Bayesian Grasping

Kenneth Y. Goldberg and Matthew T. Mason
School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

This paper describes a Bayesian approach to the problem of autonomous manipulation in the presence of state uncertainty. We model uncertainty with a probability distribution on the state space. Each plan (sequence of actions) defines a mapping on the state space and hence a posterior probability distribution. We search for a plan to optimize *expected* performance.

We apply the Bayesian framework to a grasping problem. Assuming a planar polygon whose initial orientation is described by a uniform distribution and a frictionless parallel-jaw gripper, we automatically plan a sequence of open-loop squeezing operations to reduce orientational uncertainty and grasp the object. Although many different performance measures are possible depending on the application, we illustrate the approach by searching for plans that optimize the robot's *expected throughput*.

We motivate the analysis of a frictionless parallel-jaw gripper with a mechanical arrangement that reduces friction along one gripper axis.

1. Introduction

Uncertainty is a central problem in the analysis and synthesis of manipulator programs. In particular, we are often uncertain about the initial state of the system we want to manipulate. One approach is to treat actions as mappings between *sets* of states and search for a plan that succeeds for every element in the initial set. There are two problems with this approach. First, planning is dominated by the worst-case outcomes – there is no way to construct a plan to optimize *average-case* performance. Second, planning does not fail gracefully – if we cannot find a guaranteed plan, there is no way to rank the alternatives.

In this paper we model uncertainty with a probability distribution on the state space and define a cost function on actions. We search for plans that optimize the expected outcome. We call this a *Bayesian* approach since the idea that we can specify a prior probability and use it to find the posterior probability is a fundamental tenet of Bayesian statistics (Berger, 1985).

We apply this approach to the problem of grasping a polygonal object with a frictionless parallel-jaw gripper. When the object is squeezed, it will rotate until at least one edge is aligned with the gripper's jaws. The object's orientation is kinematically constrained by the sequence of squeezing actions – sensing is performed at the end of the sequence to verify the final orientation. We search for plans that will orient and grasp the object.

Figure 1 shows an example plan. Assuming that object orientation is initially described by a uniform distribution and assuming that the two jaws make simultaneous contact, we can construct the posterior probability distribution after each step in the plan. In the example, the first step reduces the number of possible orientations to 6, the second reduces the number to 4, and the third step reduces the number of possibilities to 2. Due to gripper symmetry, there is no sequence that can orient an object beyond its 180° symmetry.

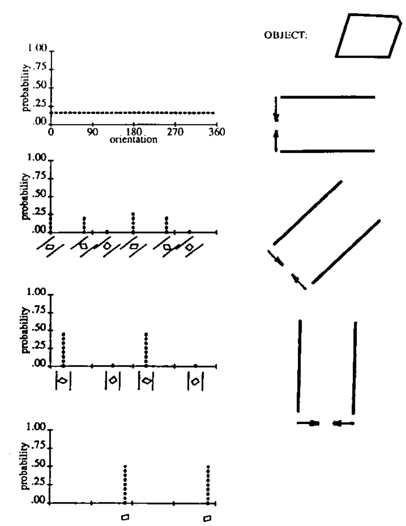

Figure 1: Evolution of the probability distribution for a 5-sided object. In this case, three squeezing steps transform a uniform distribution to a distribution with value 0.5 at each of two orientations. Symmetry in the gripper makes further improvement impossible.

The example illustrates two advantages of the Bayesian approach:

- We can consider average-case performance of plans. For example, in a manufacturing environment, the squeezing plan might be followed by a filter that accepts one orientation and rejects all others. If one time unit is required for each squeezing step and for the filtering step, then the expected time for the two-step plan is 6.2 time units while the expected time for the three-step plan is 8.0 time units. Hence we prefer the two-step plan even though it has lower probability of success.
- Although no plan is guaranteed to orient the part uniquely, we can compare a two-step plan that succeeds with probability .48 with a three-step plan that succeeds with probability .50. The extra step in the three-step plan increases the probability of success. Since the increase is small, we may not be willing to "spend" this extra step.

In this paper we consider uncertainty in state and assume that commanded actions are deterministic. The Bayesian approach can also also be used to model control uncertainty (Goldberg, 1989). The paper is organized as follows. In the next section we review related work. Next we specify our Bayesian planning framework. We then apply this framework to a class of parallel-jaw grasping problems and give numerical examples. Finally we conclude with a discussion and suggestions for future work.

2. Related Work

This section reviews work on manipulation planning with uncertainty and on grasping. Most of the previous work models uncertainty as a set of possible states, which we refer to as *possibilistic* models.

2.1. Possibilistic models of uncertainty

The mechanical designer's approach to uncertainty is to specify worst-case "tolerance" margins for each component and then guarantee performance of the assembly when these tolerances are maintained (Requicha, 1983). Taylor (1976) and Brooks (1982) used tolerances in the analysis and synthesis of manipulator programs. Lozano-Perez *et al.* (1984) describes a similar approach to fine motion planning, which was further refined by Mason (1984), Erdmann (1984), Buckley (1987), and Donald (1987). See Latombe (1989) for a review. Brost (1988) applied a possibilistic model to parallel-jaw grasping, Erdmann and Mason (1986) to tray-tilting plans, and Peshkin and Sanderson (1988) to a conveyor belt orienting system. Taylor *et al.* (1987) and Mason *et al.* (1988) extended the grasping and tray-tilting work to incorporate sensing.

For the most part, planners based on possibilistic models have adopted a "guaranteed" approach to planning: they search for plans that are guaranteed to achieve a specified goal for any of the possible initial states. One difficulty with the guaranteed approach is that "it is not always possible to find plans that are guaranteed to succeed. In the presence of model error, such plans may not even exist" (Donald, 1987, page 2). In response, Donald formalized error detection and recovery (EDR) plans that are guaranteed either to succeed or to fail recognizably. That is, we can accept plans that are not guaranteed to succeed as long as we are guaranteed to recognize when they fail. Donald has suggested that EDR plans might be compared based on their *probability* of failure.

2.2. Probabilistic models of uncertainty

Probabilistic techniques are commonly used in industrial automation. The vibratory bin feeder presents randomly oriented parts to a mechanical filter that admits only parts with a desired orientation. Boothroyd *et al.* (1972) developed a probabilistic model of this process, where the probability of each orientation is related to the surface area underneath that orientation. Erdmann (1989) explored randomized manipulation strategies and described many other cases where injecting randomness can improve average-case performance.

Strategies that optimize an expected performance criterion are encountered in both decision theory (DeGroot, 1970; Berger, 1985) and stochastic optimal control theory (Stengel, 1986; Bertsekas, 1987). In decision theory the optimal strategy is a function of sensory data. Stochastic optimal control theory generally treats state uncertainty as additive noise, often Gaussian, as in the well-studied class of L-Q-G problems (Linear system, Quadratic loss function, Gaussian noise). We borrow notation from both disciplines.

Probabilistic models have been used to combine information from multiple sensors. Smith and Cheeseman (1986) applied estimation theory to combine multiple observations for robot navigation. Durrant-Whyte (1988) argued that Bayesian probability theory offers a unified approach to combining and transforming uncertain geometric models. Hager (1988) considered computational issues in sensor fusion and developed a finite-element implementation of Bayes' theorem. See also (Cheeseman, 1985; Cameron and Durrant-Whyte, 1989; Hager and Mintz, 1989; Hutchinson and Kak, 1989).

2.3. Other Models of Uncertainty

In highly unstructured environments such as Mars or the average household, where interactions are non-repetitive, planning can be avoided (Lumelsky, 1987) or preceded by active *exploratory* procedures such as contour following (Allen, 1987; Stansfield, 1987; Koutsou, 1988; Bajcsy *et al.*, 1989). There is also an extensive body of research on high-level planning with uncertainty described in the artificial intelligence literature; Genesereth and Nilsson (1987) gives further references.

2.4. Grasping

Hanafusa and Asada (1977) used three frictionless fingers to grasp shapes in the plane, using the system potential to locate stable and convergent grasps. Mason (1986a) explored parallel-jaw grasping of shapes in the plane, and the underlying mechanics of pushing. Brost (1988) developed a systematic approach to plan parallel-jaw grasping operations; starting with an interval of possible object orientations, Brost's operations simultaneously grasp the object and eliminate the uncertainty in a single grasp. Such a grasp does not always exist. Taylor *et al.* (1987) and Mason *et al.* (1988) extended Brost's system to multi-step plans, and also incorporated sensing of jaw separation. See Pertin-Troccaz (1989) for a recent review of research on grasping.

3. Bayesian Framework

In this section we formalize the problem of choosing the best plan as a statistical optimization problem. We shall restrict ourselves to cases where the state and action spaces are finite. The six-tuple $< \Theta, f_1, \mathcal{A}, G, C, \theta^* >$ defines an instance of a planning problem.

- Θ, a set of states (configurations).
- $f_1 : \Theta \rightarrow [0,1]$, a prior probability distribution on Θ.
- $\mathcal{A}$, a set of actions (commands).
- $G : \Theta \times \mathcal{A} \rightarrow \Theta$, a transfer function that maps a state and an action onto a next state.
- $C : \mathcal{A} \rightarrow \Re$, a cost function for actions.
- θ^*, a desired final state.

For a known initial state, θ_1, each action $a \in \mathcal{A}$ defines a posterior state, $\theta_2 = G(\theta_1, a)$. A *plan* is a sequence of actions, $(a_1, a_2, \ldots, a_n)$ followed by a verification step. We will use the symbol a to denote either a single action or a plan. An n-step plan defines a composite function that maps θ_1 to θ_{n+1}. We define the cost of a plan, $C(a)$, as the sum cost of its actions plus the verification step. We represent uncertainty in the initial state with a prior probability distribution on the state space, $f_1(\theta_1)$. A plan defines a posterior probability distribution, $f_n(\theta_n|a) = \sum_{\{\theta_1|\theta_n=G(\theta_1,a)\}} f_1(\theta_1)$.

The verification stage at the end of the plan causes it to iterate until a desired state is achieved. We define the *expected cost* of a plan as the cost of the plan times the expected number of iterations. The expected number of iterations for a plan that succeeds with probability p is $1/p$. If the desired state is θ^*, we can define the probability that a plan is successful as $p = f_n(\theta^*|a)$. Then the expected cost of the plan is $C(a)/p$. A Bayes' plan is one with minimal expected cost (note that there may be more than one Bayes' plan).

4. Application to Parallel-Jaw Grasping

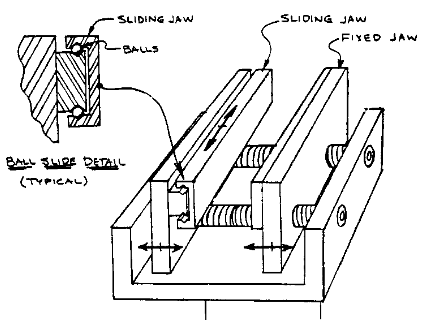

Figure 2: A "frictionless" parallel-jaw gripper. A linear bearing reduces friction between the jaws (drawing by Ben Brown).

4.1. Assumptions

In the remainder of the paper we apply the Bayesian framework to the problem of grasping with a parallel-jaw gripper, which we think of as a two-dimensional problem. We assume that:

1. The gripper has two linear jaws arranged in parallel.
2. Objects to be grasped are known rigid planar polygons.
3. The object's initial position is unconstrained as long as it lies somewhere between the two jaws. The object remains between the jaws throughout grasping; hence any polygon is equivalent to its convex hull.
4. All motion occurs in the plane and is slow enough that inertial forces are negligable. The scope of this *quasi-static* model is discussed in (Mason, 1986b) and (Peshkin, 1986).
5. The direction of squeezing is always perpendicular to the jaws.
6. There is zero friction between object and the jaws.
7. Both jaws make contact simultaneously (pure squeezing). Once contact is made between a jaw and the object, the two surfaces remain in contact throughout the grasp. A grasp continues until further motion would deform the object.

The first four assumptions were used by Brost (1988), Taylor *et al.* (1987), and Mason *et al.* (1988). The latter three assumptions simplify the analysis and improve the combinatorics of the search. By restricting gripper motion to be perpendicular to the gripper jaws, we obtain a one-dimensional space of actions to search. By using a frictionless gripper (see Figure 2) we eliminate "wedging" of the object, leaving a finite set of stable object orientations. By assuming simultaneous contact by the jaws (pure squeezing) we eliminate the difficult analysis of pushing motions, and obtain predictions that are independent of the support friction.

We next specify the six-tuple $< \Theta, f_1, \mathcal{A}, G, C, \theta^* >$ that defines an instance of a parallel-jaw grasping problem.

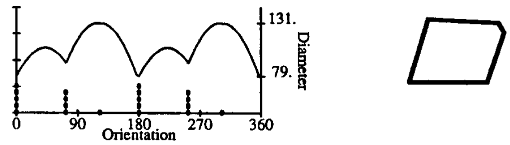

Figure 3: The diameter function for the five-sided object shown at the right in its zero orientation. During a squeeze, the object rotates so as to reduce the diameter, terminating when the diameter reaches a local minimum. The probability distribution after the first squeezing step is shown as a dotted histogram.

4.2. The Transfer Function G

In this paper we assume that the jaws make simultaneous contact and neither jaw loses contact—pure squeezing. The effect of a squeeze is predicted using the *diameter function*, which gives the minimum jaw separation as a function of the orientation of the object relative to the jaws. For polygonal objects the diameter function is piecewise sinusoidal (see Figure 4.1). During a squeeze, the object passively rotates to reduce the diameter, terminating when the diameter reaches a local minimum. If we assume a linear spring force between the jaws, the square of the diameter defines a potential function where local minima are "stable" in the sense that small deviations produce restoring forces. The diameter function has a period of π, so that it is impossible to remove a 180 degree ambiguity in object orientation through squeezing alone. See Appendix A for details on the diameter function.

A different ambiguity arises when the object's orientation is close to a local maximum in the diameter function. With a frictional gripper, the object would be wedged. Even with a frictionless gripper, an unstable equilibrium exists at the local maximum. Assuming that the prior probability density of orientations is continuous, the initial squeeze will encounter an unstable equilibrium with probability zero. Thereafter we avoid ambiguous actions.

Note that the analysis of grasp mechanics using the diameter function is not limited to polygonal objects. Any two-dimensional object can be analyzed if we can compute its diameter function.

4.3. The State Space Θ and Prior Probability f_1

The space of object orientations is the uncountable set of all planar angles, i.e. the half-open interval from zero to 2π. After the first squeeze step the object rotates into one of its stable orientations (where at least one edge of the object is aligned with the gripper, see above). The state space for the planning problem is this finite set of object orientations.

We consider the initial orientation of the object to be a random variable on the space of rotations. After the first squeezing action, the probability density for the initial orientation yields a discrete probability distribution on the set of stable orientations based on the integral between local maxima in the diameter function. In the absence of information that favors any subset of initial object orientations, we might assume that the initial orientation of the object is uniformly distributed, $f_1(\theta) = \frac{1}{2\pi}$. This is equivalent to assuming a random gripper orientation for the first squeezing step.

4.4. The Action Space, $\mathcal{A}$

A squeezing action orients the open gripper around the object and then closes the jaws until further motion would deform the object. The action space is the set of all gripper orientations, i.e. the half-open interval of angles from zero to 2π. Fortunately, this continuum can be partitioned into a finite number of equivalent intervals. For an n-sided polygonal object, there are at most n^2 unique actions. Suppose we have just squeezed the object. There are at most $2n$ different possible orientations, corresponding to the local minima of the diameter function. Now, consider the effect of a second squeeze, and how it varies with hand rotation. All that matters is which local minima fall into which intervals between local maxima. As the hand rotates, the effect of a squeeze changes only when a local minimum crosses a local maximum. There are at most n^2 of these changes. Each step in the squeezing plan is an action from this set.

4.5. Cost Function C and Desired Final State θ^*

We simply assume that each squeezing step takes one time unit as does the verification step.

Once we determine the object's orientation relative to the gripper, we can always achieve a desired final orientation by rotating the gripper. So in this case we choose the desired final state for each plan to be one of the states that has maximal probability.

5. Finding a Bayes' Plan

We use breadth-first search to look for an optimal plan. The space of plans defines a tree where each node is a probability distribution over the state space (also called a *hyperstate* (Astrom, 1987)). The root is the probability distribution resulting from the first (random) squeeze. The second level of nodes corresponds to the outcome of the two-step strategies. We stop searching when we encounter a node that assigns probability .5 to any orientation, since this is the best we can do. If, for example, such a node occurred at depth 3, its expected cost would be $(3+1)/.5 = 8$ time units, where the additional time unit is for the verification step. We would then examine each of the nodes at depth 2 and find the node that assigns maximal probability to any orientation. The path leading to this node corresponds to the best two-step plan. For any plan, the verification step is used to verify that the object is in one of the maximally-probable orientations.

We tested the planning algorithm on several familiar shapes as well as on 2000 randomly generated polygons (We generate 10 random points in the unit square and find their convex hull. The average number of sides was 5.9.) For each polygon we terminated the breadth-first after generating 1000 nodes. Only 1% of the polygons in the test set required this many nodes. In these cases we ran a second search using a best-first heuristic that hillclimbed on the node with the maximal probability. In each case the best-first search found a strategy using fewer than 236 nodes. Shorter, less certain plans gave lower expected time than the most certain plans in about one third of the cases (737/2000). Table 1 shows the results for some familiar shapes.

object shape	# sides	1-step	2-step	3-step	4-step
triangle	3	10.8	8.7	8.0	-
rectangle	4	6.3	6.0	-	-
ic-chip *	5	5.0	6.0	-	-
house *	5	8.0	9.2	9.4	10.0
example (Fig.1) *	5	7.4	6.2	8.0	-
allen-wrench *	6	9.1	7.4	8.0	-
tuning fork *	6	4.9	6.0	-	-
key *	11	9.1	6.2	8.2	-
hand	13	8.2	8.1	8.0	-

Table 1: Expected time for squeezing plans. Values are in time units where each squeezing step takes one unit. The last entry to the right in each row corresponds to the plan that minimizes orientational uncertainty. A (*) indicates objects where shorter, less certain plans have higher throughput.

6. Discussion

6.1. How Good is the Probability Model?

Of course the analysis depends on our model of probability. The need for probabilistic assumptions is at the core of a centuries-old controversy between Bayesian and frequentist statisticians (Berger, 1985). There is a similar issue in the average-case vs. worst-case analysis of algorithms (Karp, 1986). We must be careful to avoid claiming quantitative results based on ad-hoc assumptions. In this paper we address a specific problem in robot grasping and use geometry as a basis for our model of actions. The assumption of uniform prior distribution of orientations can be justified by *injecting* a random twist of the gripper prior to the first squeezing operation. This is analogous to the random pivot selection used to justify the average-case analysis of the Quicksort algorithm (Knuth, 1973, volume 3). Of course, in some applications objects will have biased orientations. We can take this into account in the prior distribution. If necessary the prior distribution can be obtained empirically.

6.2. Sensing/Acting

The goal of planning in the presence of uncertainty is to constrain the final state. This can be accomplished either by sensing, acting (Erdmann and Mason, 1986), or a combination of both. A sequence of squeezing actions is a sensorless plan where state is constrained by actions alone.

Probability theory offers a unified framework for comparing the effect of sensing with the effect of actions. Both sensing and actions have the effect of transforming a probability distribution over the state space. Perhaps we can quantify the contributions of sensing and acting in terms of their effect on the probability distribution, for example with an "information" measure such as entropy.

Our Bayesian framework can be extended to incorporate sensing, such as a simple sensor that returns jaw diameter, as in Taylor *et al.* (1987). The resulting plans would have branches, with different sensor values determining the actual path taken. The prior probability distribution can be used to estimate the likelihood of taking each path, so that we can compute an expected plan length. Also, we can borrow from the theory of Bayesian estimation to cope with noisy sensors.

If sensors are used then we can also use this framework to treat the problem of finding an optimal strategy for disambiguating among multiple objects, a form of recognition. In this case we add another dimension to the state space as in Donald's extended configuration space (Donald, 1987). The probability distribution at each node is then defined over this new state space and our cost function can reflect a desire to identify which object we are gripping from among a set of objects.

6.3. Computational Complexity

Statistical optimization often requires substantial computation. An exhaustive search with branching factor $O(n^2)$ is clearly intractable, although for the polygons we tested with fewer than 10 sides, the planner ran in a matter of seconds. A result by Natarajan (1989) can be used to show that a squeezing strategy can be found in polynomial time. We are working on a backchaining method that finds a strategy in time $O(n^3)$. For a manufacturing setting, we can perform the optimization off-line and amortize computation time over hundreds of execution cycles.

6.4. Other Extensions

We could extend the Bayesian framework to handle noisy actions. In this paper we assume that both jaws make contact simultaneously so that the the transfer function is deterministic. However when there is uncertainty in the control (how the command is physically carried out) or in the mechanics (how the object reacts) the transfer function is no longer deterministic. We can treat this form of uncertainty by modelling the transfer function with a conditional probability distribution (Goldberg, 1989). In the grasping application we could relax the assumption of pure squeezing by modelling the pushing phase, either with a worst-case analysis as in (Mason, 1986a) and (Brost, 1988) or using a model of conditional probability.

We could also extend the Bayesian framework to include a utility or payoff function on the state space. This would allow us to rank outcomes rather than making a sharp distinction between the desired final state and all others. For example, we could include a measure of grasp quality in our expected performance criterion.

A The Diameter Function

Let a two-dimensional object be described with a continuous curve in the plane, C. The distance between two parallel lines of support varies with the orientation of the lines. We define the diameter function, $d(\theta)$, to be the distance between parallel lines of support at angle θ. The maximum of this function is known as the *diameter* of the set of points in C (Preparata and Shamos, 1985).

- The diameter function is continuous: $\Delta d \to 0$ as $\Delta\theta \to 0$.
- The diameter function for C is equal to the diameter function for the convex hull of C.
- The diameter function has period π.

For an n-sided convex polygon, the diameter function can be represented as a list of piecewise sinusoids. Transitions between sinusoids can only occur when an edge is aligned with θ, so there are at most $2n$ transition angles.

Preparata and Shamos (1985) describe a linear-time algorithm for computing the diameter of a convex polygon with n sides; it proceeds by enumerating the set of all pairs of vertices that admit parallel supporting lines. There are at most $3n$ such pairs. Each pair defines a *chord* of length l_i and angle θ_i. The longest chord gives the diameter of the polygon.

To find the diameter function, sort this list of chords by increasing angle θ_i. If two chords have the same angle, discard the shorter chord. Also sort the list of polygon edges by angle and discard duplicates (corresponding to parallel edges). Every adjacent pair of edge angles in this sorted list corresponds to a sinusoid in the diameter function, $d_i(\theta) = l_i|\sin(\theta_i - \theta)|$, where l_i and θ_i are taken from the longest chord in the interval orthogonal to the interval between edges. Finding the longest chord in each interval requires a single sweep through the sorted list of chords.

Finding the convex hull and sorting dominates the running time, so we can compute the diameter function in time $O(n \log n)$. A simple reduction from SET DISJOINTNESS can be used to show that this running time is optimal.

Acknowledgements

This work was supported by the National Science Foundation under grant DMC-8520475.

We thank Randy Brost, John Canny, Lin Chase, Bruce Donald, Mike Erdmann, and Russ Taylor for helpful discussions. We also thank members of the manipulation lab: Dina Berkowitz, Alan Christiansen, Brian Welcker, and Yu Wang. We also thank Merrick Furst for a discussion that germinated the gripper design in Figure 2.

References

Allen, P. K., 1987. *Robotic Object Recognition Using Vision and Touch.* Kluwer Academic Publishers.

Astrom, K. J., 1987. Adaptive feedback control. *Proceedings of the IEEE, 75(2).*

Bajcsy, R., Lederman, S., and Klatzky, R., 1989. *Machine Systems for Exploration and Manipulation: A Conceptual Framework and Method of Evaluation.* GRASP Lab Tech Report MS-CIS-89-03, University of Pennsylvania.

Berger, J. O., 1985. *Statistical Decision Theory and Bayesian Analysis.* Springer-Verlag.

Bertsekas, D. P., 1987. *Dynamic Programming: Deterministic and Stochastic Models.* Prentice-Hall.

Boothroyd, G., Redford, A. H., Poli, C., and Murch, L. E., 1972. Statistical distributions of natural resting aspects of parts for automatic handling. *SME Manufacturing Engineering Transactions, 1:72.*

Brooks, R. A., 1982. Symbolic error analysis and robot planning. *International Journal of Robotics Research, 1(4),* 29–68.

Brost, R. C., 1988. Automatic grasp planning in the presence of uncertainty. *The International Journal of Robotics Research,* . Also appeared in Proceedings of the IEEE International Conference on Robotics and Automation, San Francisco, CA, April, 1986.

Buckley, S. J., 1987. *Planning and Teaching Compliant Motion Strategies.* PhD thesis, MIT.

Cameron, A. and Durrant-Whyte, H., 1989. *A Bayesian Approach to Optimal Sensor Placement.* Robotics Research Report OUEL 1759/89, University of Oxford, U.K.

Cheeseman, P., 1985. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence,* IJCAI, Los Angeles, CA.

DeGroot, M. H., 1970. *Optimal Statistical Decisions.* McGraw-Hill.

Donald, B. R., 1987. *Error Detection and Recovery in Robotics.* Springer-Verlag.

Durrant-Whyte, H., 1988. *Integration, Coordination, and Control of Multi-Sensor Systems.* Kluwer.

Erdmann, M. A., 1984. *On Motion Planning with Uncertainty.* Master's thesis, MIT.

Erdmann, M. A., 1989. *On Probabilistic Robot Strategies.* PhD thesis, MIT.

Erdmann, M. A. and Mason, M. T., 1986. An exploration of sensorless manipulation. In *IEEE International Conference on Robotics and Automation.*

Genesereth, M. and Nilsson, N., 1987. *Logical Foundations of AI.* Morgan-Kaufmann.

Goldberg, K. Y., 1989. Probabilistic grasping. Thesis Proposal, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA. 15213.

Hager, G. D., 1988. *Active Reduction of Uncertainty in Multi-Sensor Systems.* PhD thesis, University of Pennsylvania CIS.

Hager, G. D. and Mintz, M., 1989. Sensor modeling and robust sensor data fusion. *Fifth International Symposium on Robotics Research,* .

Hanafusa, H. and Asada, H., 1977. Stable prehension by a robot hand with elastic fingers. *Trans. of Soc. of Inst. and Control Engineers, 13(4).*

Hutchinson, S. A. and Kak, A. C., 1989. Planning sensing strategies in a robot work cell with multi-sensor capabilities. *IEEE Transactions on Robotics and Automation, 5(6).*

Karp, R. M., 1986. 1985 turing award lecture: combinatorics, complexity, and randomness. *CACM,* , 97–117.

Knuth, D. E., 1973. *The Art of Computer Programming.* Addison-Wesley.

Koutsou, A., 1988. *Object Exploration Using a Parallel-Jaw Gripper.* GRASP Lab Tech Report MS-CIS-88-48, University of Pennsylvania.

Latombe, J., 1989. *Motion Planning with Uncertainty: On the Preimage Backchaining Approach,* pages 53–70. MIT Press. edited by O. Khatib, J. J. Craig, and T. Lozano-Perez.

Lozano-Perez, T., Mason, M. T., and Taylor, R. H., 1984. Automatic synthesis of fine-motion strategies for robots. *International Journal of Robotics Research, 3(1),* 3–24.

Lumelsky, V. J., 1987. Algorithmic and complexity issues of robot motion in an uncertain environment. *Journal of Complexity, 3,* 146–182.

Mason, M. T., 1984. Automatic planning of fine motions: correctness and completeness. In *proceedings, International Conference on Robotics,* IEEE Computer Society, Atlanta, GA.

Mason, M. T., 1986a. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research, 5(3).*

Mason, M. T., 1986b. On the scope of quasi-static pushing. In Faugeras, O. and Giralt, G., editors, *The Third International Symposium on Robotics Research,* MIT Press.

Mason, M. T., Goldberg, K. Y., and Taylor, R. H., 1988. *Planning Sequences of Squeeze-Grasps to Orient and Grasp Polygonal Objects.* Technical Report CMU-CS-88-127, Carnegie Mellon University, Computer Science Dept., Schenley Park, Pittsburgh, PA 15213.

Natarajan, B. K., 1989. Some paradigms for the automated design of parts feeders. *International Journal of Robotics Research, 8(6),* 98–109. Also appeared in Proceedings of the 27th Annual Symposium on Foundations of Computer Science, 1986.

Pertin-Troccaz, J., 1989. *Grasping: A State of the Art,* pages 71–98. MIT Press. edited by O. Khatib, J. J. Craig, and T. Lozano-Perez.

Peshkin, M. A., 1986. *Planning Robotic Manipulation Strategies for Sliding Objects.* PhD thesis, CMU Dept. of Physics.

Peshkin, M. A. and Sanderson, A. C., 1988. Planning robotic manipulation strategies for workpieces that slide. *IEEE Journal of Robotics and Automation, 4(5).*

Preparata, F. P. and Shamos, M. I., 1985. *Computational Geometry: An Introduction.* Springer-Verlag.

Requicha, A. A. G., 1983. Toward a theory of geometric tolerancing. *IJRR, (2(4)).*

Smith, R. C. and Cheeseman, P., 1986. On the representation and estimation of spatial uncertainty. *IJRR, 5(4),* 56–68.

Stansfield, S. A., 1987. *Visually Guided Haptic Object Recognition.* PhD thesis, Univ. Of Pennsylvania Dept. of CIS.

Stengel, R. T., 1986. *Stochastic Optimal Control.* John Wiley.

Taylor, R. H., 1976. *A Synthesis of Manipulator Control Programs from Task-Level Specification.* PhD thesis, Stanford AIM-282.

Taylor, R. H., Mason, M. M., and Goldberg, K. Y., 1987. Sensor-based manipulation planning as a game with nature. In *Fourth International Symposium on Robotics Research.*