

Generalizing the Analysis of PRM

Andrew Ladd

Lydia E. Kavraki

Department of Computer Science

Rice University

Houston, TX, 77005

aladd|kavraki@cs.rice.edu

Abstract

This paper presents a novel analysis of the probabilistic roadmap method (PRM) for path planning. We formulate the problem in terms of computing the transitive closure of a relation over a probability space and give a bound in terms of the number of intermediate points for some path and the probability of choosing a point from a certain set. Explicit geometric assumptions are not necessary to complete this analysis and consequently it provides some unification of previous work as well as generalizing to new path planning problems, two of which, 2k-dof kinodynamic point robots and deformable robots with force field control are presented in this paper.

1 Introduction

Planning a collision-free path for a rigid or articulated robot to move from an initial to a final configuration in a static environment is a central problem in robotics and has been the topic of extensive research over the last decade [1, 8, 16]. The complexity of the problem is high and several versions of it have been shown PSPACE-hard [16]. Interesting applications and extensions of the problem exist in planning for robots that can modify their environments [9, 18] and flexible robots [3], planning for graphics and simulation [15], planning for virtual prototyping [6], and planning for medical [20] and pharmaceutical [7] applications.

This paper concentrates on the analysis of PRM [14, 11]. Since 1994, when PRM was invented, several researchers have reported on the excellent performance of the planner for robots with many degrees of freedom, several variations of the method have been developed (e.g., [2]), several planners that bare resemblances with PRM have been introduced (e.g., [10]), and several extensions of the basic path planning problem have been solved with PRM-based methods (e.g., [18]).

The experimental success of the planner has motivated many researchers to seek a theoretical basis for explaining its performance and relative successes in this direction have been reported, among others, in [5, 13, 12, 23, 10, 21]. This paper presents a further extension in this direction by using the mechanism of measure theory [4].

1.1 Previous Work

The techniques discussed in this section can be roughly classified as one of two types: isolation of a single path and space covering arguments. This paper is interested chiefly in the former type as it is the approach we take.

Path isolation method

PRM is a randomized algorithm for constructing paths by concatenating simple pieces of path together. The paths are computed by a local planner which is not complete and fails in some cases. If the local planner is sufficiently powerful and some path exists, the path can be eventually guessed. PRM turns out to be surprisingly effective in practice and seems to be exploiting the property that there are many paths between two configurations in most robotics applications.

A common technique for analyzing PRM is to consider motion planning for point robots in open subspaces of \mathbf{R}^k . An example of this kind of this analysis can be found in [12]. A single path is isolated and then analyzed by tiling with simple shapes or *buckets*. [5], [12], [22] and [14] use this approach. The simple shapes used in these papers are ϵ -balls where $\epsilon > 0$ is related to the *path clearance*. The point choosing function is assumed to have distributions proportional to the volume of the balls. A bound in terms of path clearance and the measure of an open ϵ -ball is obtained. The probability of failure is shown to decrease exponentially with the number of guesses.

An extension to this technique can be made for small

time locally controllable robots [24] such as car-like robots and tractor trailer robots. The property exploited is that for every $\epsilon > 0$, there exists $\delta > 0$ such that for any point within δ distance, it can be reached by taking a path within ϵ . A path with ϵ clearance can thus be tiled with δ -balls. Again the probability of failure was shown to decrease exponentially. A proof for car-like robots that cannot drive backwards was also achieved in [23].

Space covering method

In [13], a notion of ϵ -goodness was described in attempt to obtain the rate of total coverage of the free space. A space is ϵ -good if every point in the free space can ‘see’ more than an ϵ fraction of the space with the local planner. The spaces discussed were simply connected compact sets with measure 1. The notion of *milestone* was introduced to describe the intermediate points in the computed paths. This paper showed how to bound the number of milestones needed to achieve some probability of success. [25] showed an example of $\frac{2}{9}$ -good space which required an infinite number of milestones to capture the entire space. Some evidence that the disconnected fraction of an ϵ -good space will decrease asymptotically to 0 was given.

Hsu explored an extension to ϵ -goodness with his work on expansive spaces [9]. This work further formalized the notions of reachability and made use of measure in some sense. Concisely, for a connected set of points S , the β -LOOKOUT(S) is the subset of S whose points ‘see’ using the local planner more than a β fraction of the set of points which can be ‘seen’ from S . A space is (α, β) -EXPANSIVE if the β -LOOKOUT(S) is always larger than an α fraction of the measure of S for every connected subset S of the points reachable from any point in free space. Again this work provides a bound on the number of milestones required to generate a path. This work was useful to shed insight on the probabilistic completeness of kinodynamic planners [10].

1.2 Key Concepts and Motivation

The key difference in the treatment in our paper is the abstract reformulation of typical path planning problems to isolate the essential properties which allow us to analyze PRM. We capture the properties of a path planning task for a single path in terms of two simple constants. We relate these constants to previous analyses to suggest ways that they can be bounded in practice. Certain notions from an analysis such as [12] will be replaced with more general ideas.

The simple geometric shapes, e.g., ϵ -balls of [12], tiling a path can be replaced with sets of strictly positive measure. They are not necessarily connected, open or even infinite. The probability distribution of the point choosing function can be replaced with a computable probability measure. The configuration spaces can be replaced with the more vague notion of state space without any explicit geometric assumptions. The predicate of reachability and the local planner are formalized with binary relations with the local planner being a subset of the reachability relation. The reachability is assumed to be transitive and both are assumed to have certain measurability properties. Our reformulation of the PRM advances the state-of-the-art PRM analysis by working with spaces with weaker mathematical structure and by using geometry less explicitly. Although for some systems, an analysis using particular geometric facts might yield better bounds than what we give, we are not aware of treatment to date which has achieved this. Our approach unifies some of the existing work and gives a framework for approaching new path planning problems.

2 Problem Formulation

The operation of the PRM algorithm can be summarized as follows.

Algorithm 2.1. PRM planner.

1. Generate N points (configurations) from the state space (C -space) at random.
2. Use the local planner to build a directed graph with the points as vertices and the edges meaning local path reachability.
3. Given a query (x, y) , connect x and y to the graph and use graph search to find a path.
4. If a path is found return it otherwise return FAILURE.

We seek to connect the value of N required to guarantee sufficiently low probability of error in PRM with the interaction between the local planner and the state space.

Take a set X to be the free state space for a robot. By analogy, this is taken to mean the entire set of distinct and allowable states the robot can assume. For example, this might be the C -space with some extra information as to the state of the robot which is relevant to the planning problem, e.g., time, memory contents,

velocity or the amount of fuel left. The path reachability relation is transitive, i.e., if x reaches y and y reaches z then x reaches z . This is a natural assumption which expresses what is intuitively understood by state space and path reachability. Note that symmetry and reflexivity are not enforced. If X encodes time, for example, path relations would be necessarily asymmetric.

The local planner can also be thought of as a binary relation over X . This relation, which we will call R , is not necessarily transitive. This method has a hope of success if any valid path can be broken down into a finite sequence of states x_1, \dots, x_n such that $x_1 R \dots R x_n$. If x_2, \dots, x_{n-1} are present in the roadmap then a query of “does x_1 reach x_n ?” will be answered correctly. PRM can also return this sequence, from which the path can be reconstructed and executed [14].

PRM might be successful in the case where the transitive closure of our local planner R , denoted \bar{R} , is the path reachability relation. Cases where the closure and the path relation do not agree will be discarded - the algorithm fails in these cases. This leads to an abstract rephrasing of a general path planning problem with PRM.

3 Notation

As stated earlier, the set X will be the set of distinct and valid states the robot can assume. The set $\Sigma \subset 2^X$ is a σ -algebra for X . For example, a natural choice for this would be Borel algebra in the case where X has a topology [19]. The function $\mu : \Sigma \rightarrow [0, 1]$ is a probability measure on (X, Σ) . μ is taken to represent the distribution of the random sample function on X . If α is the random variable indicating a point chosen from X at random by the sampler and A is a measurable subset of X , $P(\alpha \in A) = \mu(A)$. In short, (X, Σ, μ) is a probability space.

For this discussion it will be necessary to extend the measure on X by finite dimensions in order to consider n -ary reachability relations over the space. The operator \otimes denotes generation of σ -algebras via the usual construction over rectangles (elements of Cartesian products of σ -algebras). The canonical σ -algebra and probability over X^n will be $\Sigma^n = \Sigma \otimes \dots \otimes \Sigma$ and $\mu_n : \Sigma^n \rightarrow [0, 1]$. μ_n is uniquely defined by its action on rectangles, i.e., $\mu_n(A_1 \times \dots \times A_n) \mapsto \prod \mu(A_i)$.

The local planner is described by a relation, R , over the set X . This relation will have the additional restriction that it is measurable, in other words $R \in \Sigma^2$.

This a natural assumption which will not inhibit the study of ‘reasonable’ planning problems. The notation of this object is given by the identity $xRy \Leftrightarrow (x, y) \in R \in \Sigma^2$. x reaches y is meant by xRy .

Another representation for R is as the characteristic function for the set R , which is more convenient for our purposes:

$$\chi_R(x, y) := \begin{cases} 0 & \text{for } (x, y) \notin R, \\ 1 & \text{for } (x, y) \in R. \end{cases}$$

The preimage of the above function is R , i.e., $R = \chi_R^{-1}(1)$. This function is measurable and can be easily extended to n -ary analogues as follows. The family of functions $f_n : X^n \rightarrow \{0, 1\}$ is given by

$$f_n(x_1, \dots, x_n) \mapsto \prod_{i=1}^{n-1} \chi_R(x_i, x_{i+1}).$$

In other words, $f_n(x_1, \dots, x_n) = 1$ iff $x_1 R \dots R x_n$. Another useful version of this function will be written f_n^{xy} and defined as $f_n^{xy}(x_1, \dots, x_n) := f_{n+2}(x, x_1, \dots, x_n, y)$. It follows directly that these functions are well-defined and measurable.

A second kind of closure can be formulated in a similar way. Intuitively, we want any membership in \bar{R} to be determinable by a random point guessing process¹. The binary relation \hat{R} will be taken to mean $x\hat{R}y$ when a random point guessing process has positive probability of proving that $x\bar{R}y$. Here is a formal construction of \hat{R} . Let

$$F_n(x, y) := \int_{X^n} f_n^{xy}(\cdot) d\mu_n.$$

Note that F_n is a measurable function from X^2 to $[0, 1]$. Now \hat{R} can be written

$$\hat{R} := \bigcup_{n=0}^{\infty} F_n^{-1}((0, 1]).$$

By their constructions, $R \subseteq \hat{R} \subseteq \bar{R}$ and each relation is an element of Σ^2 .

4 A Bound for N

The main result presented in this paper is a bound on the expected number of points needed to be generated in order to determine membership in \bar{R} . This bound

¹Any algorithm which guesses points with μ and tries to construct paths from input x and y with R .

holds if and only if $\widehat{R} = \bar{R}$. We know if this assumption does not hold, then random point guessing processes will fail on some queries. The method of proof will be to reduce the problem of finding a path between two particular points to a standard problem in discrete probability; the following theorem will be used.

Theorem 4.1 (Coupon Collector[17]). *To win a prize in a contest held by a breakfast cereal company, it is necessary to obtain at least one of each of n coupons in the boxes. The coupons are placed in the boxes according the uniform distribution, one per box. The expected number of boxes one must buy to get the prize is*

$$E(N) = n \log n.$$

For the problem considered in this paper, the existence of certain buckets (cereal boxes) of strictly positive measure will be shown. These buckets will be such that guessing at least one point from each ensures that PRM has computed a path.

Theorem 4.2. *If it is possible to randomly guess a path with positive probability, the expected number of random points needed to be generated for PRM to successfully compute $x\bar{R}y$ from R , $E(N)$, satisfies the following inequality*

$$E(N) \leq \frac{\log n}{p},$$

for constants n and p .

Proof. It is possible to randomly guess a path with positive probability if and only if $\widehat{R} = \bar{R}$.

Suppose $x\bar{R}y$ for some $x, y \in X$. Since $\widehat{R} = \bar{R}$, there is n such that $F_n(x, y) > 0$, using the earlier definition of F_n . In other words, $A = (f_n^{xy})^{-1}(1)$ is such that $\mu_n(A) > 0$. It follows that there is a rectangle $A_1 \times \cdots \times A_n \subseteq A$ such that $\mu(A_i) > 0$ for each i . Notice that for any sequence x_1, \dots, x_n with $x_i \in A_i$, $xRx_1R \cdots Rx_nRy$ holds.

PRM will certainly succeed if a point from each A_i can be guessed. Since they each have positive measure this will eventually happen, however we would like to obtain a bound as well.

Each A_i can be thought of as a Coupon Collector bucket. For an illustration look at Figure 1. We will ignore points that land outside of the distinguished $A_1 \times \cdots \times A_n$ in order to obtain an overestimation of $E(N)$. Also, we will assume that the A_i are disjoint - otherwise if a point is guessed which is in multiple buckets it can be randomly reassigned to a single bucket to obtain an overcount of $E(N)$.




Figure 1: A free path in state space and an illustration of the coupon buckets

Let $p = \min \mu(A_i)$ and conservatively take all the buckets to have measure p . Again this produces an overcount and we conclude that with probability at least np a point in at least one bucket is guessed.

We will define a sequence of random variables Y_i counting the number of guesses before the i th point in some bucket is guessed. The random variable T will count the number of guesses in any bucket required to obtain at least one guess in each bucket. It follows that $E(N) \leq E\left(\sum_{i=1}^T Y_i\right)$. Note that the Y_i are independent and identically distributed. Furthermore, the Y_i are independent of T . These observations allow us to conclude that $E\left(\sum_{i=1}^T Y_i\right) = E(T) \cdot E(Y_1)$.

By Coupon Collector we know that $E(T) = n \log n$ (since there are n buckets) and $E(Y_1) = 1/np$, the inverse of the probability of landing in some bucket. The final inequality follows. \square

Corollary 4.3. *After guessing N points, we can write the probability of not having guessed the path, P , as*

$$P \leq n(1-p)^N.$$

This can be seen by applying the union bound and it clearly decreases exponentially when $p > 0$.

The limiting part of bound of 4.2 is p . A useful additional result would be to show that for some particular path planning that p could be lower bounded in terms of the state space or some attribute of a given path. Our analysis allows for the use of sets which may not have simple geometric shapes thus extending the palette of sets of positive measure that we could use to find a lower bound for p .

5 $2k$ -dof Kinodynamic Robots

In this section, we begin by summarizing the known single path analysis of a k -dof holonomic robot and then discuss issues surrounding the extension to $2k$ -dof kinodynamic robots.

The workspace for this example will be a k -manifold $W \subset [0, 1]^k$. The state space is the workspace, $X = W$. The random sample function has the distribution induced by the Borel measure on $[0, 1]^k$ normalized to be a probability measure.

A fully holonomic robot operating in this workspace and the local planner connects points with a straight line. This analysis closely parallels [12] and [5].

The buckets (as in Th. 4.2) for k -dof point robots can be constructed and used to get more explicit bounds on $E(N)$. Let $\mu(B_\delta(\cdot))$ be the measure of an open δ -ball in $[0, 1]^k$. Suppose we have some path with ϵ path clearance and length L . Then

$$E(N) \leq \frac{(\log L - \log \epsilon)\mu(X)}{\mu(B_{\epsilon/2}(\cdot))}.$$

Estimating $E(N)$ is made more difficult by considering velocities.

The workspace for the extension is a k -manifold $W \subset [0, 1]^k$ and the state space of the robot is $X = W \times (-1, 1)^k$, which encode position and velocity. The robot can be controlled by applying a constant acceleration in the range $(-1, 1)$ for a constant non-zero time period. In 1-d, the local planner that we use tries to connect position (x_1, v_1) with (x_2, v_2) by taking

$$t = \frac{2(x_2 - x_1)}{v_1 + v_2}, \quad a = \frac{v_2 - v_1}{t}.$$

In k -d, we solve each dimension independently and cases where either acceleration is too large, singularities arise or time is non-positive are not solutions. We refer to this local planner as R in the following.

We aim to show that PRM succeeds but that ball tiling arguments as in [12] are inappropriate for this problem. The shapes of the buckets we need are disconnected and dependent on the input points.

The claim we will establish is that the solution space around a given xRy has positive measure. Together with closure of open sets under intersection, we show this fact is enough to conclude that $\hat{R} = \bar{R}$.

Let $Y \subset \mathbf{R}^4 = \{(x, v, y, w) : (x, v)R(y, w)\}$. For every $\bar{x} = (x, v, y, w) \in Y$, $|x - y| > 0$ and $|v \pm w| > 0$, so $\exists \epsilon > 0$ such that $B_\epsilon(\bar{x}) \subset Y$. It follows that Y is 4-manifold.



Figure 2: solution space for a single bucket path

The point guessing distribution (μ) that we use has positive measure on open sets. Suppose $(x_1, v_1)R(x_2, v_2)R(x_3, v_3)$. There exists $\epsilon > 0$ such that $(x'_2, v'_2) \in B_\epsilon(x_2, v_2)$ is such that $(x_1, v_1)R(x'_2, v'_2)R(x_3, v_3)$. Since $\mu(B_\epsilon((x_2, v_2))) > 0$, we have a probabilistically complete path planner and Th.4.2 applies.

In this example, n and p depend on the input points independently of the obstacles. Given a C^1 path between two points, a path made of piecewise constant second derivatives which arbitrarily well approximates the first path can be found. In Figure 2, we can see a possible solution space for a planning problem with a single intermediate milestone in 1-d, the vertical axis being velocity and horizontal axis being position. The shape depends heavily on the start and finish points and is disconnected. The measure of the set, however, is a significant fraction of the measure of smallest disc which encloses all of the points.

6 Deformable Robots

In this section, we consider motion planning with deformable robots controlled by force fields. This section will sketch how to show probabilistic completeness of the path planner. There will be little emphasis on the control and simulation of parametric deformables, an interesting topic on its own.

The robot we consider in this section will be a deformable operating in a k -dimensional compact workspace. It will be controlled by an external force field.

For the sake of simplicity, suppose the configuration space is the set of all C^2 curves embedded into a k -manifold workspace $W \subset [0, 1]^k$ which satisfy some constraints on total deformation energy and local strain energy. The state space X is W together with a C^1 velocity field on the curve. The robot can be controlled by applying C^0 force fields to the curve. We can subdivide the curve recursively (say in two pieces). This will form a lattice \mathcal{L} of subdivision

topologies. For each $\lambda \in \mathcal{L}$, we have a state space X_λ which is a m -manifold for some m which represents the curve’s constrained deformation, embedding and velocity field in terms of a finite parameter set (where the curve is obtained by interpolation). To each $\lambda \in \mathcal{L}$, we assign a probability $p_\lambda > 0$ such that $\sum_{\lambda \in \mathcal{L}} p_\lambda = 1$. Also, an operator \vee on every pair $\lambda, \lambda' \in \mathcal{L}$ can be defined so that $\lambda \vee \lambda'$ is the simplest common subdivision topology.

For two states $x, y \in X$, suppose there is a path between them. We will now sketch an approximation scheme for the path, discuss what kind of properties the local planner must have and show how we can compute the path with PRM. More specifically, we construct X' with an associated measure and local planner R such that PRM succeeds and implies paths in X . We will now show a path planning result for a space which is *not* a manifold which, to our knowledge, has not been achieved to date.

We rely on several reasonable assumptions. The subdivision scheme we propose must be of the type where the curve represented by some subdivision topology and parameters must be the limit of the subdivision process. The family of curves and primitive paths must also be sufficiently rich to approximate any given curve arbitrarily well when taken under finite composition, i.e., \bar{R} is path reachability. Finally, we assume that queries are made with representable pairs (x, y) . Let R_λ be the local planner which connects points in X_λ . We must first show that X_λ with its probability measure μ_λ and with local planner R_λ is probabilistically complete. Recall that X_λ is an m -manifold. For any $z_1, z_2, z_3 \in X_\lambda$ such that $z_1 R_\lambda z_2 R_\lambda z_3$, we define Y as the set of points $z'_2 \in Y$ where $z_1 R_\lambda z'_2 R_\lambda z_3$. It is now sufficient to show Y is also an m -manifold. We could conclude that $\mu_\lambda(Y) > 0$, then it follows $p > 0$ (in the sense of Th. 4.2) for any path with a finite number of milestones.

We construct the state space $X' = \bigcup_{\lambda \in \mathcal{L}} X_\lambda$ with probability measure taken by the product σ -algebra and measure constructions. Measures are weighted for each λ by p_λ . The local planner for points z, z' (with subdivision λ and λ' respectively) works by reinterpreting z and z' as points in $X_{\lambda \vee \lambda'}$ and using its corresponding local planner. It is easy to show that this constructs a probability space. Furthermore, it follows that PRM is probabilistically complete on X' with local planner R .

Suppose $\gamma : [0, 1] \rightarrow X$ is the path between x and y and this path has ϵ clearance. Since the subdivisions can generate arbitrarily good approximations to

points in X , there exist points $x_1, \dots, x_n \in X'$ such that $x R x_1 R \dots R x_n R y$ and the new path is within ϵ of γ .

We have shown that, under reasonable assumptions, given a path for our deformable robot we can construct a path which is within $\epsilon > 0$ that can be found using PRM without fixing a parametrization a priori. This shows that a generic path planner could be constructed for this problem and that the probability of failure of the planner would tail exponentially with respect to the number of guesses. Since the approximation space we constructed is not a manifold, we also note that we succeeded in showing path planning results in non-manifold spaces without sacrificing the aspects of PRM that make it desirable to implement in practice. This kind of analysis was not possible with previous frameworks.

7 Summary of Results

We reformulated the robot path planning problem in terms of probability spaces, measures and the computation of the transitive closure of a given measure. We showed that if it was possible to guess a path at random then by using PRM the probability of failing to find an existing path would decrease exponentially. This bound was given in terms of two intuitive constants n and p . The expected number of guesses required to find a path was shown to be logarithmic in n and inversely proportional to p . We showed that PRM succeeds for $2k$ -dof point robots and gave strong arguments towards success for deformable robots controlled by force fields.

In the examples, we exhibited asymmetric, disconnected bucket constructions to how our analysis offers alternatives to simple geometric tiling shapes. We also described path planning results for non-manifold, non-parametric spaces. We believe that our treatment provides a framework for analysis of path planning with more physically realistic spaces where geometry, dynamics and computational efficiency concerns require state spaces which have multiple and redundant representations.

Acknowledgments

Andrew Ladd is partially supported by FCAR 70577 and by NSF IRI-970228. Work on this paper by Lydia Kavradi was supported by NSF CAREER Award IRI-970228, a Whitaker grant, an ATP Award and a Sloan Fellowship. Andrew would like to thank Gre-

gory Dudek for introducing him to this topic, robotics in general and for many discussions on randomized path planning.

References

- [1] P. K. Agarwal, L. E. Kavraki, and M. Mason. *Robotics: The Algorithmic Perspective*. A K Peters, Natick, MA, 1998.
- [2] N. Amato, B. Bayazit, L. Dale, C. Jones, and D. Vallejo. Obprm: An obstacle-based prm for 3d workspaces. In P. Agarwal, L. Kavraki, and M. Mason, editors, *Robotics: The Algorithmic Perspective*. AK Peters, 1998.
- [3] E. Anshelevich, S. Owens, F. Lamiroux, and L. Kavraki. Deformable volumes in path planning applications. In *IEEE Int. Conf. Robot. & Autom.*, pages 2290–2295, 2000.
- [4] P. Billingsley. *Probability and Measure*. Wiley-Interscience, 1995.
- [5] R. Bohlin. *Motion Planning for Industrial Robots*. Licentiate thesis, Chalmers Univ. of Technology, 1999.
- [6] H. Chang and T. Y. Li. Assembly maintainability study with motion planning. In *Proc. IEEE Int. Conf. on Rob. and Autom.*, pages 1012–1019, 1995.
- [7] P. W. Finn and L. E. Kavraki. Computational approaches to drug design. *Algorithmica*, 25:347–371, 1999.
- [8] K. Gupta and A. P. del Pobil. *Practical Motion Planning in Robotics*. John Wiley, West Sussex, UK, 1998.
- [9] D. Hsu. *Randomized Single-Query Motion Planning In Expansive Spaces*. PhD thesis, Department of Computer Science, Stanford University, 2000.
- [10] D. Hsu, R. Kindel, J. Latombe, and S. Rock. Randomized kinodynamic motion planning with moving obstacles. In *International Workshop on Algorithmic Foundations of Robotics (WAFR)*, 2000.
- [11] L. Kavraki. *Random Networks in Configuration Space for Fast Path Planning*. PhD thesis, Stanford University, 1995.
- [12] L. Kavraki, M. N. Kolountzakis, and J.-C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Trans. Robot. & Autom.*, 14(1):166–171, February 1998.
- [13] L. Kavraki, J. Latombe, R. Motwani, and P. Raghavan. Randomized query processing in robot motion planning. In *Proc. ACM Symp. on Theory of Computing*, pages 353–362, 1995.
- [14] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. & Autom.*, 12(4):566–580, June 1996.
- [15] Y. Koga, K. Kondo, J. Kuffner, and J. Latombe. Planning motions with intentions. *Computer Graphics (SIGGRAPH'94)*, pages 395–408, 1994.
- [16] J. Latombe. *Robot Motion Planning*. Kluwer, Boston, MA, 1991.
- [17] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge, 1996.
- [18] C. Nielsen and L. Kavraki. A two level fuzzy PRM for manipulation planning. In *IEEE/RSJ Int. Workshop on Intelligent Robots & Systems (IROS)*, Japan, 2000.
- [19] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, 1976.
- [20] A. Schweikard, R. Tombropoulos, L. Kavraki, J. Adler, and J.-C. Latombe. Treatment planning for a radiosurgical system with general kinematics. In *Proc. of the IEEE International Conference on Robotics and Automation*, pages 1720–1727, San Diego, CA, 1994.
- [21] S. Sekhavat, P. Svestka, J.-P. Laumond, and M. H. Overmars. Multilevel path planning for nonholonomic robots using semiholonomic subsystems. *Int. J. Robot. Res.*, 17:840–857, 1998.
- [22] P. Švestka. *Robot Motion Planning using Probabilistic Road Maps*. PhD thesis, Utrecht University, the Netherlands, 1997.
- [23] P. Švestka and M. Overmars. Motion planning for car-like robots using a probabilistic learning approach. *Int. J. of Robotics Research*, 16(2):119–143, 1997.
- [24] P. Svestka and M. H. Overmars. Coordinated motion planning for multiple car-like robots using probabilistic roadmaps. In *IEEE Int. Conf. Robot. & Autom.*, pages 1631–1636, 1995.
- [25] P. Valtr. On galleries with no bad points. Technical Report DIMACS Technical Report 96-55, DIMACS Center, Rutgers University, Piscataway, New Jersey, 1996.