International Workshop on Information Processing in Mobile Robots, Technische Universität München, Germany, March 6 - 8, 1991.

DYNAMIC CONTROL OF ROBOT PERCEPTION USING STOCHASTIC SPATIAL MODELS

Alberto Elfes

Department of Computer Sciences
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598
E-Mail (Internet): ELFES@IBM.COM

Abstract

Robot perception has traditionally been addressed as a passive and incidental activity, rather than an active and task-directed activity. Consequently, although sensor systems are essential to provide the information required by the decision-making and actuation components of a robot system, no explicit planning and control of the sensory activities of the robot is performed. This has lead to the development of sensor modules that are either excessively specialized, or inefficient and unfocused in their informational output. In this paper, we develop strategies for the dynamic control of robot perception, using stochastic sensor and spatial models to explicitly plan and control the sensory activities of an autonomous mobile robot, and to dynamically servo the robot and its sensors to acquire the information necessary for successful execution of robot tasks. We discuss the explicit characterization of robot task-specific information requirements, the use of information-theoretic measures to model the extent, accuracy and complexity of the robot's world model, and the representation of inferences about the robot's environment using the Inference Grid, a multi-property tesselated random field model. We describe the use of stochastic sensor models to determine the utility of sensory actions, and to compute the loci of observation of relevant information. These models allow the development of various perception control strategies, including attention control and focussing, perceptual responsiveness to varying spatial complexity, and control of multi-goal perceptual activities. We illustrate these methodologies using an autonomous multi-sensor mobile robot, and show the application of dynamic perception strategies to active exploration and multi-objective motion planning.

1 Introduction

Traditionally, most approaches to robot perception have cast the sensing and perceptual activities of a robot system in what is fundamentally a passive rôle: although the robot planning and action stages depend fundamentally on sensor-derived information, no explicit planning and control of the perceptual activities themselves is performed. Rather, the robot's sensing subsystems acquire data that incidentally happens to be available at the moment of sensory observation during the perception/planning/action cycle. No explicit connection is made between the sensor data acquired and the information required by the robot for the successful execution of a given task (Fig. 1). Robot systems that embody this *passive perception* approach tend to fall into two categories: those where the sensory subsystems are highly specialized and "hardwired" to the feedback control loop that handles the execution of a specific class of robot tasks; and those that embody

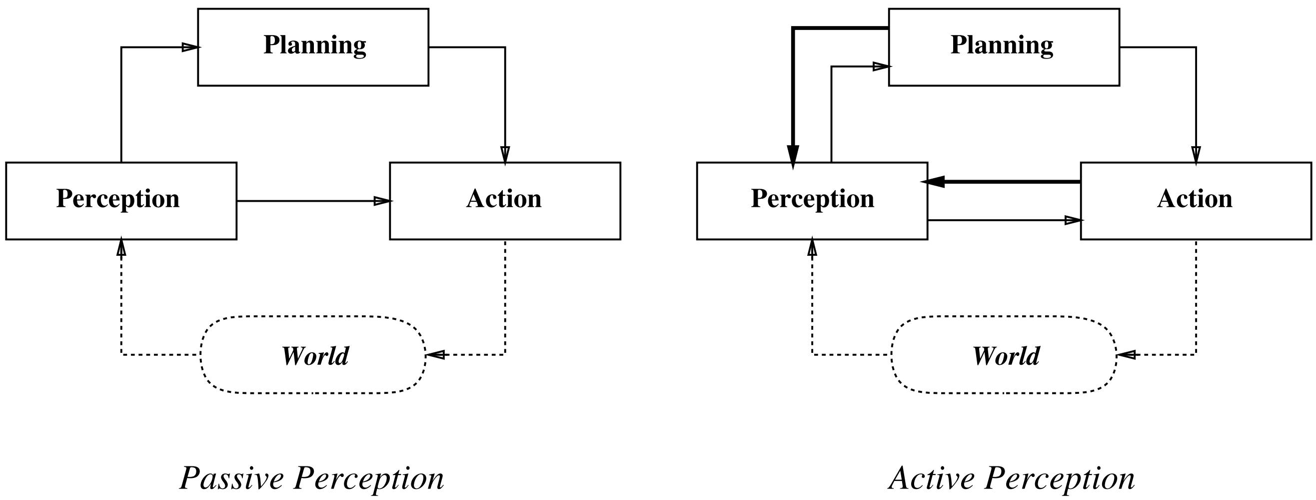

Figure 1: Passive and Active Control of Robot Perception.

"general-purpose" sensor understanding and world-modelling subsystems. Typically, the former category cannot reconfigure its sensory subsystems to adapt to different classes of tasks or to accomodate sensor degradation; while the latter category is populated by notoriously large and ineffective systems that are unfocused in their informational output, have slow reaction times, and are inefficient in the use of their sensory and computational resources. Because there is no explicit characterization of the information really needed, over- or undersampling of the data may occur, and data irrelevant to the task at hand may be acquired.

More recently, researchers have identified the limitations intrinsic to the formulation of robot perception as a passive activity, and have argued for viewing it as an active process (Fig. 1). Bajcsy [2] discusses the need for *active sensing*, defined as the application of modelling and control strategies to the various layers of sensory processing required by a robot system. Elfes has developed an approach for *active mapping* using the Occupancy Grid framework [12, 14, 10]; this framework uses estimation-theoretic and Markov Random Field models to compose information from multiple sensors and multiple points of view, thereby addressing in a robust way the underconstrainedness of the sensor data. Aloimonos suggests that the composition of information from multiple views can be used to handle several ill-posed problems in Computer Vision [1]. Related research includes the development of recursive estimation procedures for robot perception [20, 15, 18]; methods for *active sensor control*, where specific parameters of the sensor system (such as camera aperture, focal distance, or sensor placement) can be changed under computer control [19, 25]; development of theoretical foundations for coordination, integration and control of sensor systems [8, 17, 16]; and generation of optimal and adaptive sensing strategies [7, 6, 25].

In this paper, we discuss *Dynamic Perception*, a framework for dynamic and adaptive planning and control of robot perception in response to the information needs of an autonomous robot system as it executes a given mission. The Dynamic Perception framework uses stochastic and information-theoretic sensor and world models to identify the information needs of the robot, plan the acquisition of task-specific data, dynamically servo the robot and its sensors to acquire the information needed for successful execution of the task, update the information acquisition goals as the robot progresses through sequences of tasks, and integrate the sensory tasks with the mission-specific tasks to be performed by the robot. We describe the use of

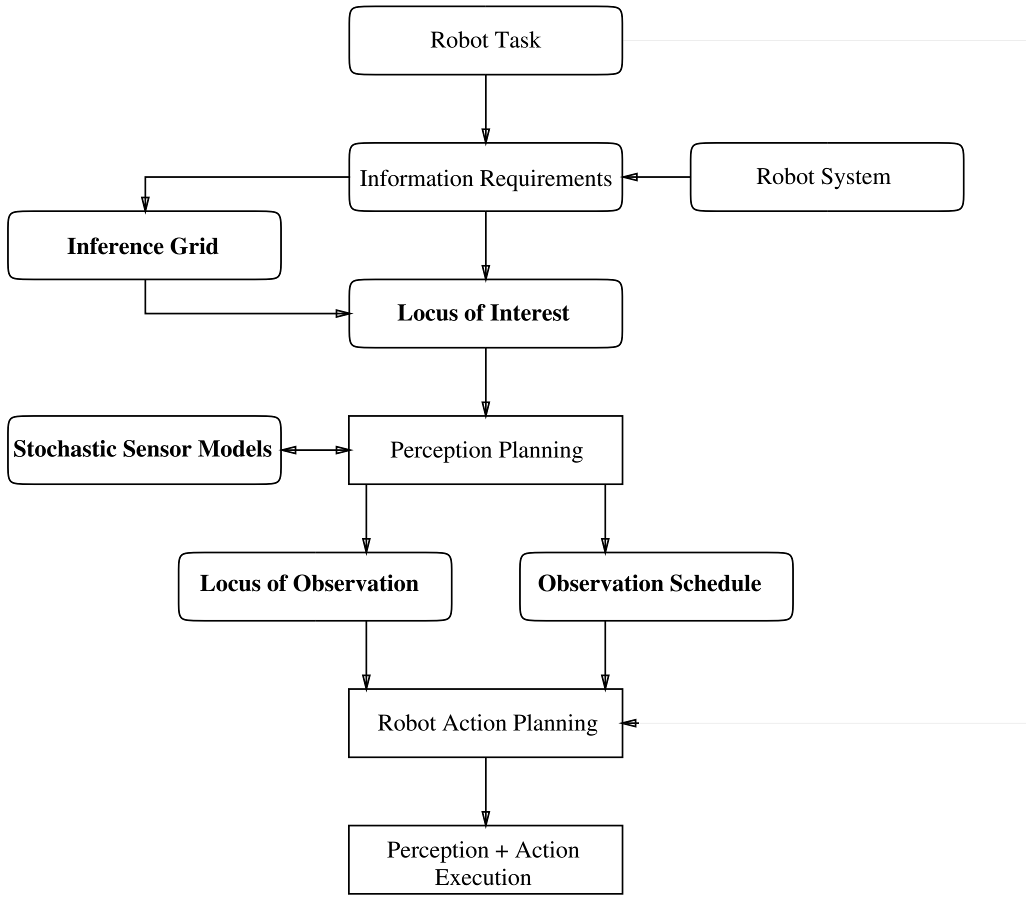

Figure 2: Components of the Dynamic Robot Perception Framework.

a multi-property tesselated random field model, the Inference Grid, to encode inferences about the robot's environment. We illustrate the explicit characterization of robot task-specific information requirements, and the use of information-theoretic measures to model the extent, accuracy and complexity of the robot's world model. We also discuss the use of stochastic sensor models to evaluate the utility of sensory actions, and to compute the loci of observation of relevant information. We illustrate our approach using an autonomous multi-sensor mobile robot, and show the application of dynamic perception strategies to *active exploration* and *integrated motion planning*, combining perception and navigation goals. A more extensive discussion, with further experimental results, can be found in [13].

2 Dynamic Robot Perception

The work presented here is part of a long-term research effort that addresses the development of *agile* and *robust* autonomous robot systems, able to execute complex, multi-phase missions in unknown and unstructured real-world environments, while displaying real-time performance [10, 13]. For planning, scheduling, execution and monitoring purposes, a specific robot *mission* is separated into *phases*, which are in turn decomposed into sets of *tasks*. These tasks can be scheduled in parallel and/or sequentially, depending on the nature of the component activities. As tasks are scheduled for detailed planning and execution by the robot, the information requirements posed by these tasks are used to dynamically update the perceptual goals of the robot. By maintaining explicit models of the task goals and the perceptual goals, the Dynamic Perception framework allows a closer integration between task planning and perception planning, and between task execution and perception execution. Using descriptions of the information required for successful execution of a specific task, appropriate sensors are selected, sensing strategies are formulated, and the robot's locomotion and sensing activities are planned so as to maximize the recovery of relevant information and accomplish the robot's task. Overall, the system's behaviour can be described as *servoing on required information*, as well as on the robotic task itself. We note that in control-theoretic terms this can be phrased as a problem in *dual control* [4, 13].

Some of the issues that have to be addressed in the context of the Dynamic Perception framework include *Sensor Modelling*, *Information Modelling*, *Perception Planning*, and the *Integration of Perception and Action*. This will allow us to enable an autonomous mobile robot with a number of important capabilities, including *attention control*, or the ability to efficiently acquire and process relevant data; *attention focussing* from larger sensing areas to smaller regions of specific interest; and development of *optimal information acquisition* and *exploration strategies*.

3 Components of the Dynamic Perception Framework

We have chosen to address the concerns outlined above through the development of estimation- and information-theoretic models. Contrary to the *ad hoc* AI-based methods still used in much of Computer Vision work, estimation-theoretic models have a long history of success and have been widely applied in signal processing and control tasks [4]. In this section, we describe some of the components of the Dynamic Perception framework. In particular, we discuss the *stochastic sensor models* developed to describe the robot sensors; the use of a stochastic multi-property tesselated spatial representation, the *Inference Grid*, to encode inferences about the robot's world; the development of various information metrics to measure the robot's knowledge of the environment, and of complexity measures to determine the spatial variability of the environment; the use of *mutual information* to measure the utility of sensory actions; and the computation of the *Locus of Interest* and the *Locus of Observation* of relevant information [13]. The flow of computation between these components is illustrated in Fig. 2.

3.1 Stochastic Sensor Models

To enable us to reason about the perceptual capabilities of an autonomous robot system, mathematical models are needed to describe the behaviour of the various sensors of the robot. In previous work, we have discussed sensor models that are stochastic in nature and have shown their use in the recursive estimation of a tesselated spatial random field model, the Occupancy Grid [14, 10]. While the specific sensor models developed will depend on the properties being measured, the physical characteristics of the sensor systems and the environment of operation of the robot, we have found that the class of models briefly described below can be tailored to a number of different range sensor systems.

The stochastic range sensor models used for the experimental component of our research take into account the target detection probability, the sensor noise that corrupts the range measurements, and the variation of range uncertainty with distance. The uncertainty in the measurement $\mathbf{r}$ of a detected target T positioned at $\mathbf{z}$ from the sensor is modelled by the pdf $p(\mathbf{r} \mid det(T, \mathbf{z}))$ expressed as:

$$p(\mathbf{r} \mid det(T, \mathbf{z})) = \widetilde{G}(\mathbf{r}, \mathbf{z}, \Sigma(\mathbf{z})) \tag{1}$$

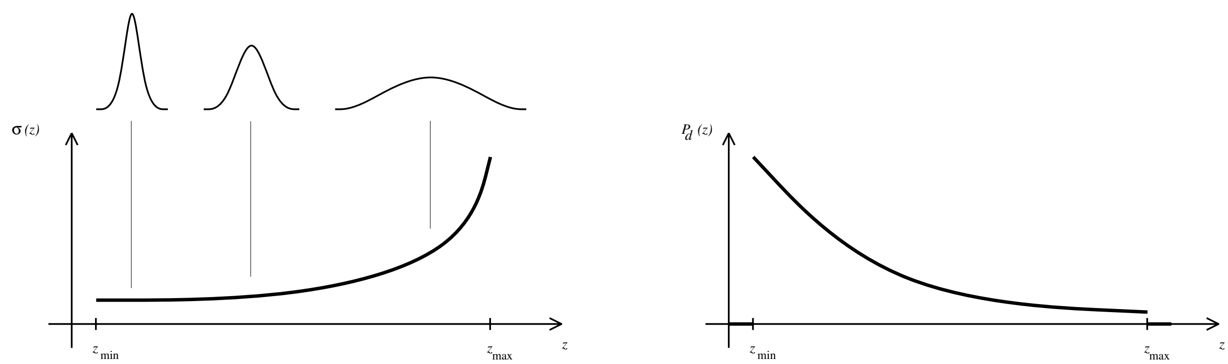

Figure 3: Sensor Range Variance and Detection Probability as a Function of Distance to the Target. These are typical curves for range estimation only.

where $det(T, \mathbf{z})$ is the detection event of T at $\mathbf{z}$, $\Sigma(\mathbf{z})$ describes the variation of range noise with distance, and $\widetilde{G}()$ is a corrupted Gaussian distribution, given by

$$\widetilde{G}() = (1 - \epsilon)\; G_1() + \epsilon\; G_2() \tag{2}$$

The distributions $G_1()$ and $G_2()$ are Gaussian, with $G_1()$ modelling the normal behaviour of the sensor, and $G_2()$ occasional gross errors. The parameter ϵ weighs the relative contributions of both terms.

To model the sensor detection behaviour, we use the target detection probability, $P[det(T, \mathbf{z}) \mid \exists\, T \text{ at } \mathbf{z}](\mathbf{z}) = P_d(\mathbf{z})$, and the false alarm probability, $P[det(T, \mathbf{z}) \mid \nexists\, T \text{ at } \mathbf{z}](\mathbf{z}) = P_f(\mathbf{z})$ (known as a Type I error [9]). The missed detection probability (Type II error) is of course given by $P[\neg det(T, \mathbf{z}) \mid \exists\, T \text{ at } \mathbf{z}](\mathbf{z}) = 1 - P_d(\mathbf{z})$.

The target localization probability, $p(\exists\, T \text{ at } \mathbf{z} \mid \mathbf{r})$, can be computed using Bayes estimation as:

$$p(\exists\, T \text{ at } \mathbf{z} \mid \mathbf{r}) = \frac{p(\mathbf{r} \mid \exists\, T \text{ at } \mathbf{z})\; P[\exists\, T \text{ at } \mathbf{z}]}{p(\mathbf{r} \mid \exists\, T \text{ at } \mathbf{z})\; P[\exists\, T \text{ at } \mathbf{z}] + p(\mathbf{r} \mid \nexists\, T \text{ at } \mathbf{z})\; (1 - P[\exists\, T \text{ at } \mathbf{z}])} \tag{3}$$

where

$$p(\mathbf{r} \mid \exists\, T \text{ at } \mathbf{z}) = p(\mathbf{r} \mid det(T, \mathbf{z}))\; P_d(\mathbf{z}) + p(\mathbf{r} \mid \neg det(T, \mathbf{z}))\; (1 - P_d(\mathbf{z})) \tag{4}$$

and

$$p(\mathbf{r} \mid \nexists\, T \text{ at } \mathbf{z}) = p(\mathbf{r} \mid det(T, \mathbf{z}))\; P_f(\mathbf{z}) + p(\mathbf{r} \mid \neg det(T, \mathbf{z}))\; (1 - P_f(\mathbf{z})) \tag{5}$$

This class of stochastic sensor models can be applied to a large variety of range sensors, including infrared and laser scanners, sonar sensors, and stereo systems. The required parameters can be obtained through analysis of the physical characteristics of the sensor and through calibration (see, for example, [10, 20, 19, 21]). A typical set of curves showing the dependency of the detection probability and of the range variance on the distance to the target being imaged is given in Fig. 3, while Fig. 4 shows a plot of $p(\mathbf{r} \mid \exists\, T \text{ at } \mathbf{z})$. These stochastic sensor models are used in perception planning to determine the *Locus of Observation*, from which the robot can acquire spatial information of relevance to specific perceptual goals.

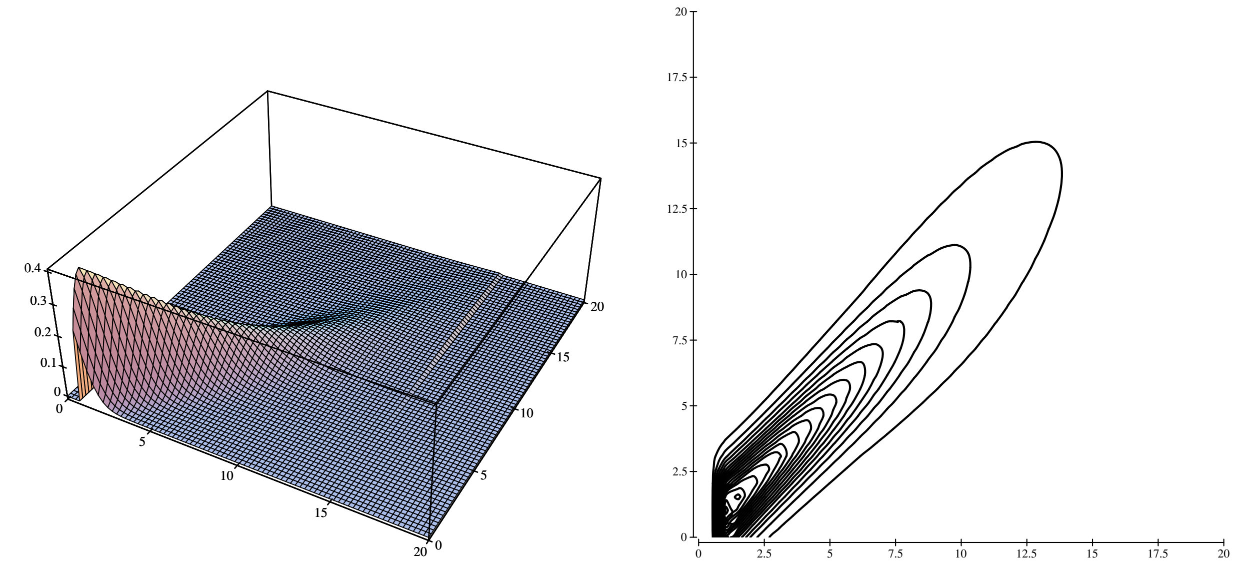

Perspective Plot **Contour Plot**

Figure 4: Stochastic Range Sensor Model. The function $p(\mathbf{r} \mid \exists\, T \text{ at } \mathbf{z})$ for range estimation only is shown, in both a perspective plot and a contour plot.

3.2 Inference Grids

In previous work, we have developed an approach to spatial robot perception and navigation called the *Occupancy Grid* framework [12, 14, 10]. The Occupancy Grid is a multi-dimensional discrete random field model that maintains probabilistic estimates of the occupancy state of each cell in a spatial lattice. Recursive Bayesian estimation mechanisms employing stochastic sensor models allow incremental updating of the Occupancy Grid using multi-view/multi-sensor data, as well as composition of multiple maps, decision-making, and incorporation of robot and sensor position uncertainty. The Occupancy Grid framework provides a unified approach to a variety of problems in the mobile robot domain, including autonomous mapping and navigation, sensor integration, path planning under uncertainty, motion estimation, handling of robot position uncertainty, and multi-level map generation. It has been successfully tested on several mobile robots, operating in real-time in real-world indoor and outdoor environments [10].

For the Dynamic Perception work we have generalized the Occupancy Grid representation, and have developed a spatial model called the *Inference Grid*. The Inference Grid is a multi-property Markov Random Field defined over a discrete spatial lattice. By associating a random vector with each lattice cell, the Inference Grid allows the representation and estimation of multiple spatially distributed properties. For robot perception and spatial reasoning purposes, typical properties of interest may include the *occupancy* state of a lattice cell, as well as its *observability* and *reachability* by the robot (Fig. 5). For visual perception, properties such as surface *color* or *reflectance* may be estimated, while for navigation purposes properties such as terrain *traversability* or region *connectedness* may be of relevance. Properties may be independent, or derivable from other properties.

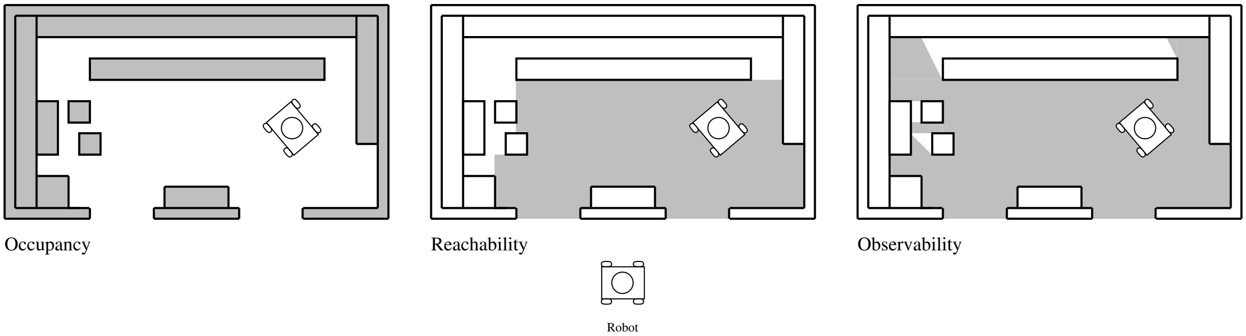

Figure 5: Some Properties of the Inference Grid: Occupancy, Reachability, Observability. The shaded areas of each map indicate regions that are occupied, reachable by the robot, and observable by the robot.

The occupancy state, $s(C)$, of a cell C of the Occupancy Grid (which is now a component of the Inference Grid) is modelled as a discrete random variable with two states, *occupied* and *empty*. Given a sensor range reading r, and a stochastic sensor model $p(r \mid z)$, the recursive estimation of the Occupancy Grid is done using Bayes' theorem to obtain the cell state probabilities $P[s(C) = \textsc{occ} \mid r]$ (see [14, 10, 11]). An example of the Occupancy Grid is shown in Fig. 6.

The reachability and observability properties can also be treated as binary stochastic variables, and are estimated using the cell occupancy probabilities, $P[s(C) = \textsc{occ} \mid M]$, stored in the Occupancy Grid M. Cell *reachability*, $\Xi(C)$, is computed using $P[s(C) = \textsc{occ} \mid M]$ as:

$$P[\Xi(C) \mid M] = \prod_{\forall\, Z \,\in\, \eta} (1 - P[s(Z) = \textsc{occ} \mid M]) \;\text{ s.t. }\; \eta = \arg \min_{\forall \xi(M)} \Gamma(\xi) \tag{6}$$

where $\xi(M)$ is a trajectory defined over the Occupancy Grid, $\Gamma(\xi)$ is the cost of the trajectory, and η is the minimum-cost robot trajectory, computed using a dynamic programming formulation [10, 4]. Cell *observability*, $\Psi(C)$, is estimated using $P[\Xi(C) \mid M]$ and the sensor model of Eq. 3 as:

$$P[\Psi(C) \mid M] = \max_{\forall\, Z \,\in\, M} P[det(C) \mid s(C) = \textsc{occ} \wedge \pi(R) = Z]\; P[\Xi(Z) \mid M] \tag{7}$$

where $\pi(R)$ denotes the position and orientation of the robot.

3.3 Information Measures

To guide the perceptual activities of a robot, we need metrics to evaluate the robot's world knowledge. The specific metric used depends on the robot task and the particular kind of information required for successful execution of the task. For tasks that involve spatial reasoning and navigation, we require measures of the extent and accuracy of the robot's sensory maps, while for target localization or shape recovery, precise surface position information is needed.

Map Uncertainty

For the Inference Grid model discussed above, an intuitive way of expressing the uncertainty in the spatial information encoded in cell occupancy estimates obtained from the sensor data is given by the cell uncertainty function [10]:

$$U(C) = 1 - 4\left(P[s(C) = \textsc{occ} \mid M] - \frac{1}{2}\right)^2 \tag{8}$$

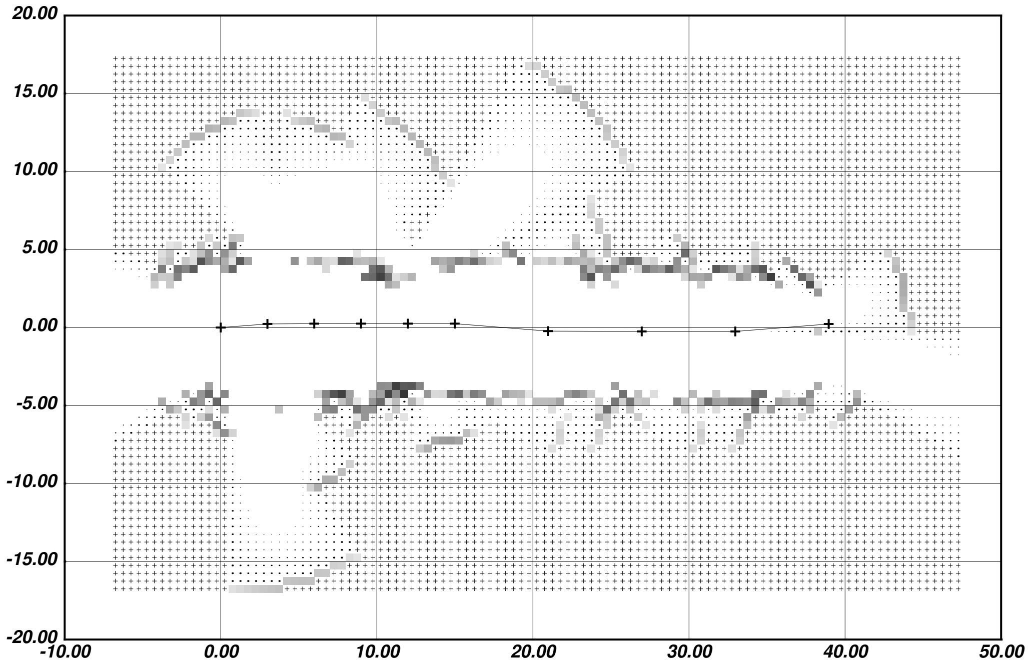

Figure 6: An Example of the Occupancy Grid. The map shows an Occupancy Grid built by a mobile robot using a sonar range sensor array. The robot is moving along a corridor, from left to the right. The gray cells correspond to high occupancy probability areas, while the areas marked with "·" correspond to cells with high emptyness probability. Areas not observed by the robot are identified using "+".

A more generally used metric of uncertainty is the entropy of a random variable [5, 24]. The Inference Grid cell uncertainty can be measured using the entropy of the cell occupancy state estimate as (Figs. 7 and 8):

$$E(C) = -\sum_{s_i} P[s_i(C) \mid M] \, \log P[s_i(C) \mid M] \tag{9}$$

Using the cell entropy, the uncertainty over a region W of the Inference Grid can be computed as:

$$E(W) = \sum_{\forall\, C_i \in W} E(C_i) \tag{10}$$

This definition allows us to determine upper and lower bounds on the uncertainty of the region W:

$$0 \leq E(W) \leq \#(W) \tag{11}$$

where $\#(W)$ is the cardinality of the region W. To obtain an entropy measure that is independent of the region size, the average entropy $\overline{E}(W)$ is defined as:

$$\overline{E}(W) = \frac{E(W)}{E_0(W)} \tag{12}$$

where $E_0(W) = \#(W)$ is the maximum entropy of W.

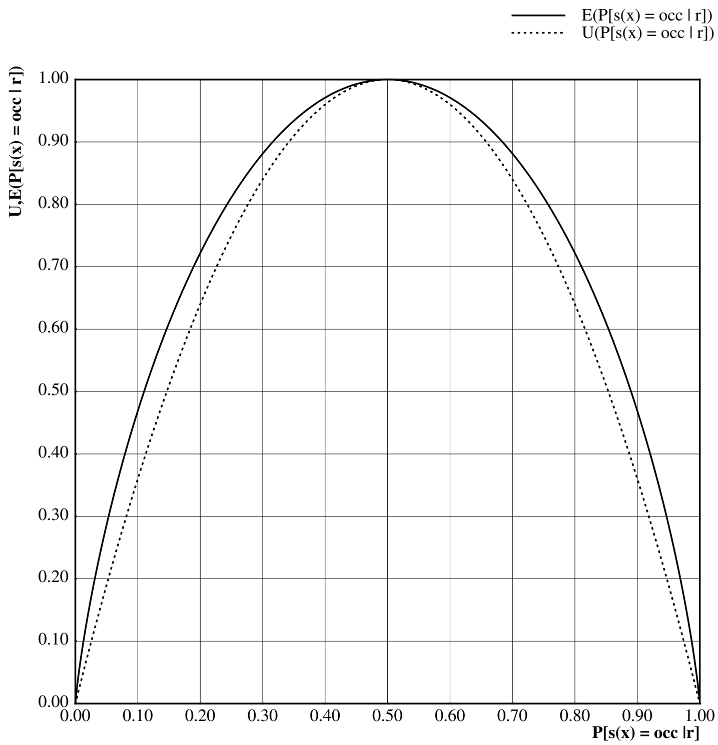

Figure 7: Occupancy Grid Cell Uncertainty Measures. The cell uncertainty function $U(C)$ and the cell entropy function $E(C)$ are shown.

Target Localization Uncertainty

For precise localization or shape recovery tasks, the error probabilities for target detection and the variance in the position estimate of detected features, such as surfaces or vertices, can be directly used as quantitative uncertainty measures. Consider a range sensor whose measurements are corrupted by Gaussian noise of zero mean and variance σ^2, modelled by the pdf:

$$p(r \mid z) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(\frac{-(r-z)^2}{2\sigma^2}\right) \tag{13}$$

The uncertainty in surface localization is given directly by σ, and the Type II error probability $1 - P_d(\mathbf{z})$ provides a direct measure of detection uncertainty.

3.4 Information Provided by a Sensing Action

The information $\Delta I(\alpha_k)$ added to the Inference Grid by a sensing action α_k can be defined directly in terms of the decrease in uncertainty caused by that sensing action:

$$\Delta I(\alpha_k) = -(E_k(W) - E_{k-1}(W)) \tag{14}$$

where $E_{k-1}(W)$ is the entropy of a region W of the Occupancy Grid before the sensing action α_k, and E_k is the entropy of W after the sensing action α_k. This measure is known as *Mutual Information* [24].

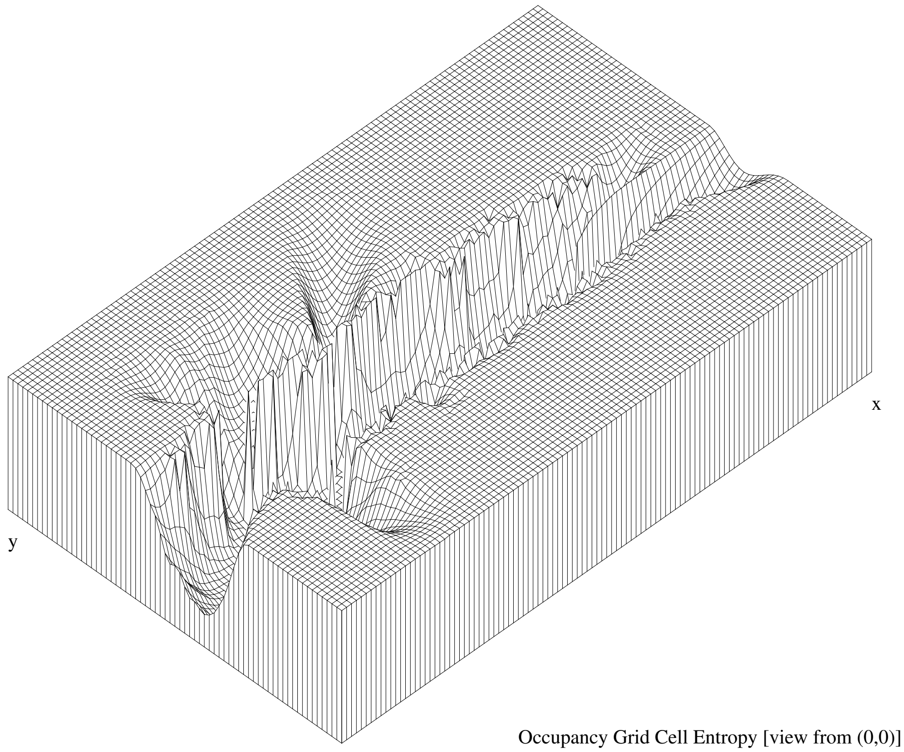

Figure 8: Entropy of the Occupancy Grid. The cell entropies for the Occupancy Grid of Fig. 6 are shown.

Eq. 14 serves as an indicator of the efficiency of a specific sensor or sensing action, and can be used to implement "stopping rules" such as those used in statistical analysis [3]. Fig. 9(a) shows how the average entropy $\overline{E}(M)$ of an OG M decreases for the case of a static robot, where multiple scans are performed from a fixed location. In contrast, Fig. 9(b) shows how the average entropy decreases for a mobile robot that is taking sensor scans from multiple locations, as it moves around in the environment. The graphs make explicit and express in quantitative terms what would be intuitively expected, namely, that it is more useful in terms of world knowledge acquisition to integrate data obtained from multiple sensing locations than from a single view. Note that, for non-biased robot sensors, we have $\lim_{k\to\infty} I(\alpha_k) = 0$. In the case of a static robot, the average entropy $\overline{E}_k(W)$ will tend towards $\lim_{k\to\infty} \overline{E}_k(W) = 1 - \#(\omega)/\#(W)$, where $\omega \subset W$ is the region observable by the sensors from the robot's location, and W is the total extent of the region of interest. For the case of a moving robot, the average entropy $\overline{E}_k(W)$ will tend towards $\lim_{k\to\infty} \overline{E}_k(W) = 1 - \#(\Omega)/\#(W)$, where $\Omega \subset W$ is the region observable by the robot as its explores its environment, and W is again the total extent of the region of interest. This behaviour can be observed in Figs. 9(a) and (b).

As we already mentioned, in localization problems the quality of an estimate is usually measured by its variance. If n measurements are taken using a sensor described by the model of Eq. 13, the sample variance of the estimated parameter $\hat{r}$ will be $\sigma_r^2 = \sigma^2/n$, and the information added by sensing action α_k is given

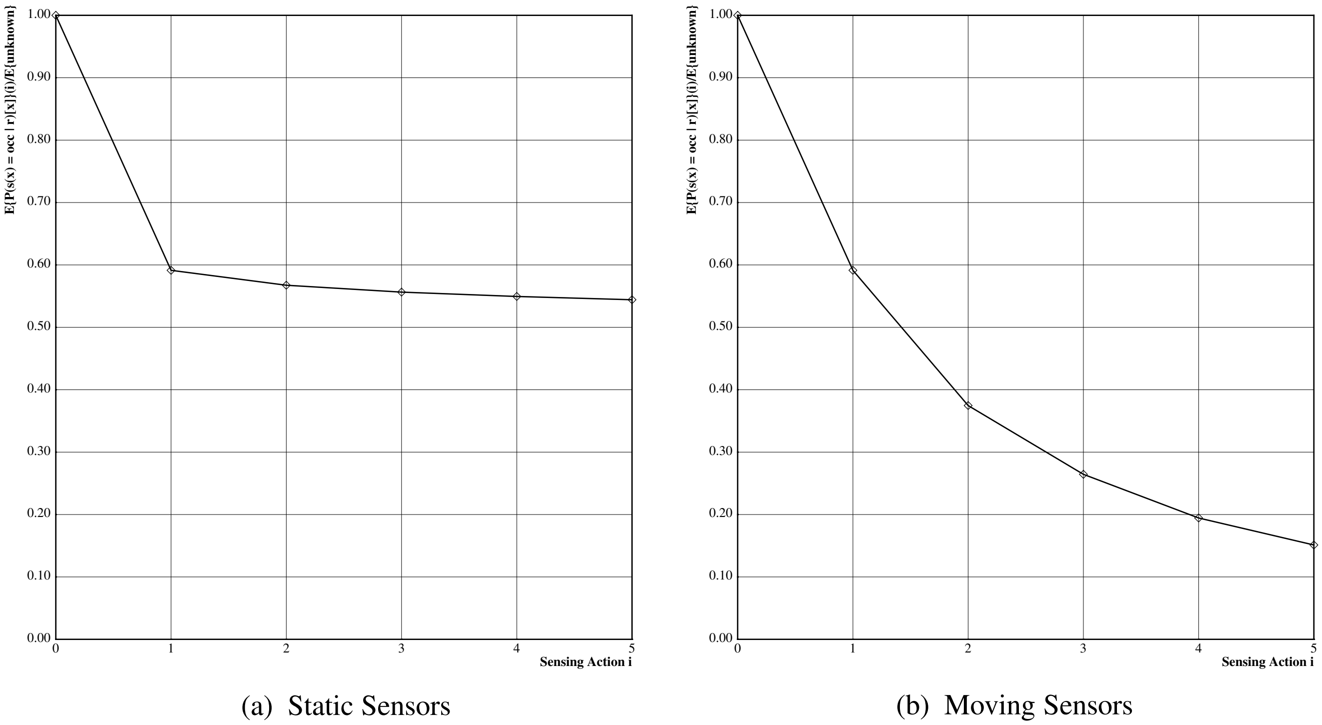

Figure 9: Occupancy Grid Entropy Change With Sequential Sensing Actions. Graphs (a) and (b) both show the change in OG uncertainty, measured as the average entropy $\overline{E}(M)$, as range sensor data is acquired sequentially and used to update the Occupancy Grid M. For graph (a), the robot and its sensors are static, and data is being acquired from a single view. In graph (b), the robot and its sensors are moving, and data is being acquired from multiple views. For this case, it can be seen that the uncertainty decay rate is faster, and that the asymptotic uncertainty limit is lower.

by:

$$-\Delta\sigma_k^2 = \frac{\sigma^2}{k\,(k-1)} \tag{15}$$

It is straightforward to associate utility functions to the information measures mentioned above (Eqs. 14 and 15). Similarly, cost functions can be associated with the effort and risk involved in performing a sensing action [4]. For planning and decision-making purposes, expected values of ΔI_k and $-\Delta\sigma_k^2$, as well as of the sensing action costs, can be used to select optimal single-stage sensory actions, compute limited lookahead multi-stage sensing strategies, and determine stopping or termination conditions for sensor data acquisition [13]. An example of a single-stage optimal sensory action choice for attention control is shown in Fig. 13.

3.5 Spatial Complexity

An additional metric of importance for dynamic control of robot perception is a quantitative measure of the complexity of the robot's world. For exploration and navigation purposes, for example, it is intuitive that the robot should investigate more carefully and do more frequent measurements of areas that have high spatial complexity, while spending less sensory and computational effort in “uninteresting” regions. A straightforward measure of the spatial complexity of a region W of an OG is given by the mean zero-crossing rate (for a given threshold) computed over W. Another measure is provided by interpreting the sensing activity of the robot as a spatial sampling process performed over the robot's environment. Given the tesselated nature of the Inference Grid model and its representational similarity to images [10, 11], a position-dependent Fast Fourier Transform (FFT) $\mathcal{F}_W[M(x, y)]$ can be computed over a finite-size window W of the Occupancy Grid M to obtain the spatial frequency spectrum of a specific region. Appropriate window functions include the rectangular window and the Hamming window [22]. The spatial frequency spectrum gives us a metric of the spatial complexity of the regions being explored by the robot, and allows us to derive an optimal sensing strategy, by performing the sensing actions at spatial intervals determined using the Nyquist (or optimal sampling) rate [22]:

$$\Delta x = \frac{\pi}{\omega_x} \qquad \text{and} \qquad \Delta y = \frac{\pi}{\omega_y} \tag{16}$$

where ω_x and ω_y are the band limits of the spatial frequency spectrum. An example of the use of spatial complexity measures to plan the locomotion of a mobile robot is shown in Fig. 6. Using spatial variability estimates of its immediate surroundings as constraints on the distances between data acquisition stops, the robot is able to respond to the complexity of its environment: it stops more frequently in regions of high spatial variability, such as the two open doors on either side of the left portion of the corridor, and speeds up when the corridor becomes “dull”.

4 Strategies for Dynamic Control of Robot Perception

We now turn to the application of the stochastic and information-theoretic models, discussed in the previous sections, in the development of strategies for dynamic control of robot perception. We will discuss the *Locus of Interest* and the *Locus of Observation*, outline methods for attention control and attention focussing, and illustrate the application of these strategies in autonomous robot exploration and in the integrated planning of robot navigation and robot perception.

4.1 Task-Directed Perception

Application scenarios for autonomous mobile robots require the execution of a variety of tasks related to spatial perception, reasoning and navigation, such as motion planning, detection and inspection of spatially distributed features, object recognition and pose determination, grasp planning, etc. In the work discussed here, the connection between robot task and robot perception is done by explicitly mapping the information needs of the task on the Inference Grid. Typical perceptual tasks may include observing a spatial feature with some minimum detection probability, localizing a spatial feature with some bound on the positional uncertainty, or selecting specific regions of interest, so that non-pertinent sensor data can be ignored.

Perception Constraints and the Locus of Interest

We define the *Locus of Interest* as a region specified on the Inference Grid that is fundamentally relevant for a specific robot task. It is determined by having the task define a utility function $R(W)$ over a region W of the Inference Grid, which measures the relevance of knowledge about W for successful task execution. Consequently, the Locus of Interest $L\,I$ defined by task τ_i is computed as the region of the Inference Grid that exceeds a utility threshold u_t:

$$L\,I(M, \tau_i) = \{\forall C \in M \text{ s. t. } R(C) \geq u_t\} \tag{17}$$

Examples of the Locus of Interest for some specific robot tasks are shown in Fig. 10. In addition to region

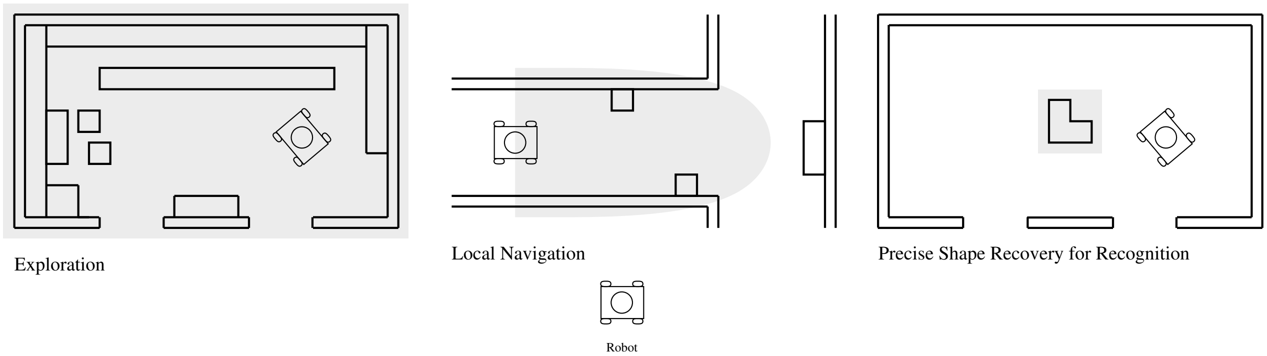

Figure 10: Locus of Interest for Several Robot Tasks. The regions of interest to the robot for *exploration*, *local navigation* and *precise shape recovery* tasks are shown as shaded areas.

selection, tasks can impose additional constraints on the information to be acquired. For example, *Detection Constraints* and *Localization Constraints* of the form $det(T, \mathbf{z}) \geq D_t$ and $\sigma(z) \leq \sigma_t$ can be defined, where D_t and σ_t are detection and range uncertainty thresholds (see example in Fig. 12). Note that if multiple goals are being pursued by the robot, the corresponding task-defined Loci of Interest and perceptual constraints can be merged and simultaneously represented on the same Inference Grid. It should also be mentioned that the robot system itself may impose LIs derived from robot-specific operational tasks.

The Locus of Observation

After determining the Locus of Interest, we can compute the *Locus of Observation*. The Locus of Observation is the configuration space region where the robot has to position a selected sensor or set of sensors to acquire the information needed by the robot's tasks and specified in the LI. The LO is computed as:

$$LO(M, \Theta[L\,I(M)]) = \{\forall C \in M \text{ s. t. } \pi(R) \in C \wedge \Theta[L\,I(M)] \geq \lambda\} \tag{18}$$

where $\Theta[L\,I(M)]$ is a predicate or utility function defined over the perceptual constraints, and which has to be above a threshold λ to be satisfactory. Fig. 11 shows the limits imposed on the region of operation of the stochastic sensor model of Eq. 4 by detection and range uncertainty constraints. An illustration of the Locus of Observation is given in Fig. 12. Industrially used mobile platforms, such as AGVs or sentry robots, frequently rely on the detection of specific landmarks, such as active beacons radiating on a known signature frequency, to determine the location of the vehicle or select the next segment of the path to be traversed. The shaded areas in Fig. 12 correspond to the Loci of Observation, and indicate the regions from which the beacons can be observed with $det(T, \mathbf{z}) \geq D_t$.

4.2 Attention Control and Attention Focussing

In the Dynamic Perception framework, *Attention Control* is performed through the selection of regions in the Inference Grid that have both *high current relevance* as measured by the Locus of Interest and are *observable* as measured in the Inference Grid. An example of Attention Control is given in Fig. 13. The task is exploration, and the robot has already partially mapped the region of interest. After a perception planning cycle during which the observability of the cells of the Inference Grid was estimated, four regions were selected for further exploration. These regions correspond to windows that have both high average Occupancy Grid entropy and high average probability of being observable.

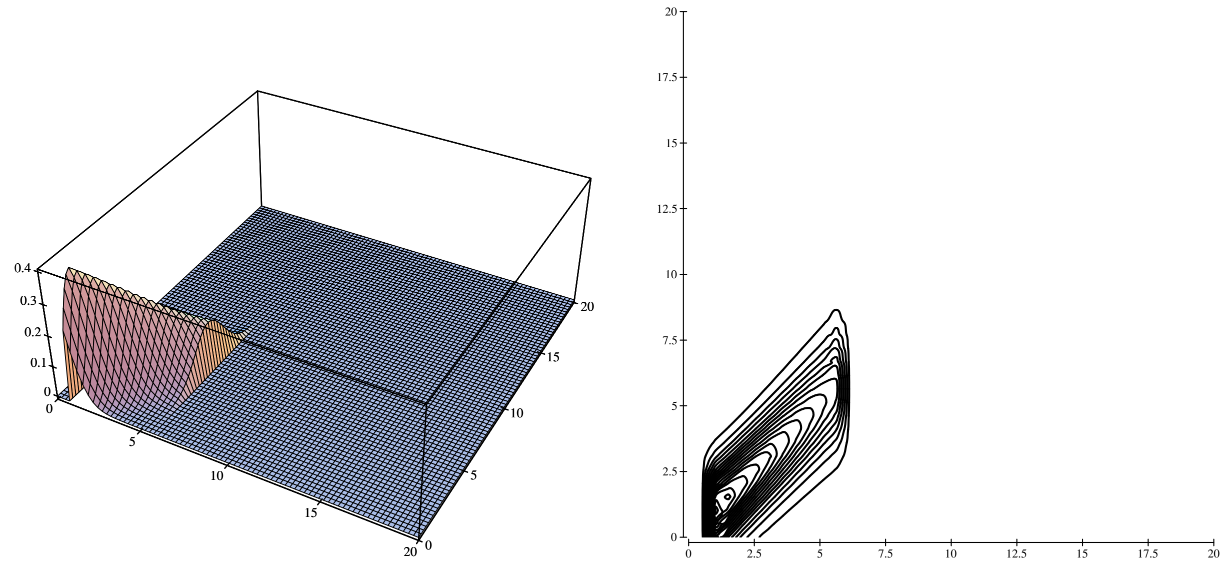

Perspective Plot **Contour Plot**

Figure 11: Imposing Perceptual Constraints on the Stochastic Sensor Model. The stochastic sensor model of Eq. 4 is shown with the operational limits imposed by the perceptual constraints $det(T, \mathbf{z}) \geq D_t$ and $\sigma(z) \leq \sigma_t$.

Given the uncertainty intrinsic to sensory observations and execution of robot actions, it is generally not possible to compute exact *a priori* measures of observability, sensor information, etc., and many of the measures used for perception planning are estimates or expected values of the actual parameters. Therefore, most scenarios require an iterative refinement approach, since at the beginning of the sensing and world modelling activity the robot will have only an incomplete map and partial information to work with. This situation leads naturally to the need for *Attention Focussing*, where the Locus of Interest covers a large area during the preliminary reconnaissance stage, but is narrowed down to more specific regions as more information becomes available. This is again illustrated in the exploration scenario presented in Fig. 13, where after a general reconnaissance phase the attention of the mapping system is now being narrowed to specific regions to be explored further.

5 Applications

Robot Exploration

We now outline two concrete scenarios. The first is the *exploration scenario*, whose components were already discussed in previous sections. As illustrated in Figs. 6 and 13, the Dynamic Perception framework allows the robot to react to the complexity of its environment, as well as to determine optimal exploration strategies, by reasoning explicitly about what the robot needs to know, what it does know, and what it doesn't know about its environment. A second exploration example is given in Fig. 14. It should be mentioned that the exploration problem we address can be seen as a stochastic generalization of the deterministic *Art*

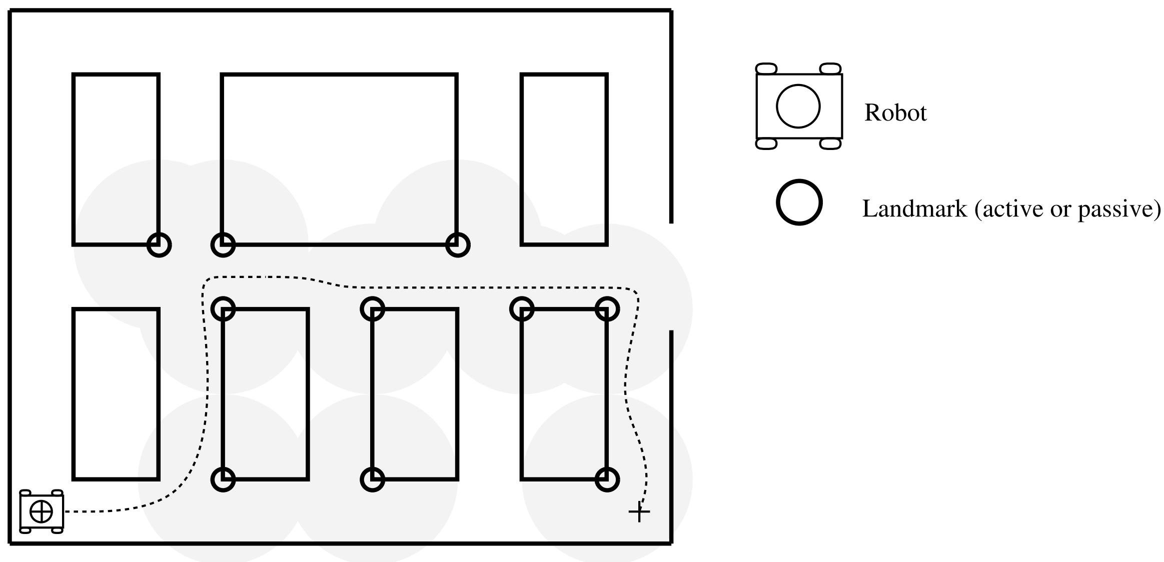

Figure 12: Locus of Observation for Landmark-Based Navigation. An environment instrumented with beacons (active landmarks) is shown. The shaded regions correspond to the Loci of Observation.

Gallery Problem [23].

Integration of Navigation and Perception

The second scenario involves *integrated perception and navigation planning*. As discussed in section 1, robot perception and robot task planning have generally been treated as separate stages of the robot's cycle of operation. This is clearly seen in the area of robot motion planning. Path planning methods have generally been limited to planning safe trajectories from a given starting point to a given goal, avoiding obstacles and taking into account the kinematic and dynamic characteristics of the robot. Other concerns, such as robot registration constraints, perceptual requirements, or environment complexity, are ignored.

The Dynamic Perception framework allows integrated navigation planning, where both perceptual and locomotion requirements can be taken into account. Simple obstacle-avoidance path planning is performed on the Inference Grid as the minimization of a dual-objective cost function [10]:

$$\min_{\mathbf{P}} \ f(\mathbf{P}) = w_d \, \text{length}(\mathbf{P}) + w_c \sum_{\forall C \in \mathbf{P}} \Gamma(C) \tag{19}$$

where $\mathbf{P}$ is the robot path, w_d is the path length weight, w_c is the cell cost weight, and $\Gamma(C) = f_c(P[s(C) = \text{OCC}])$ is the cell traversal cost, defined directly in terms of the Occupancy Grid cell state estimates.

Integrated navigation planning for perceptual and locomotion tasks can be performed on the Inference Grid as the minimization of a multi-objective cost function:

$$\min_{\mathbf{P}} \ f(\mathbf{P}) = \sum_i w_i \, c_i(M) \tag{20}$$

where the $c_i(M)$ are cost functions representing various perception and locomotion requirements, computed on the Inference Grid M, and w_i is the weight vector. An example of this approach is shown in Fig. 14,

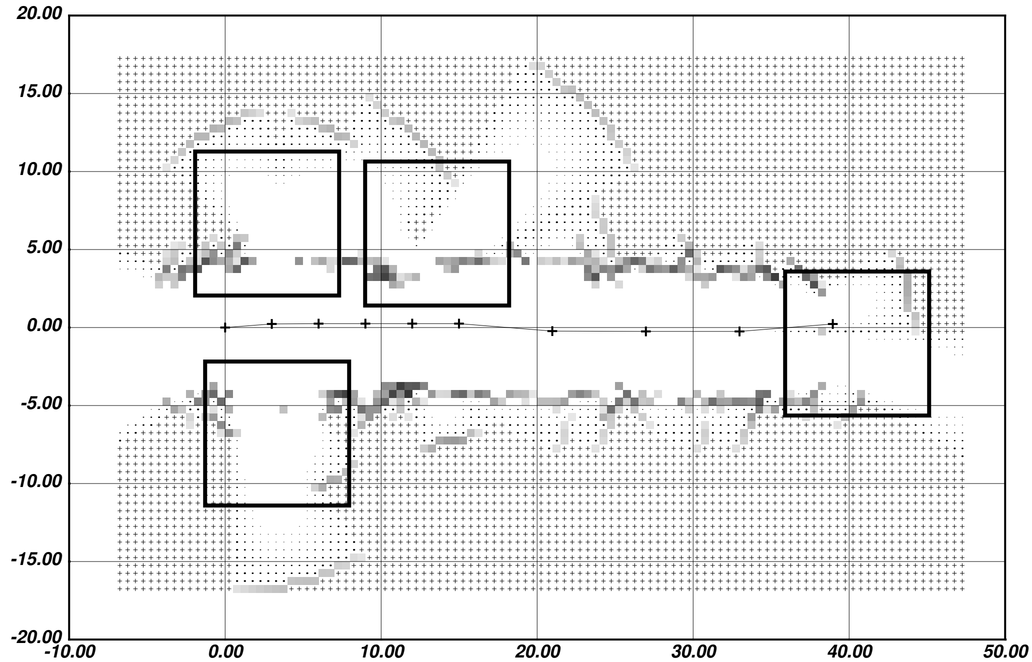

Figure 13: Controlling the Attention of the Robot. The same Occupancy Grid of Fig. 6 is shown. The robot's task is exploration, and the four areas to be investigated next are shown superimposed on the map.

where different behaviours of a mobile robot are obtained by varying the relative importance of two tasks: *exploration*, where extent and accuracy of the resulting map are important, and *courier navigation*, where finding the shortest distance to the goal is essential.

6 Conclusions

The *Dynamic Robot Perception* framework outlined in this paper stresses the active and adaptive control of the perceptual activities of an autonomous robot. This is done by explicitly determining the evolving information requirements of the different tasks being addressed by the robot, and by planning appropriate sensing strategies to recover the information required for successful completion of these tasks. We discussed the development of strategies for dynamic control of robot perception that emphasize the use of stochastic and information-theoretic sensor interpretation and world modelling mechanisms, and explored the connection between specification of a task and its information requirements.

We have performed an initial experimental validation of the components of the Dynamic Perception framework discussed in this paper. Currently, our research group is finishing the software structure and sensor interfaces for a more powerful mobile robot. This vehicle is based on an omni-directional platform, and is equipped with a number of sensors, including infrared proximity sensors, a sonar sensor array, and an optical rotating range scanner. It will be used to conduct more extensive experimental work.

Acknowledgments

The author wishes to thank José Moura for useful insights into the information control literature, and Ingemar Cox for making me aware of the Art Gallery literature. The research discussed in this paper was

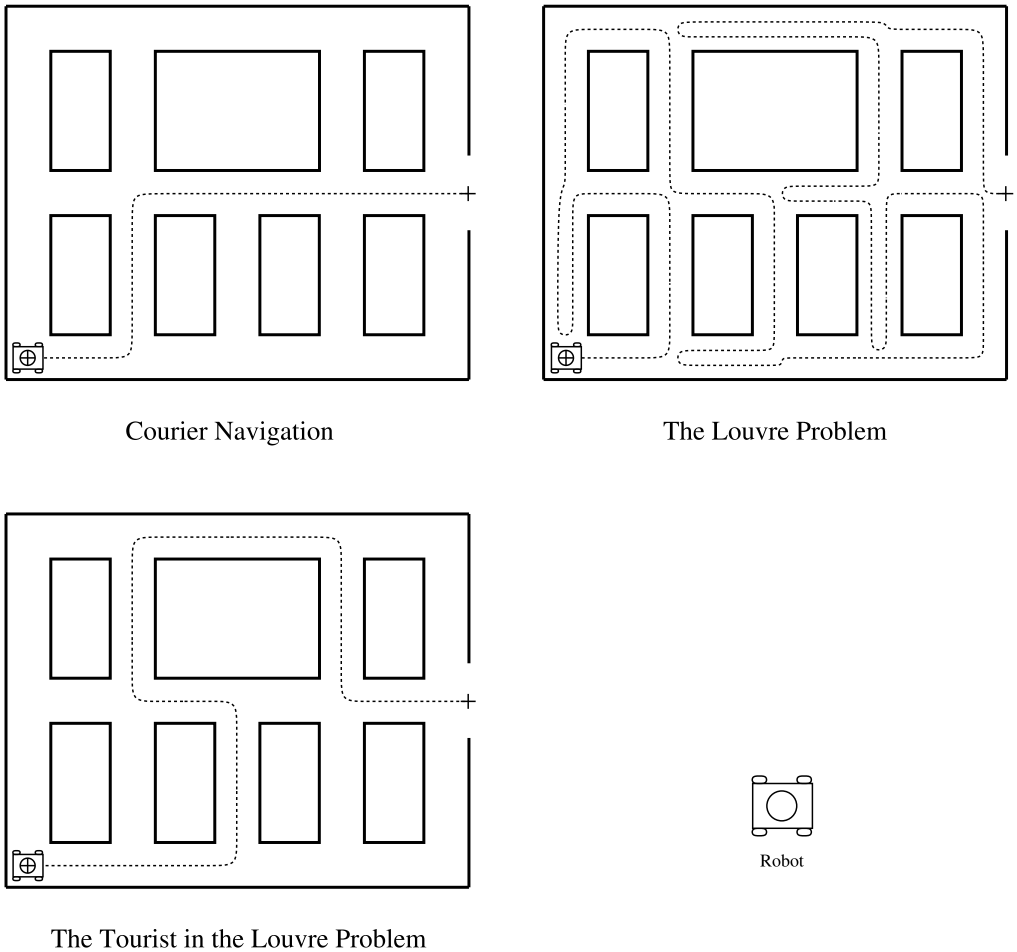

Figure 14: Integration of Perception and Navigation Tasks. The maps show the behaviour of a mobile robot in three cases: 1. Courier navigation tasks, when the goal is finding the fastest route. 2. Exploration tasks, when the primary goal is careful mapping of the robot's area of operation. 3. Integration of courier navigation and exploration, when both activities are given comparable importance, and the robot automatically adjusts its behaviour accordingly.

performed at the Intelligent Robotics Laboratory, Computer Sciences Department, IBM T. J. Watson Research Center. It incorporates some results from research performed by the author during his association with the Mobile Robot Lab, Robotics Institute, Carnegie-Mellon University.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations.

References

[1] J. Aloimonos, I. Weiss, and A. Bandyophadyay. Active Vision. *International Journal of Computer Vision*, 1(4), January 1988.

[2] R. Bajcsy. Active Perception. *Proceedings of the IEEE*, 76(8), August 1988.

[3] J. O. Berger. *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag, Berlin, 1985. Second Edition.

[4] D. P. Bertsekas. *Dynamic Programming: Deterministic and Stochastic Models*. Prentice-Hall, Englewood Cliffs, NJ, 1987.

[5] R. E. Blahut. *Principles and Practice of Information Theory*. Addison-Wesley, Reading, MA, 1988.

[6] T. Dean, K. Basye, and M. Lejter. Planning and Active Perception. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, DARPA, 1990.

[7] T. Dean et al. Coping with Uncertainty in a Control System for Navigation and Exploration. In *Proceedings of the Eight National Conference on Artificial Intelligence*, AAAI, Boston, MA, July 1990.

[8] H. F. Durrant-Whyte. *Integration, Coordination, and Control of Multi-Sensor Robot Systems*. *Kluwer International Series in Engineering and Computer Science*, Kluwer Academic Publishers, Boston, MA, 1988.

[9] J. L. Eaves and E. K. Reedy. *Principles of Modern Radar*. Van Nostrand Reinhold, New York, 1987.

[10] A. Elfes. *Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation*. PhD thesis, Electrical and Computer Engineering Department/Robotics Institute, Carnegie-Mellon University, May 1989.

[11] A. Elfes. Occupancy Grids: A Stochastic Spatial Representation for Active Robot Perception. In *Proceedings of the Sixth Conference on Uncertainty and AI*, AAAI, Cambridge, MA, July 1990.

[12] A. Elfes. Sonar-Based Real-World Mapping and Navigation. *IEEE Journal of Robotics and Automation*, RA-3(3), June 1987.

[13] A. Elfes. *Strategies for Dynamic Robot Perception Using a Stochastic Spatial Model*. Research Report, IBM T. J. Watson Research Center, 1991. In preparation.

[14] A. Elfes. A Tesselated Probabilistic Representation for Spatial Robot Perception and Navigation. In *Proceedings of the 1989 NASA Conference on Space Telerobotics*, NASA/Jet Propulsion Laboratory, JPL, Pasadena, CA, January 1989.

[15] E. Grosso, G. Sandini, and M. Tistarelli. 3-D Object Reconstruction Using Stereo and Motion. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(6), November/December 1989.

[16] G. Hager. *Information Maps for Active Sensor Control*. Technical Report MS-CIS-87-07, GRASP Lab, Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, February 1987.

[17] G. Hager and M. Mintz. Estimation Procedures for Robust Sensor Control. In *Proceedings of the 1987 AAAI Workshop on Uncertainty in Artificial Intelligence*, AAAI, Seattle, WA, July 1987.

[18] M. Hebert, T. Kanade, and In So Kweon. 3-D Vision Techniques for Autonomous Vehicles. In *Analysis and Interpretation of Range Images*, Springer-Verlag, Berlin, 1990.

[19] E. P. Krotkov. *Active Computer Vision by Cooperative Focus and Stereo*. Springer-Verlag, Berlin, 1989.

[20] L. H. Matthies. *Dynamic Stereo Vision*. PhD thesis, Computer Science Department, Carnegie-Mellon University, 1989.

[21] G. L. Miller and E. R. Wagner. An Optical Rangefinder for Autonomous Robot Cart Navigation. In *Proceedings of the 1987 SPIE/IECON Conference*, SPIE, Boston, MA, November 1987.

[22] A. V. Oppenheim and R. W. Shafer. *Discrete-Time Signal Processing*. *Prentice Hall Signal Processing Series*, Prentice Hall, Englewood Cliffs, NJ, 1989.

[23] J. O'Rourke. *Art Gallery Theorems and Algorithms*. Volume 3 of *International Series of Monographs on Computer Science*, Oxford University Press, Oxford, 1987.

[24] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York, 1984.

[25] H.-L. Wu and A. Cameron. A Bayesian Decision Theoretic Approach for Adaptive Goal-Directed Sensing. In *Proceedings of the Third International Conference on Computer Vision*, IEEE, Osaka, Japan, December 1990.