# Vision-Motion Planning of a Mobile Robot considering Vision Uncertainty and Planning Cost

Jun Miura and Yoshiaki Shirai
Dept. of Computer-Controlled Mechanical Systems, Osaka University
Suita, Osaka 565, Japan
Email: jun@mech.eng.osaka-u.ac.jp
URL: http://www-cv.mech.eng.osaka-u.ac.jp/"jun/

## Abstract

This paper proposes a planning method for a vision-guided mobile robot under vision uncertainty and limited computational resources. The method considers the following two trade-offs: (1) granularity in approximating a probabilistic distribution vs. plan quality, and (2) search depth vs. plan quality. The first trade-off is managed by predicting the plan quality for a granularity using a learned relationship between them, and by adaptively selecting the best granularity. The second trade-off is managed by formulating the planning process as a search in the space of feasible plans, and by appropriately limiting the search considering the merit of each step of the search. Simulation results and experiments using a real robot show the feasibility of the method.

## 1 Introduction

There has been an increasing interest in autonomous mobile robot which recognizes an environment with vision and moves without guidance of human operators. A key to realize such a robot is the ability to generate a plan of vision and motion operations so that a robot may efficiently reach the destination. To design a planning algorithm for such a robot, we have to consider the following two issues: limited computational resources and uncertainty in visual data. These two issues are closely related; planning based on uncertain data usually requires more computation than planning without uncertainty because multiple possible outcomes of actions should be considered, and therefore, the limitation of computational resources tends to be critical.

One of the useful tools for planning under uncertainty is statistical decision theories [Berger, 1985]. Several works applied statistical decision theory to vision and/or motion planning tasks (e.g., [Hutchinson and Kak, 1989] [Cameron and Durrant-Whyte, 1990] [Dean *et* al., 1990] [Miura and Shirai, 1993]). One drawback of these approaches is that a large branching factor of a search tree, which is determined not only by the number of possible actions but also that of possible situations that arise due to the uncertainty of sensory data, makes a planning process computationally expensive.

Regarding the limitation of computational resources, many works have recently been focusing on the concept of *limited rationality* [Russell and Wefald, 1991], in which the cost of planning is explicitly considered and the time for object-level planning is allocated so that the overall utility including both plan efficiency and planning cost is maximized. Some of examples are: flexible computation [Horvitz, 1990], decision-theoretic meta-level control of (object-level) reasoning [Russell and Wefald, 199l], and expectation-driven iterative refinement (EDIR) using anytime algorithms [Boddy and Dean, 1989].

This paper is concerned with a vision-motion planning of a mobile robot considering vision uncertainty and planning cost. Fig. 1 shows an example problem. Our mobile robot is going to the destination while avoiding obstacles. There is a route which passes the narrow space (we call it the *gate)* between the board and the partition; however the passability of the gate is initially unknown due to the uncertainty of visual data. The detour passing through the hallway is known to be passable, although it is longer. The robot estimates the gate width with stereo vision to determine the passability. The planner determines a set of observation points which efficiently navigates the robot to the destination. For this problem, we propose a planning method combines a decision-theoretic approach with consideration of planning cost.

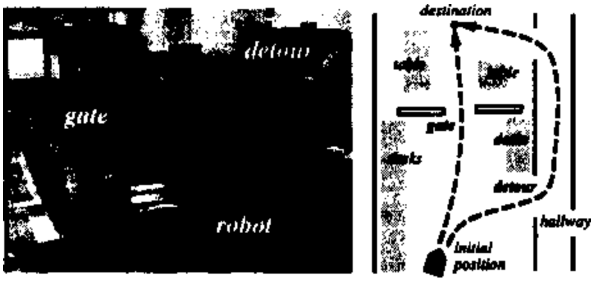


Figure 1: An example planning problem.

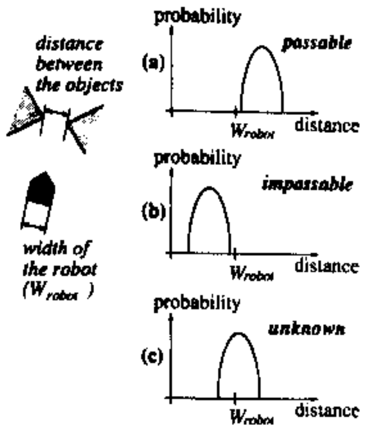


Figure 2: Three possible state of the gate.

# 2 Basic Planning Strategy

## 2.1 Plan Representation

An action is composed of a movement to the next observation point and an observation at that point. A state is represented by the current estimate of the gate width and the current robot position. Due to the uncertainty in observation results, the robot cannot determine the gate width deterministically but obtains its probabilistic distribution.

After an observation, the robot classifies the state of the gate into one of the three states (see Fig. 2): if the robot width is smaller than the minimum value of the probabilistic distribution of the gate width, the gate is passable; if the robot width is larger than the maximum value of the probabilistic distribution, the gate is impassable; otherwise the passability is unknown.

Since the actual state after an observation depends on the observation result and cannot be determined beforehand, a subplan is generated for each possible state, which is predicted using the uncertainty model of vision. Fig. 3 shows an example plan for the problem shown in Fig. 1. Such a plan is represented by an AND/OR tree; an OR node corresponds to selection of an action; an AND node corresponds to a possible state. The quality of a plan is measured in terms of its execution cost, which is the *expectation* of the total execution time for movement and observation.

The leaves of an AND/OR tree are either *terminal node* or *open node.* At a terminal node, the passability is decided without uncertainty and, thus, the final action (i.e., passing the gate or taking the detour) is also decided. At an open node, since the passability is unknown, the final action has not been decided yet. A plan candidate is refined by expanding (making subplans for) one of its open nodes. In expansion of an open node, the possible range of the gate width is discretized with some *granularity,* and a subplan is generated for each discretized state.

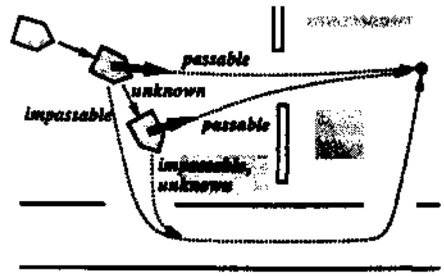


Figure 3: An example plan. Dotted arrows indicate possible movements after observation. Bold arrows indicate observation of the gate.

## 2.2 Computational Trade-offs to be Considered

The planning method is designed to deal with the following two computational trade-offs:

search depth vs. plan quality: This trade-off has been investigated by several researchers (e.g., DTA* by Russell and Wefald [1991]). We consider this trade-off in generating a multi-step plan.

granularity vs. plan quality: A finer granularity for discretization improves plan quality, but it increases the planning cost. This trade-off has little been considered, although it is important in planning under uncertainty.

The first trade-off is managed by formulating the planning process as an iterative refinement process [Boddy and Dean, 1989], and by appropriately limiting the iteration. To cope with the second trade-off, we represent the relationship between granularity and the expectation of the reduction of a plan cost (we call this expectation a *plan improvement)* as a *performance profile* [Dean and Boddy, 1988] [Zilberstein, 1993], and then determine the best granularity by examining the performance profile and the cost of node expansion. With consideration of these trade-offs, the planner tries to minimize the total cost of plan generation and plan execution.

# 3 Formulation as Iterative Refinement

## 3.1 Easily-Obtainable Feasible Plan

In an iterative refinement framework, the planner searches the space of feasible plans (executable plans) for the final plan. This formulation entails an *easily-obtainable* feasible plan for any open node. There are two such feasible plans. One of them is to take the detour from the current position; this feasible plan is usually costly. Thus we use the other feasible plan:

> *The robot moves from the current position to the position just before the gate[1]. If the gate*

[1] At this position, the robot is assumed to be able to measure the gate width without uncertainty; this position is called the *zero-uncertainty point,* indicated as $x^*$.

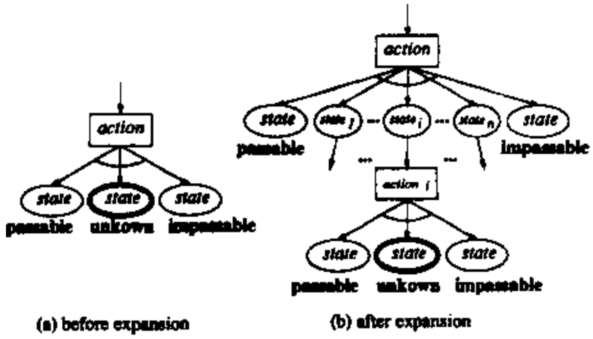


Figure 4: Expansion of an open node of a plan candidate. Ellipses drawn with bold lines indicate open nodes.

> *is passable, the robot passes it; if not, the robot takes the detour from that position.*

Each plan candidate has the *temporary cost,* $C^{temp}$, which is obtained by *temporarily* assigning the above feasible subplan to all of its open nodes.

## 3.2 Selection of Open Node to Expand

In the planning process, the most promising open node is selected and expanded. Before expansion, an unknown state is treated as one open node and has a feasible plan with it (see Fig. 4(a)). The expansion of the open node consists of searching for the best action for each discretized state and assigning the feasible subplan to newly generated open nodes (see Fig. 4(b)).

The open node to be expanded is determined as follows. Suppose we can predict the maximum plan improvement (i.e., cost reduction) $\Delta C$ of a plan candidate, which will be obtained by expanding its best open node[2]. Then, the *merit* of expanding the node is given by subtracting the expansion cost $C^{exp}$ from the predicted plan improvement $\Delta C$.

Let $C_i^{temp}$, $\Delta C_i$, and $C_i^{exp}$ denote the temporary cost, the predicted plan improvement, and the expansion cost of a plan candidate $PC_i$, respectively. Then, we define the *new cost* $C_i^{new}$ as:

$$C_i^{new} = C_i^{temp} - (\Delta C_i - C_i^{exp}). \tag{1}$$

Let $C_{FP^*}$ be the cost of the incumbent (the best feasible plan among those which have been obtained so far). During the planning, the plan candidates are kept whose new costs are less than $C_{FP^*}$. Among the candidates, the plan candidate $PC_{i^*}$ which is to be expanded next is determined as:

$$PC_{i^*},\ i^* = \arg\min_i C_i^{new}. \tag{2}$$

The planner iteratively selects $PC_{i^*}$ and expands its best open node. This planning process has the *anytime* [Dean and Boddy, 1988] property. The iteration stops when $C_{i^*}^{new}$ is larger than $C_{FP^*}$.

[2]The method to predict the plan improvement will be described in Section 4.

## 3.3 Consideration of Meta-Planning Cost

The above algorithm will expand a node as long as $C_{i^*}^{new}$ is less than $C_{FP}$* even if their difference is very small. Expansion in such a case, however, may be useless if the meta-planning cost is high.

Thus, we slightly modify the termination condition. That is, if the difference is less than the meta-planning cost, the iteration process stops and the best feasible plan is returned as the final plan. The cost of meta-planning is, at present, considered to be constant.

## 3.4 Deciding Only the Next Action

When the objective of the planner is not to generate a whole plan but to select the best next action (in a dynamic environment, for example), the algorithm is altered as follows.

Let $C_{Fp^*A}$ be the cost of the best feasible plan which starts with action *A.* If the smallest value of $C_i^{new}$ of plan candidates which start with actions other than *A* is larger than $C_{FpA}$*, the planning process terminates and returns *A* as the best next action. This strategy is similar to that of DTA* [Russell and Wefald, 1991].

# 4 Determining the Best Granularity

We have mentioned that the granularity in discretizing the ranges of random variables (the gate width in our case) directly affects both the plan improvement and the expansion cost. The selection of the best granularity is, therefore, crucial to managing the trade-off between planning cost and plan efficiency.

This paper proposes to represent the relationship between the granularity and the predicted plan improvement as a performance profile, and to calculate the best granularity and the merit of expansion simultaneously.

Currently we equally divide the range of a variable; thus, the granularity is specified with the number of divisions. Let $PI(n)$ denote the predicted plan improvement for granularity $n$.

The best granularity $n^*$ is given by

$$n^* = \arg\max_n \{PI(n) - C^{exp}(n)\}, \tag{3}$$

where $C^{exp}(n)$ is the cost of expansion with granularity n, which is defined as follows:

$$C^{exp}(n) = N_{cand} * C^{exam} * n, \tag{4}$$

where $N_{eand}$ is the number of action candidates; $C^{exam}$ is the cost required for examining one action candidate. Fig. 5 illustrates the determination of the best granularity. Once the best granularity n* is determined, the merit of expansion is calculated as $PI(n^*) - C^{exp}(n^*)$.

If there is at least one OR node between the root node and an open node, the probability of reaching the open node is less than one. In this case, we multiply *PI(n),* which is originally generated for the case that the reaching probability is one, by the current reaching probability.

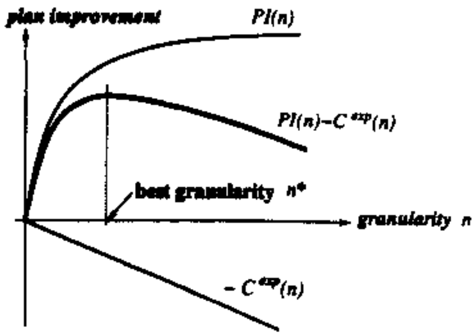


Figure 5: Determining the best granularity. This figure is based on [Horvitz, 1990].

Every time an open node is generated (by expansion of its parent), the best granularity for the node is determined. We include the cost of determining the best granularity in the cost of expanding the parent node.

# 5 Deriving Performance Profile

Derivation of a performance profile (PP) is one of the important issues in anytime algorithm-based approaches. In some cases, PPs may be obtained from the structural analysis of the problem; in other cases, PPs may be obtained from experimental data. It is, however, usually difficult to obtain PPs for complex problems only with one of these methods. We, thus, derive a PP through structural *and* experimental analysis of the planning problem in the following steps:

1. Analyze the structure of the planning problem and extract important problem parameters which can reasonably characterize the PP.
2. Construct a *generalized* PP using the parameters obtained above; a generalized PP has coefficients to be estimated.
3. Calculate actual PPs for an enough number of problem parameter sets, and adjust the coefficients of the generalized PP so that the PP fits well to the actual data set.

## 5.1 Problem Analysis

The plan improvement is the difference of temporary costs of a plan candidate before and after expanding one of its open nodes. Let us calculate the plan improvement using an example situation shown in Fig. 6.

Suppose that the robot is initially at $x_0$ and the next observation point is $x_1$. The open node under consideration is the state that the gate's passability is unknown after the observation at $x_1$. The feasible plan before expansion is to go from $X_1$ directly to the zero-uncertainty point $x^*$, where the robot can measure the gate width without uncertainty. Expansion of this open node, i.e., selection of a second observation point $x_2$ for each discretized state results in a new feasible plan.

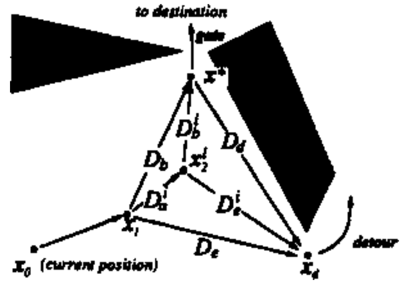


Figure 6: An example situation.

Using the cost of movement between points defined in Fig. 6 (for example, $D_a^i$ is the cost of movement from $x_1$ to $x_2^i$), the plan improvement $\Delta C$ is described as (see Appendix A):

$$\Delta C = \sum_{i=1}^{n} \begin{pmatrix} P_{x_2^i,ng}(D_b + D_d - D_a^i - D_e^i) \\ -P_{x_2^i,ud} C_v \\ -(P_{x_2^i,ok} + P_{x_2^i,ud})(D_a^i + D_b^i - D_b) \end{pmatrix}, \quad (5)$$

where $C_v$ is the cost of one observation; $P_{x_2^i,ok}$ ($P_{x_2^i,ng}$, $P_{x_2^i,ud}$) is the probability that the gate's state is passable (impassable, unknown) after observation at $x_2^i$. We here suppose that $x_2^i$ indicates the best second viewpoint for the $i$th divided state.

This equation represents the trade-off between the effect of visual information to be obtained by observation at $x_2^i$ and the cost of observation including that of movement to the observation point. The first term inside the parentheses is the effect of visual information, which is gained by deciding to take the detour at $x_2^i$, instead of decision at $x^*$, in case that the gate is impassable. The second term is the cost of the extra observation at $x^*$ in case that the passability is still unknown at $x_2^i$. The third term is the extra movement cost to visit $x_2^i$.

## 5.2 Designing Generalized Performance Profile

We examine equation (5) in order to design a generalized PP which can reasonably approximate the equation.

The actual values which depend on $x_2^i$ (e.g., $P_{x_2^i,ng}$) cannot be obtained before performing the search for $x_2^i$; it may also be hard to reliably predict these values. Thus, we would like to predict the plan improvement using only parameters whose values are available.

We use an upper bound of plan improvement for constructing a generalized PP. To calculate the upper bound, we first employ the concept of the *assumption of perfect sensor information* [Miura and Shirai, 1993], which is based on the fact that a plan generated using certain information is always more efficient than plans generated using uncertain information.

Assuming that the observation result at $x_2^i$ includes no uncertainty (i.e., $P_{x_2^i,ud} \to 0$), a modified plan improvement $\Delta C'$ is given by

$$\Delta C' = \sum_{i=1}^{n} P_{x_2^i, ng}(D_b + D_d - D_a^i - D_e^i) - \sum_{i=1}^{n} P_{x_2^i, ok}(D_a^i + D_b^i - D_b). \quad (6)$$

Since $\Delta C'$ still includes values depending on $x_2^i$, we then find its maximum value. $\Delta C'$ is maximized by letting $x_2^i \to x_1$ (i.e., $D_b^i \to D_b$, $D_e^i \to D_e$, and $D_a^i \to 0$); finally, we obtain the following upper bound:

$$\Delta C_{ub} = P_{x_1 \to x^*, ng}(D_b + D_d - D_e), \quad (7)$$

where $P_{x_1 \to x^*, ng}$ is the probability that the gate is impassable after observations at $x_1$ and $x^*$, and is equal to $\sum_{i=1}^{n} P_{x_2^i, ng}$ under the assumption of perfect sensor information. We suppose that the actual plan improvement $\Delta C$ increases as $\Delta C_{ub}$ increases.

We also reasonably assume that the plan improvement is a monotonically increasing function of the granularity $n$ and asymptotically approaches some limit value.

Based on the above examination, we eventually decided to use the following generalized PP (function $PI(n)$) as an approximation of equation (5):

$$PI(n) = K(1 - e^{-k_1 n}), \quad (8)$$

$$K = k_2 \cdot \Delta C_{ub}{}^{k_3}, \quad (9)$$

where $k_1$, $k_2$, and $k_3$ are coefficients to be estimated from the experimental data.

### 5.3 Determining PP through Experiments

We found that the coefficients in $PI(n)$ varies as the initial position $x_0$ changes. Thus, we divided the work space of the robot, which is the triangle composed of $x_0$, $x^*$, and $x_d$ (see Fig. 6), into several regions and generated a PP for each divided region. The size of one region is determined empirically.

In each region, its center is selected as $x_0$. Then, we calculated actual PPs about a hundred problems with various $x_1$'s and gate widths. We limited $n$ to one of $\{1, 3, 5, 7, 9, 11\}$.

In estimating the coefficients, we first fitted equation (8) to each PP and estimated $K$ and $k_1$. Fig. 7 shows an example of actual PP and the fitted curve. Then, we examined the relationship between $K$ and $\Delta C_{ub}$ in equation (9). Fig. 8 shows a log-scaled plot of these values. By fitting a line to the plotted data, we obtained $k_2 = 0.033$ and $k_3 = 1.319$ in this case. As for $k_1$, we adopted the mean of all $k_1$'s because we could not find any strong correlation between $k_1$ and other problem parameters.

## 6 Simulation Results

Fig. 9 shows examples of generated plan for a planning problem. From the initial observation result of the gate width, a complete plan, which is an AND/OR tree of observation points, was planned off-line. Observation points are limited only to grid points set on the work space of the robot.

In order to show that the planner can adaptively determine the planning time (search depth and granularity) according to its computational power, we performed

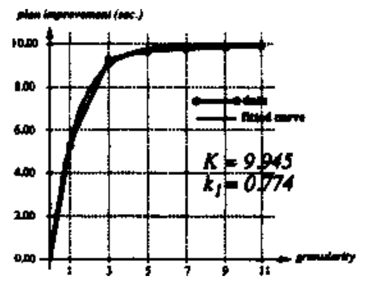


Figure 7: An example performance profile and the fitted curve.

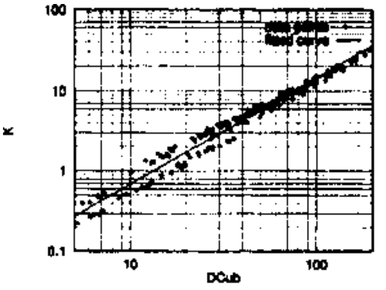


Figure 8: Log-scaled plot of data for equation (9).

simulations for the same problem with different computing powers ($C^{exam}$'s, see equation (4)). Notice that the higher the computing power is (the smaller $C^{exam}$ is), the more precise and the less costly plan is generated. The difference of the costs is, however, rather small compared with that of the computing powers. This is because that the initial feasible plan is already a good plan in this case.

We then compared, in terms of the total of planning cost and execution cost, the proposed method with the *fixed* method, which uses a fixed granularity and a fixed search depth. Fig. 10 shows a comparison result. The proposed method gives the best performance.

We conducted this comparison for about forty different problems; in about two-thirds cases, the proposed method outperformed others. In other cases, the difference between the best result and the result by the proposed method was less than one percent of the best result. Note that in order to obtain the best result with the fixed method, we had to adjust the granularity and search depth for each problem, while the proposed method always performed well without any such adjustments.

## 7 Experiments using Real Robot

This section describes preliminary results for a real planning problem shown in Fig. 1. The gate width is estimated by the two vertical segments on both sides; its probabilistic distribution is calculated using an uncertainty model of stereo vision [Miura and Shirai, 1993].

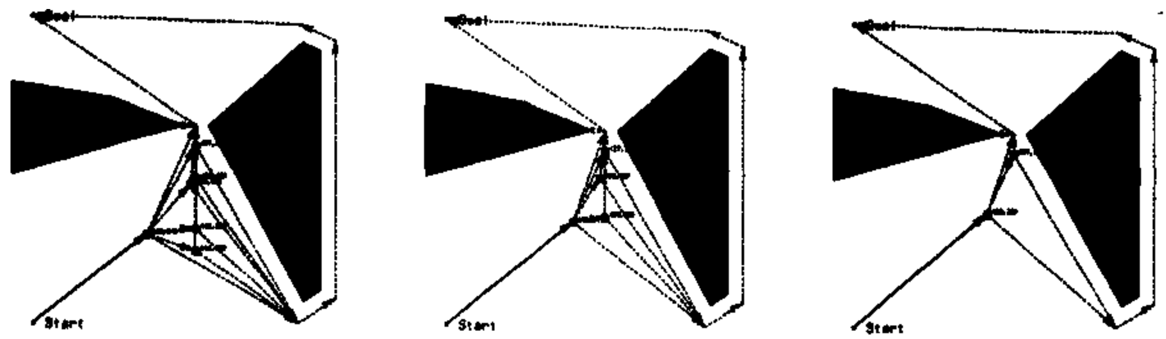


(a) $C^{exam}=1.1e\text{-}4$, $Cost=186.9$, $n^*=7$ (b) $C^{exam}=1.1e\text{-}3$, $Cost=188.6$, $n^*=3$ (c) $C^{exam}=1.1e\text{-}2$, $Cost=189.0$, $n^*=0$

Figure 9: Simulation results: circles indicate observation points; solid arrows indicate the path to the next observation point; dotted arrows indicate possible paths after observations, $n^*$ is the granularity used for discretizing the open node at the next observation point. In case of (c), the initial feasible plan is used.

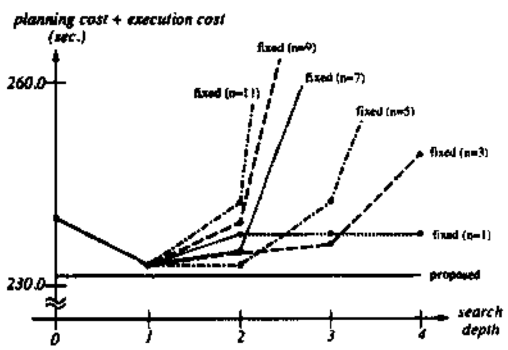


Figure 10: Comparison of the proposed method with fixed-granularity and fixed-search depth methods. Change of the total cost according to the change of the search depth is indicated for each granularity.

In the experiment, the planner decides only the next action (see Section 3.4). Then the robot moves to the planned observation point and observes the gate. The planning and action operations are iteratively performed until the passability of the gate is determined.

Fig. 11 shows the result of a trial. From the initial image (at lower-right in the figure), the robot estimated the probabilistic distribution of the gate width; the probability of the gate being passable was 0.53. Then the planner determined the next observation point as shown in the map. The robot moved to the next observation point (lower-left) and obtained the new image (upper-right). After this observation, the gate was determined to be passable; thus the robot moved forward and passed the gate (upper-left).

## 8 Conclusions and Discussion

We have proposed a planning method for a vision-motion planning of a mobile robot under vision uncertainty and limited computational resources. We managed *granularity vs. plan quality* and *search depth vs. plan quality* trade-offs by: (1) considering the relationship between granularity and plan improvement which is represented as a performance profile, and (2) formulating the process of generating a multi-step plan as an iterative refinement process. The proposed method is always comparable with the best of the methods which uses fixed granularity and fixed search depth. We have also described a method to derive a performance profile through structural *and* experimental analysis of the planning problem.

In the performance profile (PP)-based planning, the quality of PP has the largest importance. If the parameters (e.g., configuration of obstacles) of the current problem are largely different from those of problems used in derivation of PPs, the predicted plan improvement may not be accurate enough. In this paper, in order to increase the accuracy, we divided into the problem space into several regions and obtained a PP for each region. Such a strategy might be practical if a PP which is applicable to a wide problem space is hard to find.

This paper has treated a relatively simple vision-motion planning problem (a *single-gate problem).* A future work is to apply the proposed method to large planning problems which have many *gates* to observe. To solve such a large problem, it is necessary to decompose the problem into a set of single-gate problems. We are now developing an efficient decomposition method.

## Acknowledgments

This work is supported in part by Grant-in-Aid for Scientific Research from Ministry of Education, Science, Sports, and Culture, Japanese Government and by the Okawa Foundation for Information and Telecommunications, Tokyo, Japan.

## A Derivation of Plan Improvement $\Delta C$

Only the subplan for the open node at $X_1$ changes by expansion. The cost $C_1$ of the subplan before expansion

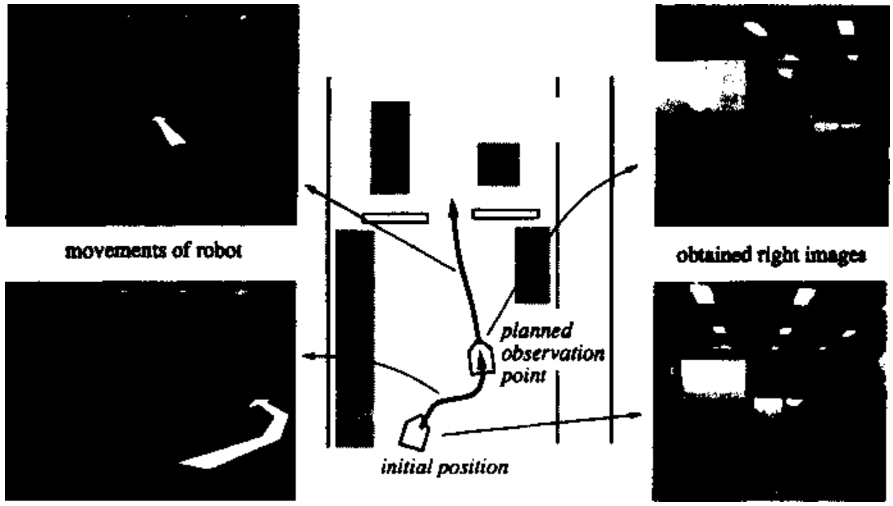


Figure 11: An experimental result: [center] the planned observation point and the trajectory taken by the robot; [lower-right] the right image obtained at the initial position; [upper-right] the right image obtained at the next observation point; [lower-left] the movement from the initial position to the next observation point; [upper-left] the movement from the next observation point into the gate.

is given by

$$\begin{aligned} C_1 &= P_{x_1,ud}(D_b + C_v) + P_{x_1 \to x^*} D_{goal} \\ &+ P_{x_1 \to x^*,ng}(D_d + D_{detour}), \end{aligned} \tag{10}$$

where $D_{goal}$ and $D_{detour}$ is the cost of movement from $x^*$ to the destination and from $x_d$ to the destination through the detour, respectively; $P_{x_1 \to x^*,ok}$ ($P_{x_1 \to x^*,ng}$) is the probability that the gate is passable (impassable) at $x^*$ in case that the passability is unknown at $x_1$. The cost $C_2$ after expansion is given by

$$C_2 = \sum_{i=1}^{n} \min_{x_2^i} C_2^i, \tag{11}$$

$$C_2^i = \begin{pmatrix} P^i_{x_1,ud}(D_a^i + C_v) + P_{x_2^i,ok}(D_b^i + D_{goal}) + \\ P_{x_2^i,ng}(D_e^i + D_{detour}) + P_{x_2^i,ud}(D_b^i + C_v) + \\ P_{x_2^i \to x^*,ok} D_{goal} + \\ P_{x_2^i \to x^*,ng}(D_d + D_{detour}) \end{pmatrix}, \tag{12}$$

where $P^i_{x_1,ud}$ is the probability of the $i$th divided state; $P_{x_2^i \to x^*,ok}$ ($P_{x_2^i \to x^*,ng}$) is similarly defined as $P_{x_1 \to x^*,ok}$ ($P_{x_1 \to x^*,ng}$). The best $x_2^i$ is selected for each state $S^i$.

The plan improvement is obtained as $C_1 - C_2$. The following relations hold among probabilities:

$$\begin{aligned} &P_{x_1,ud} = \textstyle\sum_{i=1}^{n} P^i_{x_1,ud}, \\ &P_{x_1 \to x^*,ok} = \textstyle\sum_{i=1}^{n} P_{x_2^i,ok} + \sum_{i=1}^{n} P_{x_2^i \to x^*,ok}, \\ &P_{x_1 \to x^*,ng} = \textstyle\sum_{i=1}^{n} P_{x_2^i,ng} + \sum_{i=1}^{n} P_{x_2^i \to x^*,ng}, \\ &P^i_{x_1,ud} = P_{x_2^i,ok} + P_{x_2^i,ng} + P_{x_2^i,ud}. \end{aligned} \tag{13}$$

After a series of calculation using the above relations, we eventually obtain equation (5).

## References

[Berger, 1985] James O. Berger. *Statistical Decision Theory and Bayesian Analysis.* Springer-Verlag, second edition, 1985.

[Boddy and Dean, 1989] M. Boddy and T. Dean. Solving time-dependent planning problems. In *Proceedings of IJCAI-89,* pages 979-984, 1989.

[Cameron and Durrant-Whyte, 1990] A. Cameron and H. Durrant-Whyte. A bayesian approach to optimal sensor placement. *Int. J. of Robotics Res.,* 9:70-88, 1990.

[Dean and Boddy, 1988] T. Dean and M. Boddy. An analysis of time-dependent planning. In *Proceedings of AAAI-88,* pages 49-54, 1988.

[Dean *et al*., 1990] T. Dean, T. Basye, and S. Lejter. Planning and active perception. In *Proceedings of 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control,* pages 271-276, 1990.

[Horvitz, 1990] E.J. Horvitz. *Computation and Action Under Bounded Resources.* PhD thesis, Stanford University, 1990.

[Hutchinson and Kak, 1989] S.A. Hutchinson and A.C. Kak. Planning sensing strategies in a robot work cell with multi-sensor capabilities. *IEEE Trans, on Robotics and Automat,* RA-5(6):765-783, 1989.

[Miura and Shirai, 1993] J. Miura and Y. Shirai. An uncertainty model of stereo vision and its application to vision-motion planning of robot. In *Proceedings of the 18th Int. Joint Conf. on Artificial Intelligence,* pages 1618-1623, Chambery, France, August 1993.

[Russell and Wefald, 1991] S. Russell and E. Wefald. *Do The Right Thing.* The MIT Press, 1991.

[Zilberstein, 1993] S. Zilberstein. *Operational Rationality through Compilation of Anytime Algorithm.* PhD thesis, University of California at Berkeley, 1993.