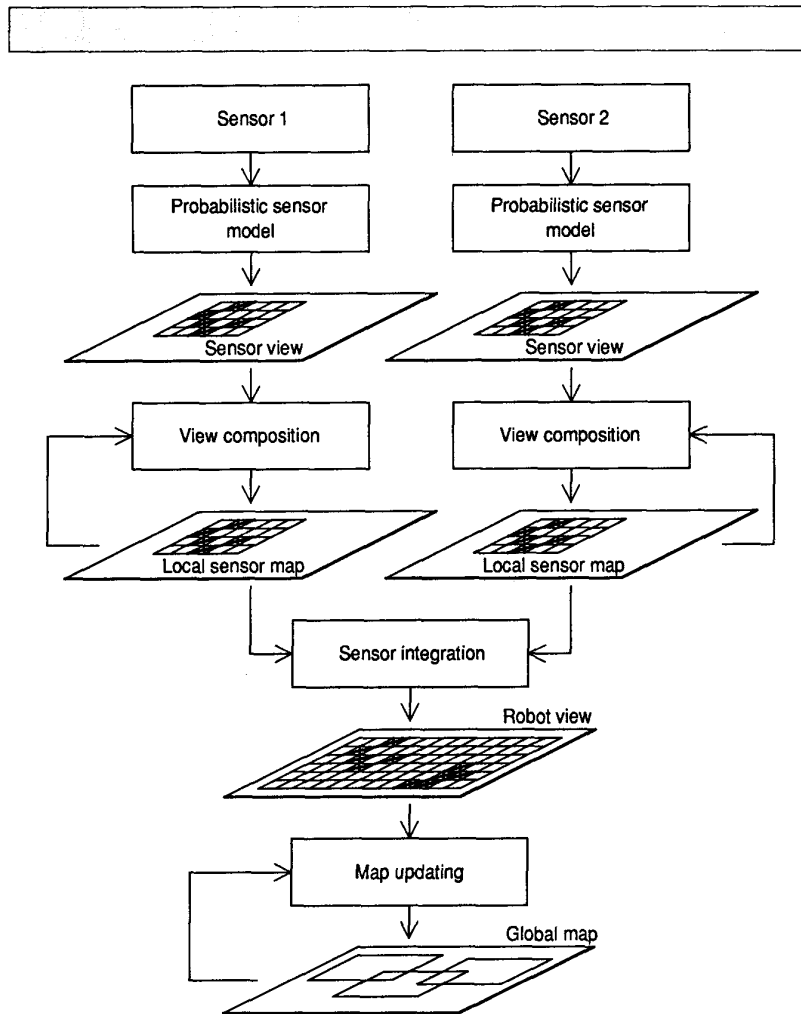

Figure 6. A framework for occupancy-grid-based robot mapping.

cells. In path planning, for example, we can compute the cost of a path in terms of a risk factor directly related to the corresponding cell probabilities.[1,3]

Characteristics of the occupancy grid approach. From the foregoing discussion, several aspects of the occupancy grid framework become evident. I have stressed the use of probabilistic sensor models to perform sensor interpretation and handle sensor uncertainty, and the use of probabilistic estimation procedures to update the occupancy grid. Consequently, no precompiled geometric models and no runtime segmentation decisions are necessary. Additionally, the use of a decision-theoretic framework makes possible statements about the optimality of the estimates.

Further note that the occupancy grid itself provides a stochastic spatial world model. The random field explicitly encodes both the spatial information and the associated uncertainty, and does not require discrete choices. It is possible to derive deterministic voxel models or higher-level geometric representations from the occupancy grid; however, the suitability of a representation is directly related to how well it describes its subject and how easily relevant information can be extracted from it. From this point of view, I argue that a number of robotic tasks can be efficiently addressed within the occupancy grid framework.[1]

This approach also has some specific implications. Due to the intrinsic limitations of sensor systems, spatial sensor interpretation is fundamentally an underconstrained problem. Within the occupancy grid framework, we achieve disambiguation of the sensor data and recovery of better world models primarily through strategies that emphasize additional sensing, rather than through the use of finer tuned heuristics or additional assumptions about the robot's environment. Instead of relying on a small set of observations to generate a world model, we compose information from multiple sensor readings taken from different viewpoints to estimate and improve the sensor-derived occupancy grids. This leads naturally to an emphasis on a high sensing-to-computation ratio and on the development of improved sensor models and active sensing strategies.

Figure 5 provides a contrast between some of the emphases in the occupancy grid approach and in the geometric paradigm, outlined earlier.

Using occupancy grids for mobile robot mapping

Reviewing some applications to the mobile robot domain will illustrate the occupancy grid framework. This section discusses the use of occupancy grids in sensor-based robot mapping. The next section provides an overview of their use in robot navigation.

One possible flow of processing for the use of occupancy grids in mobile robot mapping appears in Figure 6. The vehicle explores and maps its environment, acquiring information about the world. Data acquired from a single sensor reading is called a *sensor view*. Various sensor views taken from a single robot position can be composed into a *local sensor map*. Multiple sensor maps can be maintained separately for different sensor types, such as sonar or laser. To obtain an integrated description of the robot's surroundings, sensor fusion of the separate local sensor maps is performed to yield a *robot view*, which encapsulates the total sensor information recovered from a single sensing position. As the vehicle travels through its terrain of operation, robot views taken from multiple data-gathering locations are composed into a *global map* of the environment. This requires the registration of the robot views to a common frame of reference, an issue addressed in the next section.

For experimental validation, the framework outlined above was implemented and tested on several mobile robots in both indoor and outdoor scenarios. We will look at some results derived from experiments in sonar-based mapping and in sensor integration of sonar and single-scanline stereo.

Sonar-based mapping. Early work with sonar-based mapping[3,9] initially motivated the development of occupancy grids and led to the implementation of a mobile robot range-based mapping and navigation system called Dolphin. A variety of experiments were used to test this system.[1-3] For indoor runs, a mobile robot called Neptune was used (see Figure 7); outdoor runs were performed with a larger robot vehicle called the Terregator. More recently, a new version of the Dolphin system was installed on the Locomotion Emulator, a mobile platform designed for navigation in mining environments (see Figure 8).

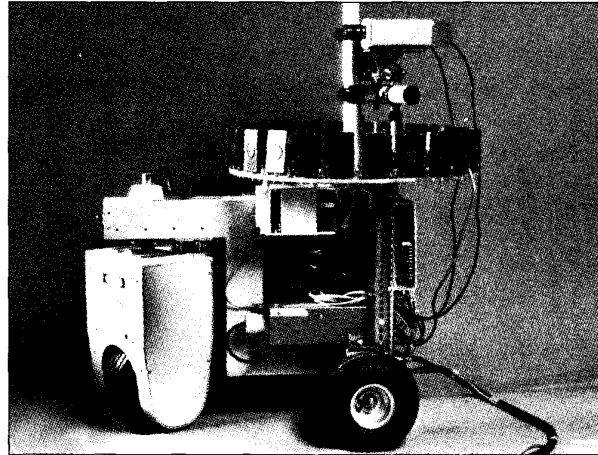Figure 9 displays a typical 2D sonar occupancy grid, while Figure 10 provides

Figure 7. The Neptune mobile robot, built by Gregg Podnar at the Carnegie Mellon University Mobile Robot Lab, shown with a circular sonar sensor array and a pair of stereo cameras. Vehicle locomotion and sensor interfaces are controlled by on-board processors, while the Dolphin mapping and navigation system runs on an off-board mainframe. This robot was used for indoor range mapping and sensor integration experiments.
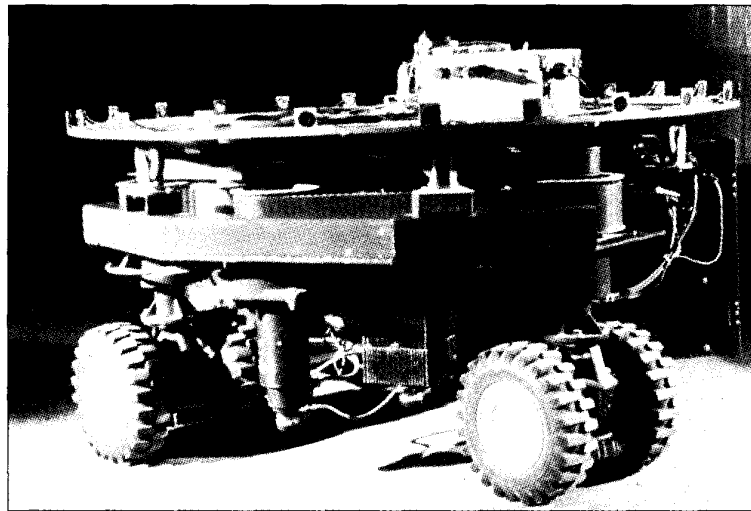

Figure 8. The Locomotion Emulator mobile robot, built at the CMU Field Robotics Center. Designed for navigation experiments in mining environments, this vehicle is capable of implementing several locomotion strategies. It is shown here with a sonar sensor array.

a 3D plot of the corresponding occupancy probabilities. Examples of other maps are given in Figure 11, which shows a sonar map obtained during navigation down a corridor, and Figure 12, which corresponds to a run in a wooded outdoor park.
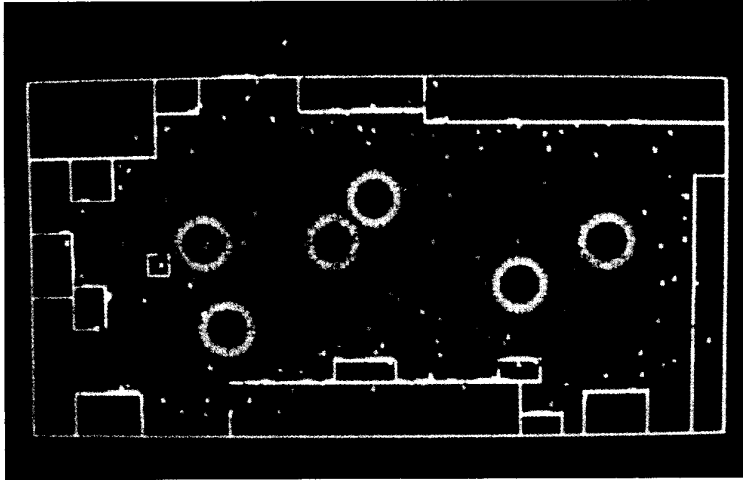
Figure 9. A two-dimensional sonar occupancy grid. Cells with high occupancy probability are represented in red, while cells with low occupancy probability are shown in blue. The robot positions from where scans were taken are shown by green circles, while the outline of the room and of major objects is given by white lines. This map shows the Mobile Robot Lab.
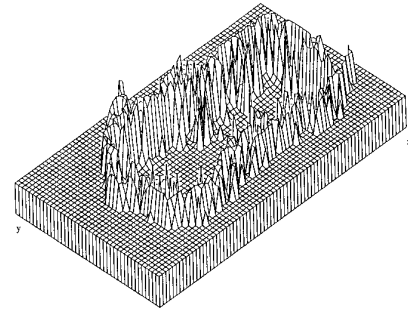
Figure 10. Occupancy grid probabilities for the sonar map. This 3D view shows the occupancy probabilities $P[s(Cx_i \, y_j) = \text{OCC}](x_i \, y_j)$ of the map in Figure 9.
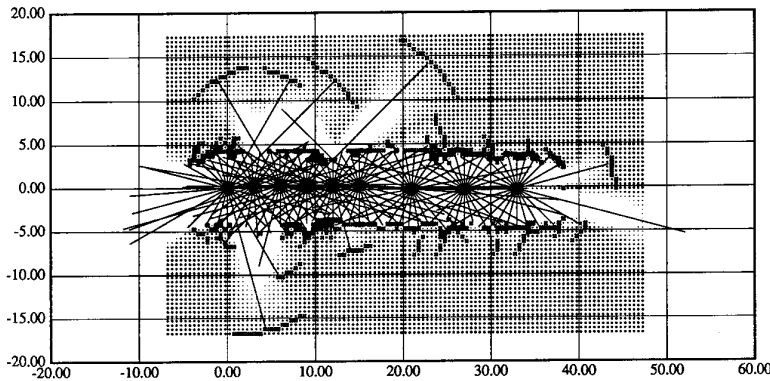
Figure 11. Sonar mapping and navigation along a corridor. Walls and open doors can be distinguished, and the resolution is sufficient to allow wall niches to be noticeable in the map. The range readings taken from each robot stop are drawn superimposed on the occupancy grid.

Sensor integration of sonar and scanline stereo. The occupancy grid framework provides a straightforward approach to sensor integration. Range measurements from each sensor are converted directly to the occupancy grid representation, where data taken from multiple views and from different sensors can be combined naturally. Sensors are treated modularly, and separate sensor maps can be maintained concomitantly with integrated maps, allowing independent or joint sensor operation. In collaboration with Larry Matthies, I have performed experiments in the fusion of data from two sensor systems: a sonar sensor array and a single-scanline stereo module that generates horizontal depth profiles.[4] For sensor integration runs, the Neptune mobile robot was configured with a sonar sensor ring and a pair of stereo cameras (see Figure 7). The independent opinion pool method, mentioned earlier, was used to combine the occupancy grids derived separately for the two sensor systems.

Figure 13 shows a typical set of maps. In general terms, we can see that the integrated maps take advantage of the complementarity of the sensors. The stereo system depends on matching high-contrast image
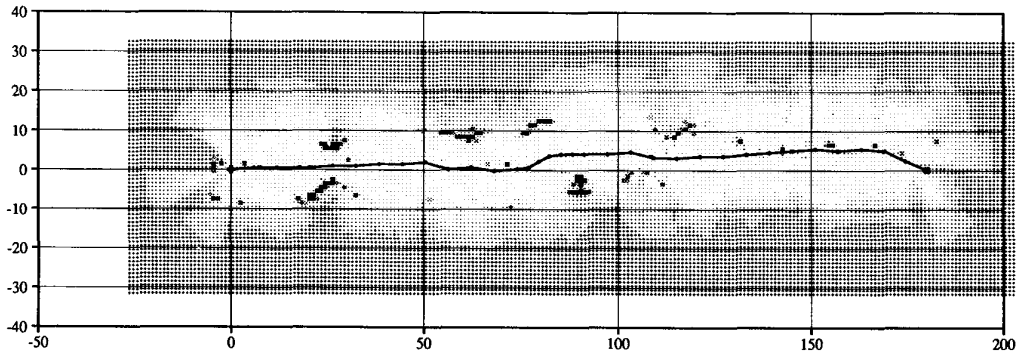
Figure 12. An outdoor run. This map shows a sonar-based outdoor run in a wooded park area. The obstacles encountered are trees.

features, so unmarked surfaces or low-contrast edges are not detected well. Stereo angular resolution is comparatively high, while the range uncertainty increases with distance. Sonar, on the other hand, detects surfaces well. But it has poor angular resolution due to the large beam width, while the range uncertainty itself is comparatively low. Some of these characteristics become noticeable in Figure 13, where sonar misses open paths due to its beam width, while stereo misses object edges due to low contrast against the background. A corrective behavior can be seen in the integrated map.

Using occupancy grids for mobile robot navigation

We now turn to some examples of the use of occupancy grids in mobile robot navigation. We briefly address issues in path planning, estimating and updating the robot position, and incorporating the positional uncertainty of the robot into the mapping process (as shown in Figure 14).

Path planning. In the Dolphin system, path planning and obstacle avoidance are performed using potential functions and an A* search algorithm. The latter operates directly on the occupancy grid, optimizing a path cost function that takes into account both the distance to the goal and the occupancy probabilities of the cells trav-

ersed.[1,3] Results of the operation of the path planner can be seen in Figures 11 and 12.

Handling robot position uncertainty. To allow the merging into a coherent model of the world of multiple views acquired by the robot from different sensing positions, we need accurate motion information to allow precise registration of the views for subsequent composition. For mobile robots that move around in unstructured environments, recovering precise position information poses major problems. Over longer distances, dead reckoning estimates are not sufficiently reliable. Consequently, motion-solving methods that use landmark tracking or map matching approaches are usually applied to reduce the registration imprecision due to motion. Additionally, the positional error is compounded over sequences of movements as the robot traverses the environment. This leads to the need for explicitly handling positional uncertainty and taking it into account when composing multiview sensor information.

To represent and update the robot position as the vehicle explores the terrain, we use the approximate transformation (AT) framework developed by Smith, Self, and Cheeseman.[10] A robot motion M, defined with respect to some coordinate frame, is represented as $\tilde{M} = <\hat{M}, \Sigma_M>$, where $\hat{M}$ is the estimated (nominal) position and Σ_M is the associated covariance matrix that captures the positional uncertainty. The parameters of the robot motion are determined from dead reckoning and inertial navigation estimates, which can be composed
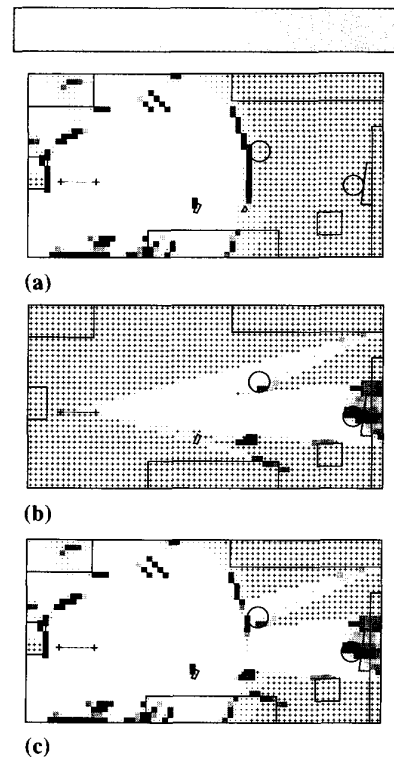

(a)

(b)

(c)

Figure 13. Sensor integration of sonar and scanline stereo. Occupancy grids generated separately for sonar (a) and scanline stereo (b), as well as the integrated map (c) obtained through sensor fusion, are shown. Occupied regions are marked by shaded squares, empty areas by dots fading to white space, and unknown spaces by + marks.
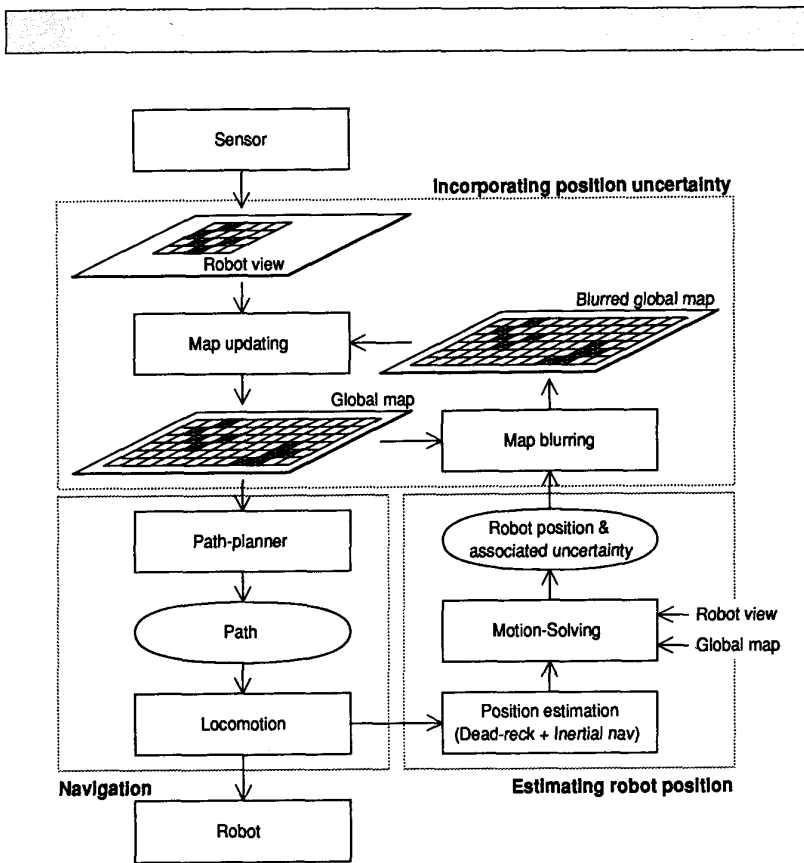
Figure 14. A framework for occupancy-grid-based robot navigation. New robot views are used to update the global map, which in turn is used by the path planner. After locomotion, the new robot position estimate is refined using a motion-solving procedure that finds an optimal registration between the robot view and the current global map. Finally, the remaining positional uncertainty is incorporated into the map updating process as a blurring operation.

using the AT merging operation, while the updating of the robot position uncertainty over several moves is done using the AT composition operation.[10]

Motion solving. For more precise position estimation, we employ a multiresolution correlation-based motion-solving procedure.[9] Increasingly lower resolution versions of the occupancy grids are generated, and the search for an optimal registration between the current robot view and the global map is done first at a low level of resolution. The result is subsequently propagated up to guide the search process at the next higher level of resolution.

Incorporating positional uncertainty into the mapping process. After estimating the registration between the new robot view and the current global map, we can incorporate the associated uncertainty into the map updating process as a blurring or convolution operation performed on the occupancy grid. We distinguish between world-based mapping and robot-based mapping.[1,2] In *world-based mapping*, the motion of the robot is related to an absolute or world coordinate frame, and the current robot view is blurred by the robot's global positional uncertainty prior to composition with the global map. If we represent the blurring operation by the symbol $\otimes$ and

the composition of maps by the symbol $\tilde{\oplus}$, the world-based mapping procedure can be expressed as

global map $\leftarrow$ global map $\tilde{\oplus}$ (robot view $\otimes$ global position uncertainty)

Since the global robot position uncertainty increases with every move, this updating procedure has the effect that the new views become progressively more blurred, adding less and less useful information to the global map. Observations seen at the beginning of the exploration are "sharp," while recent observations are "fuzzy." From the point of view of the inertial observer, the robot eventually "dissolves" in a cloud of probabilistic smoke.

For *robot-based mapping*, we estimate the registration uncertainty of the global map due to the recent movement of the robot, and the global map is blurred by this uncertainty prior to composition with the current robot view. This mapping procedure can be expressed as

global map $\leftarrow$ (global map $\tilde{\otimes}$ local position uncertainty) $\tilde{\oplus}$ robot view

A consequence of this method is that observations performed in the remote past become increasingly uncertain, while recent observations have suffered little blurring. From the point of view of the robot, the immediate surroundings (which are of direct relevance to its current navigational tasks) are "sharp." The robot is leaving, so to speak, an expanding probabilistic trail of weakening observations (see Figure 15).

Note, however, that the local spatial relationships observed within a robot view still hold. To avoid losing this information, we use a two-level spatial representation, incorporating occupancy grids and approximate transformations. On one level, the individual views are stored attached to the nodes of an AT graph (a *stochastic map*[10]) that describes the movements of the robot. On the second level, a global map is maintained that represents the robot's current overall knowledge of the world (see Figure 16). This two-level structure provides an adequate and efficient representation for various navigation tasks.

Operations on occupancy grids

We have looked at the application of the occupancy grid framework to the mobile

Robot view 0

Global map 0

Robot view 1

Global map 1

Robot view 2
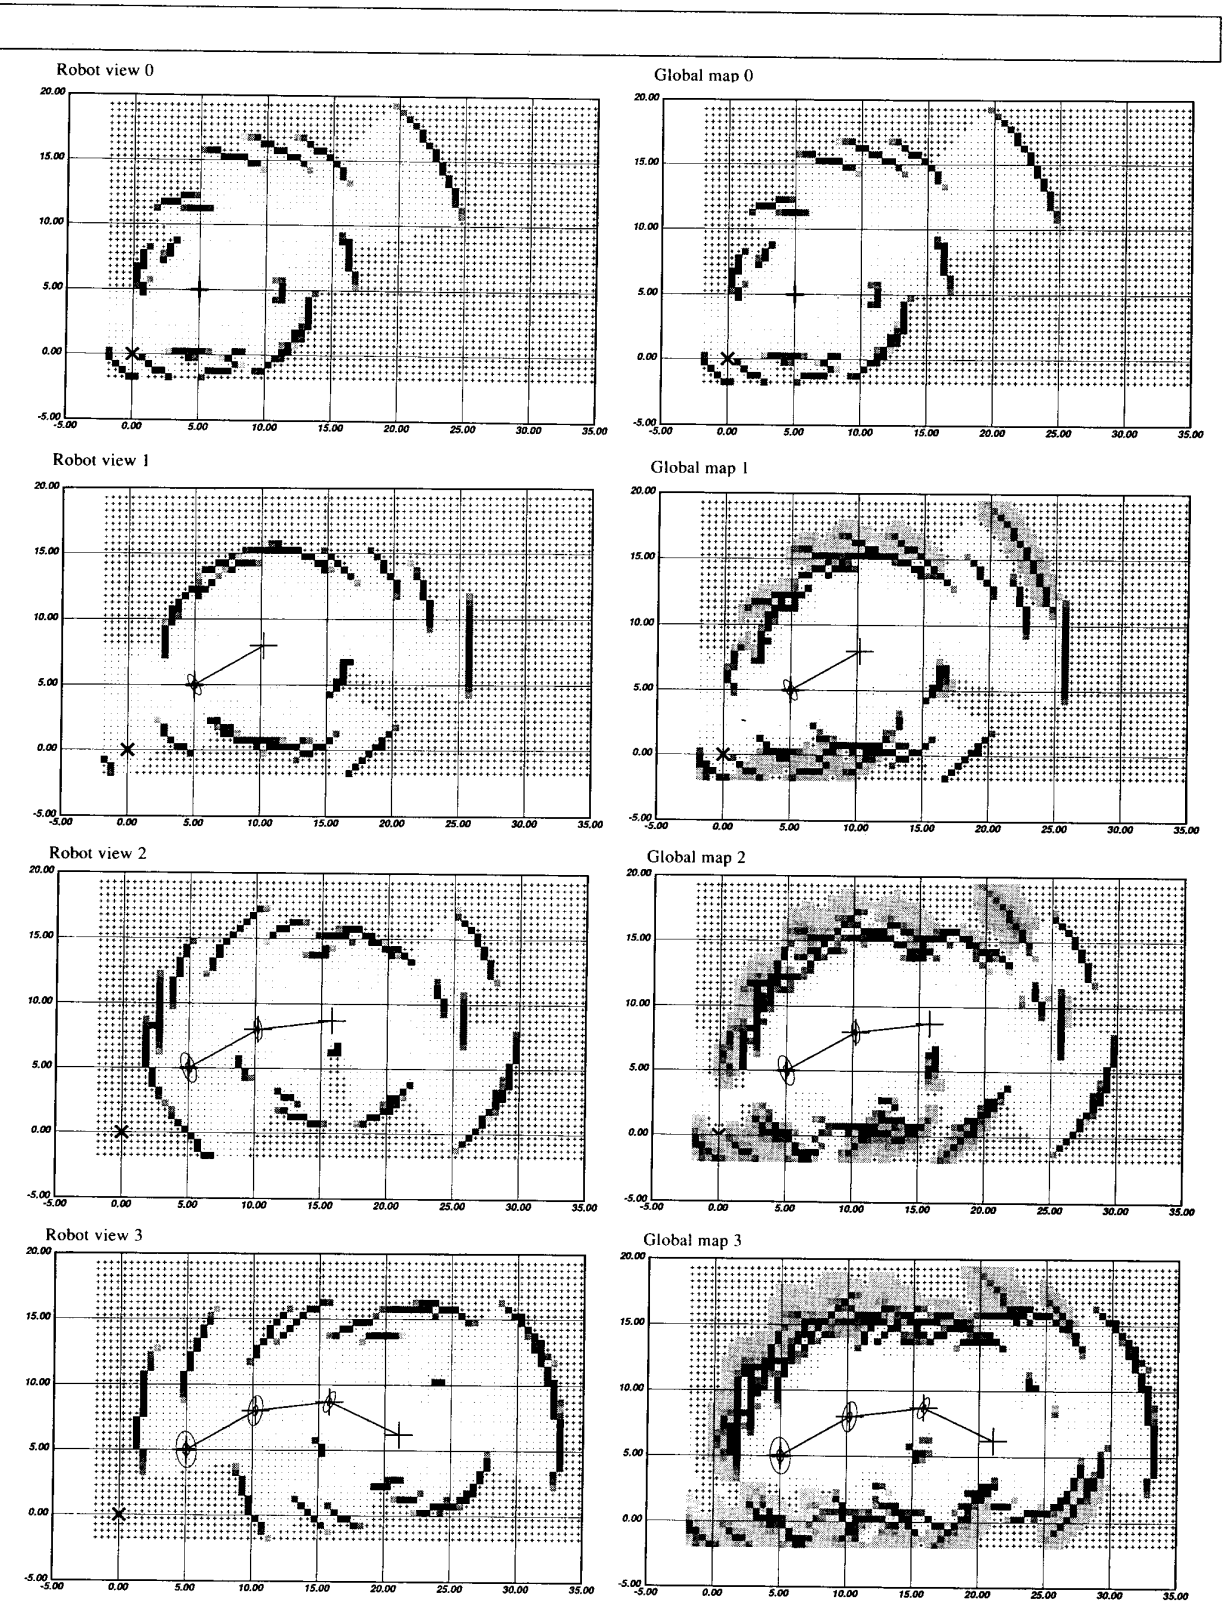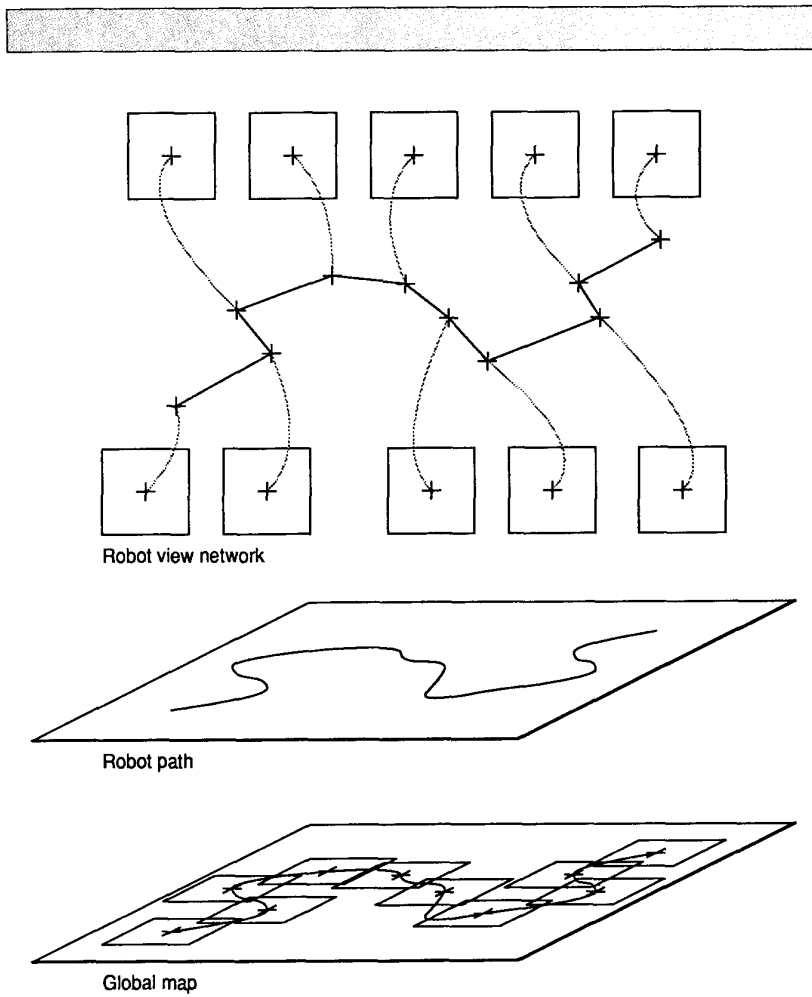
Global map 2

Robot view 3

Global map 3

Figure 15. Incorporating motion uncertainty into the mapping process. For robot-centered mapping, the global map is blurred by the back-propagated robot position uncertainty (shown using the corresponding covariance ellipses) prior to composition with the robot view.

Robot view network

Robot path

Global map

Figure 16. Maintaining a two-level spatial representation. The individual robot views are stored attached to the nodes of an AT graph describing the robot motion and are maintained in conjunction with the global map .

Table 1. Operations on occupancy grids for various robotic tasks and similar operations performed in the image processing domain.

Occupancy Grids	Images
Labeling cells as occupied, empty, or unknown	Thresholding
Handling position uncertainty	Blurring/convolution
Removing spurious spatial readings	Low-pass filtering
Map matching/motion solving	Multiresolution correlation
Obstacle growing for path planning	Region growing
Path planning	Edge tracking
Extracting occupied, empty, and unknown areas	Segmentation/region coloring/labeling
Determining object boundaries	Edge detection
Incorporating precompiled maps	Scan conversion
Prediction of sensor observations from maps	Correlation
Object motion detection over map sequences	Space-time filtering

robot mapping and navigation domain. This framework also allows us to address a number of other robot perception and spatial reasoning problems in a unified way. It is important to observe that many operations performed on occupancy grids for various robotic tasks are similar to computations performed in the image processing domain. This is a useful insight, since it allows us to take advantage of results from this context. Table 1 provides a qualitative overview and comparison of some of these operations.

Extending the occupancy grid framework

Additional issues explored within the occupancy grid framework include the recovery of geometric descriptions from occupancy grids,[3] the incorporation of precompiled maps,[1] and the use of logarithmic maps where the resolution drops with the distance to the robot.[1] Other possible applications include the prediction of sensor readings from occupancy grids and the detection of moving objects over sequences of maps. Current work is investigating other domains, such as the use of occupancy grids for laser scanner mapping, precise positioning, and navigation in mining applications using the Locomotion Emulator; the development of mapping and planning strategies that take advantage of high-level precompiled maps when available; the exploration of strategies for landmark recognition and tracking; and the recovery of 3D occupancy grids from laser rangefinders or stereo depth profiles.

We have reviewed the occupancy grid framework and looked at results from its application to mobile robot mapping and navigation in unknown and unstructured environments. The occupancy grid framework represents a fundamental departure from traditional approaches to robot perception and spatial reasoning. It supports agile and robust sensor interpretation methods, incremental discovery procedures, composition of information from multiple sensors and over multiple positions of the robot, and explicit handling of uncertainty. Furthermore, the occupancy grid representation can be used directly in various robotic planning and problem-solving activities, thereby precluding the

need for the recovery of deterministic geometric models. The results suggest that the occupancy grid framework provides an approach to robot perception and spatial reasoning that has the characteristics of robustness and generality necessary for real-world robotic applications. □

Acknowledgments

References

1. A. Elfes, *Occupancy Grids: A Probabilistic Framework for Mobile Robot Perception and Navigation*, PhD thesis, Electrical and Computer Engineering Dept./Robotics Inst., Carnegie Mellon Univ., 1989.

2. A. Elfes, "A Tesselated Probabilistic Representation for Spatial Robot Perception," *Proc. 1989 NASA Conf. on Space Telerobotics*, NASA/Jet Propulsion Laboratory, California Inst. of Technology, Pasadena, Calif., Jan. 31-Feb. 2, 1989.

3. A. Elfes, "Sonar-Based Real-World Mapping and Navigation," *IEEE J. Robotics and Automation*, Vol. RA-3, No. 3, June 1987.

4. A. Elfes and L.H. Matthies, "Sensor Integration for Robot Navigation: Combining Sonar and Stereo Range Data in a Grid-Based Representation," *Proc. 26th IEEE Conf. on Decision and Control*, Dec. 1987. Also in *Proc. 1988 IEEE Int'l Conf. on Robotics and Automation*, CS Press, Los Alamitos, Calif.

5. E. Vanmarcke, *Random Fields: Analysis*

and Synthesis, MIT Press, Cambridge, Mass., 1983.

6. J.O. Berger, *Statistical Decision Theory and Bayesian Analysis*. 2nd ed., Springer-Verlag, Berlin, 1985.

7. A.E. Bryson and Y.C. Ho, *Applied Optimal Control*, Blaisdell Publishing, Waltham, Mass., 1969.

8. D.J. Kriegman, E. Triendl, and T.O. Binford, "A Mobile Robot: Sensing, Planning, and Locomotion," *Proc. 1987 IEEE Int'l Conf. on Robotics and Automation*, CS Press, Los Alamitos, Calif., April 1987.

9. H.P. Moravec and A. Elfes, "High-Resolution Maps from Wide-Angle Sonar," *Proc. IEEE Int'l Conf. on Robotics and Automation*, CS Press, Los Alamitos, Calif., March 1985.

10. R.C. Smith, M. Self, and P. Cheeseman, "A Stochastic Map for Uncertain Spatial Relationships," *Proc. 1987 Int'l Symp. on Robotics Research*, MIT Press, Cambridge, Mass., 1987.

Alberto Elfes is a researcher with the Engineering Design Research Center and the Robotics Institute at Carnegie Mellon University. His research interests include robotics, computer vision, mobile robots, and design automation systems. His current work addresses the development of probabilistic and estimation theory approaches to robot perception and spatial modeling, and the use of spatial reasoning techniques in design automation. He has published more than 30 articles in these areas.

Elfes received the EE degree in electronics engineering in 1975 and the MS degree in computer science in 1980 from the Instituto Tecnológico de Aeronáutica, Brazil. He was granted the PhD degree in 1989 by the Electrical and Computer Engineering Department at Carnegie Mellon University. He is a member of the IEEE Computer Society, the IEEE, and the ACM.

Readers may contact Elfes at the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213.