

Information Theoretic Construction of Probabilistic Roadmaps

Brendan Burns Oliver Brock
 University of Massachusetts Amherst
 Amherst, Massachusetts 01003
 {bburns, oli}@cs.umass.edu

Abstract—Probabilistic Roadmaps (PRM) are a randomized tool for path planning in configuration spaces where exhaustive search is computationally intractable. It has been noted that the PRM algorithm’s computational cost can be greatly reduced by reducing the number of samples necessary to construct a successful roadmap. We examine the information theoretic properties of roadmap construction and propose sampling techniques based upon maximizing the information gain of the roadmap for each configuration sampled. Instead of sampling algorithms which are meant to understand the entirety of configuration space, our sampling is focused on finding configurations which facilitate roadmap construction. We show empirically that these approaches can lead to a significant reduction in the number of samples necessary to construct a useful roadmap.

I. INTRODUCTION

Path planning in high dimensional spaces is difficult. In some cases it has shown to be PSPACE hard [6]. Complete deterministic path planning is often computationally impractical. The central problem is that the size of the configuration space is exponential in the degrees of freedom of the robot. Probabilistic Roadmap (PRM) path planning techniques [5] were designed to use randomization to make configuration space path planning computationally tractable.

Since the collision checking and path planning associated with adding new samples requires most of the computation in PRM construction, reduction of the number of samples necessary for the construction of a complete roadmap is desirable. This leads to a decrease in overall computational costs.

Originally, the PRM technique used uniform random sampling to construct a roadmap in configuration space [5]. More recently there have been a number of alternative sampling strategies proposed as improvements to PRM algorithm. These can be generally characterized into deterministic approaches which guarantee a certain density of sampling [2], [3], and guided approaches which use heuristics to determine the areas of configuration space which should be examined [1], [4].

A. Basic PRM Path Planning

The basic PRM algorithm consists of construction, expansion and query steps. The construction step of the PRM algorithm operates as follows:

Given an empty roadmap, $R(V, E)$, for a specified number of samples do the following: Select an unobstructed point p in free space. Add p to the roadmap $R(V, E)$. Find the set N consisting of MaxNeighbors nearest neighbors to p . For each $n \in N$ if p and n are in different connected components in $R(V, E)$ and a straight-line path is possible between n and p , add an edge from n to p in the roadmap. After the specified number of samples have been selected, return the roadmap R for expansion.

The PRM expansion step attempts to connect disconnected components in $R(V, E)$. For a specified number of iterations, select v , the most “difficult” vertex in V . Difficulty is measured by some heuristic, usually the percentage of failed path planning attempts involving v . Next, select N , the set of v ’s nearest neighbors as above and attempt to connect v to all $n \in N$ where v and n are not in the same connected component. After a specified number of iterations, return $R(V, E)$ for path planning.

Given a roadmap $R(V, E)$ and two points p_1 and p_2 , the PRM algorithm performs path planning by first finding v_1 , the vertex in the roadmap closest to p_1 . Next, it selects v_2 equivalently for p_2 . Given v_1, v_2 and $R(V, E)$ PRM uses a standard path querying technique to find a path between v_1 and v_2 in $R(V, E)$.

B. Deterministic extensions to PRM Path Planning

It has been suggested that quasi-random deterministic approaches to sampling can improve the performance of the PRM algorithm [2], [3]. These approaches use quasi-random sequences to generate sets of configuration points which guarantee both good coverage and a bounds on the density of samples in any given region. Additionally they offer incremental refinement of the roadmap produced by the path planning algorithm.

However, the density of sampling required is determined by the most difficult area of the configuration space.

As a result, unless the configuration space is uniformly difficult deterministic quasi-random sampling will place unnecessary samples in the configuration space.

C. Guided Sampling

In contrast, guided sampling attempts to bias the choice of configurations towards those areas of configuration space which are thought to contain the configurations necessary for a complete roadmap. Guided sampling does not enforce uniform sampling over the entire configuration space. If the informed sampling is ideally designed (and this of course is the goal, however unattainable) then the roadmap will contain a minimum number of vertices. If an ideal sampling technique knew the location of these vertices, it would restrict its search to these points.

1) *Previous guided samplers:* Earlier work in information based sampling has focused upon finding areas near obstacles [1] and narrow passage-ways in configuration space [4]. These efforts are generally motivated by the knowledge that important areas in configuration space are necessarily near obstacles and the difficulty of a problem is often defined by its narrowest passage.

There are two reasons why these measures are sub-optimal. They are local methods which do not take into account the larger context of the roadmap under construction. To them, an obstacle/passageway is equally important even after a roadmap has been successfully constructed around/through it. Samples will continue to be placed near an obstacle or in a passageway long after it has been successfully incorporated into a roadmap.

Secondly, the PRM algorithm was designed to avoid the construction a complete covering of the configuration space. Since such a construction has been shown to require an exponential number of samples in the worst case [9]. However, previous methods for biased sampling in PRM construction have relied on the type of configuration space information, such as the location of obstacles, which is precisely the kind of information the original PRM algorithm was trying to avoid. Instead our methods focus upon creating successful, connected roadmaps, i.e., understanding the connectivity of the configuration space which does not require even an approximate covering to be successful.

Previous work has suggested using information gain as a guide for roadmap construction [11], [10]. In this work the robot was given the ability to explore the world with its sensors and use this information to path plan. Each decision about which area of workspace to explore with the robot's sensors was made in an effort to maximize the expected reduction in entropy of the robot's knowledge of the configuration space. Our work differs importantly from this work since we are focusing on reducing the entropy of the roadmap rather than the entropy of the configuration space.

II. ENTROPY-GUIDED PRM

In contrast to the previous methods (Section I-C) which were a priori and local criteria, we propose an adaptive sampling technique which quickly understands easy spaces, analyzes the roadmap under construction to find difficult regions and focus sample selection there. It will choose the areas in the configuration space which can lead to the greatest improvements in the success of the roadmap. To measure the contribution of a sample to the improvement of the roadmap we use the notion of information gain developed in formal information theory [7], [8].

Entropy is the measure of the stability of the predictability of a distribution. The entropy for some discrete distribution over a set of values D is defined as:

$$H(D) = - \sum_{d \in D} p(d) \log p(d)$$

We can define the entropy of a system when some variable K has a particular value k :

$$H(S|K = k) = - \sum_{d \in D} p(d|K = k) \log p(d|K = k)$$

If we consider gaining some new information K about the system, but only have a distribution over possible values of K we can define conditional entropy of S given K as:

$$H(S|K) = \sum_{k \in K} p(k) H(S|K = k)$$

Information gain represents the change in the entropy of a system as a result of gaining knowledge related to the system. For some system S and some new knowledge K , information gain is:

$$IG(S, K) = H(S) - H(S|K)$$

Entropy has a minimal value at zero, pursuing a strategy which maximizes the decrease in entropy or information gain moves us to a low entropy state as quickly as possible given our current knowledge and representation of the entropy of the system. If we design our entropy measure so that a state with low entropy corresponds to a solution to our problem (in this case path-planning) information gain can guide us along a rapid path to a solution.

A. Driving Sampling with Information Gain

In the case of roadmap building, we are trying to maximize the information gain for each sample we add to the roadmap we are constructing. To use this to inform our sampling algorithm then we must derive a sampling strategy which selects points that maximize the information gain for a particular roadmap. Given some roadmap R we know there exists some set of points

P_R which maximize the information gain for roadmap R . Our task is to design a sampling strategy X , such that $\forall R, X(R) \in P_R$. That is, for any given roadmap, the function X will select a point which maximizes the information gain for that roadmap. Due to our lack of knowledge about the configuration space we are unable to determine the point which maximizes the information gain. Instead select the point which maximize the *expected* information gain.

Thus we propose a modified PRM algorithm which proceeds as follows:

Entropy-Guided PRM

R is an empty roadmap

For iterations

Select $p = X(R)$

Add p to the roadmap in the traditional manner

B. Entropy of the Roadmap

To successfully measure the information gain of a point added to a roadmap we need to analyze the entropy of a roadmap. We must choose some distribution which is characteristic of the roadmap R . (Our choice of distribution is important since it will drive our sampling.) It is necessary that when the entropy of the distribution we select is zero, that the roadmap be complete. If this were not the case then sampling would stop prior to the completion of the roadmap.

Initially it might seem intuitive to use the distribution over the probability of successful path being produced by the roadmap. However, since a roadmap which fails 100% of the time has as low entropy as a roadmap that succeeds 100% of the time, half of the time the decreasing the entropy of this distribution provides the wrong motivation. Whenever the probability of path planning success is less than 50% our sampling algorithm is driven to succeed as little as possible. This is not desirable.

C. A better distribution

Consider instead, the distribution of connected components. Remember that the roadmap R is simply a graph with one or more connected components. Define a function $f(x, R)$, which when given some point in configuration space x returns the closest connected component R_i to which a straight-line path is possible. If a configuration is chosen which cannot connect to any connected components (such as when the roadmap is empty) it creates a new connected component containing simply that point and returns that connected component. This function f applied to a configuration chosen uniformly at random defines a distribution over all points in configuration space.

This function defines a volume which surrounds each connected component. For a connected component R_i ,

this function defines an volume of points A_i such that $\forall p \in A_i, f(R, p) = R_i$. This region is similar to Voronoi regions for the connected components.

We can use these areas to define the probability of a particular point in free configuration space being connected to a particular component. The probability is equal to the volume of the component divided by the total volume of free configuration space. The probability of a point drawn at random producing a particular connected component R_i is equal to A_i/C_{free} .

Using our previous definition of entropy, we sum over all connected components in the roadmap and the entropy of a roadmap is defined to be the entropy over this distribution which is:

$$- \sum_{R_i \in R} P(R_i) \log P(R_i) = - \sum_{R_i \in R} \frac{A_i}{C_{\text{free}}} \log \frac{A_i}{C_{\text{free}}}$$

Adding some point p to a roadmap R will result in a new roadmap R' which will contain a new connected component R_u which consist either solely of p or a combination of p with several connected components already in R . In either event, if we define R'' to be the (possibly empty) set of components which are joined by the point p , then we can define the new roadmap R' to be those components which are in R but not in R'' unioned with R_u . Given this we can define the information gain provided by adding some point p which results in the combination of connected components create R_u :

$$IG(R, p) = H(R) - H(R')$$

$$IG(R, p) = - \sum_{R_i \in R} \frac{A_i}{C_{\text{free}}} \log \frac{A_i}{C_{\text{free}}} - - \sum_{R_j \in R'} \frac{A_j}{C_{\text{free}}} \log \frac{A_j}{C_{\text{free}}}$$

$$IG(R, p) = \sum_{R_i \in R} - \frac{A_i}{C_{\text{free}}} \log \frac{A_i}{C_{\text{free}}} + \sum_{R_j \in R'} \frac{A_j}{C_{\text{free}}} \log \frac{A_j}{C_{\text{free}}}$$

$$IG(R, p) = \frac{1}{C_{\text{free}}} - \sum_{R_i \in R} A_i \log \frac{A_i}{C_{\text{free}}} + \sum_{R_j \in R'} A_j \log \frac{A_j}{C_{\text{free}}}$$

Given this definition, since the area of free configuration space is constant, information gain for the addition of some point p to a roadmap R producing some new roadmap R' is proportional to:

$$IG(R, p) \propto - \sum_{R_i \in R - R''} A_i \log \frac{A_i}{C_{\text{free}}} + \sum_{R_j \in R'} A_j \log \frac{A_j}{C_{\text{free}}}$$

D. Sampling principles

Since the entropy of our current roadmap, and hence the first term, is fixed, information gain is maximized if the second term in the sum is minimized. This observation leads us to several principles for designing entropy guided samplers.

Connecting a point to an existing component is better than creating a new component. Connecting a point to

an existing component R_i will increase the area of that component A'_i this in turn increases the $A_i \log \frac{A_i}{C_{\text{free}}}$ term in the summation which increases the value of the second term as well as the information gain.

Connecting to a larger area is better than connecting to a smaller area. Connecting to a larger area creates an even larger area, with a correspondingly larger probability of selection, as this probability increases the overall entropy of the system drops more than if two probabilities became closer (as they would if a point was added to a smaller connected component), so prefer adding points to larger components.

Connecting the largest volume of previously disjoint connected components will maximize the decrease in entropy of this distribution. Not only does it reduce the number of connected components over which the summation is performed ($|R'| < |R|$) but it also increases the probability that configurations chosen uniformly at random will be nearest the newly merged component. As this probability increases, the log of the probability becomes closer to zero.

III. AN IMPLEMENTATION OF ENTROPY-GUIDED PRM

We use the intuition developed in the previous section to implement an instantiation of an Entropy-Guided PRM algorithm. In this implementation we use connecting pairs of disconnected components in the roadmap as the sampling technique which attempts to maximize the expected information gain. Alternative entropy-guided implementations would use other methods to maximize the information gain, part of our future work is the development of these techniques.

The component-connecting instantiation of entropy-guided sampling extends the standard PRM algorithm in two ways: For each connected component R_i in a roadmap R , an volume of configuration space, A_i , is maintained. At each step when a configuration point p is added to a connected component R_i , the volume A_i is updated to include the new point p .

To sample a new configuration point, two connected components R_i, R_j are selected from the roadmap R . The selection of connected components is biased towards components which are larger and closer together. We choose pairs of connected components with a probability weighted by the expected value of connecting them. This value is given by the probability that a path is possible between the point chosen and each of the components multiplied by the sum of the volumes of both components:

$$P(\text{PathPossible} | \text{distance}(A_i, A_j)/2)(A_i + A_j)$$

This measure is directly proportional to the expected decrease in entropy since the sum of the two volumes reflects the potential decrease in entropy resulting from the connection of the two components, while the probability

of connection reflects the likelihood that such a connection will actually occur.

Once two connected components are chosen, A' , the volume of configuration space which lies between their volumes A_i, A_j is computed. A new uniform random configuration point p is drawn from A' .

There are two special cases in this point selection. The first is that some percentage of the time a new configuration point is chosen uniformly at random from the entire configuration space to ensure total coverage of the configuration space. The actual percentage of points drawn in this manner begins at fifty percent when the roadmap is initialized and decreases as the number of samples grows to a minimum of ten percent. Secondly, whenever roadmap contains a single connected component, points are also drawn uniformly from the entire configuration space. The algorithm completes when a specified number of samples have been chosen.

Although at first blush this method may seem very similar to the component expansion steps of other instantiations of PRM it differs in two important ways; First it is careful in its selection of components to connect in order to maximize the gains of each connection attempt. Second, it uses attempts to connect disconnected components as its sampling strategy throughout roadmap building rather than as a post-processing step.

A. Bounding Boxes as an approximation of components

Optimally, when we maintain A_i the volume around some connected component R_i we would calculate the true volume around each connected component and use that as the input to bias out sampling towards connecting components. The volume around each connected component is defined as the set of points p_i from which there is a straight-line path to some configuration in the connected component and no point in any other connected component is closer to p . Calculating this volume is roughly equivalent to computing the Voronoi region around each connected component. Such calculations are quite expensive. As a result it makes sense to use an approximation for the volume. We use a max-min bounding box as an approximation of the space around a connected component. A max-min bounding box is defined in each dimension by the maximum and minimum values of that dimension over all points in the component.

IV. EXPERIMENTS

Several different PRM algorithms were used for path planning in a number of worlds. For every world we ran each algorithm for a number of samples. At each number of samples we analyzed the accuracy of the roadmap which was built with that number of samples.

We define roadmap accuracy, $p(\text{success})$, to be the probability that a roadmap can find a path between two



Fig. 1. The maze with a 6-DOF robot

random points in free configuration space. To estimate the accuracy of a roadmap we choose one hundred pairs of random free configurations and attempted to path plan between them using the roadmap. The percentage of successful path planning attempts is then used as the accuracy of the roadmap. In each of the graphs shown below each point is the average over fifty independent runs of the path planning algorithm.

The first set of worlds we used to test the planners were two-dimensional planar mazes. The maze we used is shown in Figure 1. Path planning was performed in this maze for robots with six, eight and ten degree-of-freedom, free floating robots.

As can be seen from the results (Figure 2), in these environments the entropy-guided algorithm provided significant improvements in performance, especially as the degree-of-freedom of the robot increased. In the case of the ten degree-of-freedom robot neither OB-PRM, Expansive-PRM or traditional PRM have made any headway in the maze while our implementation of entropy-guided sampling has largely solved the maze.

The second set of worlds we used was also two dimensional planar worlds, however these robots were stationary.

As can be seen from the results, bounding box sampling, while performing better with statistical significance than standard and OB-PRM did not result in large improvements. This demonstrates some of the limitations of the max-min bounding box approximation used for this implementation of entropy-guided PRM. In worlds where the bounding boxes surrounding connected components have a high degree of overlap, the performance of the approximation degrades to that of standard PRM. While it is comforting to know that the max-min bounding box approximation can never do worse than standard PRM, better bounding box approximations are expected to result in the same types of improvements seen in mazes.

A. Orientation of Bounding Boxes

One concern of the bounding box strategy, especially in the maze, is that the workspace is at least partially

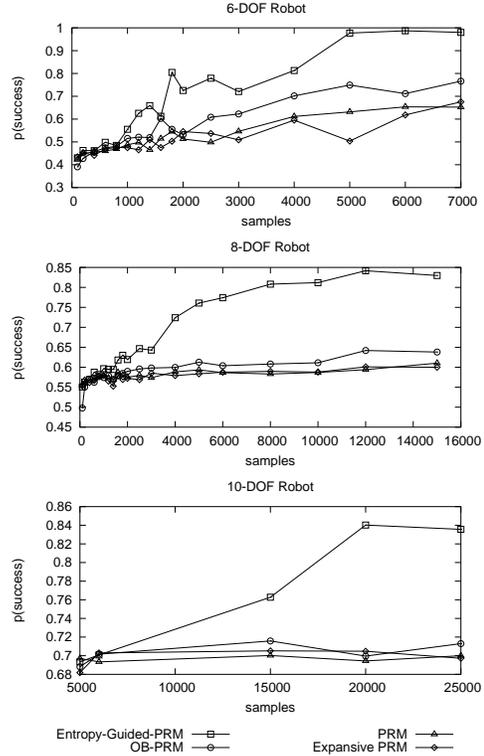


Fig. 2. Results for free-floating robots

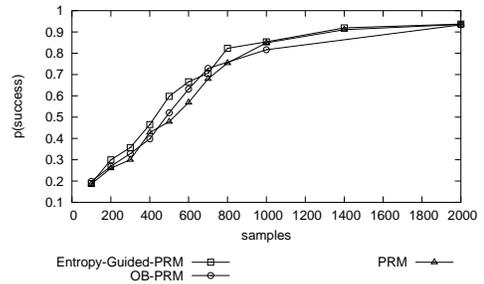


Fig. 3. Results for a 5-DOF stationary robot

oriented to the bounding boxes. To ensure that our success was not based on this alignment, we ran two experiments with the same simple u-shaped maze. In one instance the maze was oriented with the x-y axis (and thus a portion of the bounding box). In the other the maze was rotated forty five degrees.

From the graph in Figure 5 it can be seen that there is not a significant difference between the performance of all three algorithms on the two different worlds. And in both cases entropy-guided PRM outperforms the other algorithms.

B. Comparison to deterministic techniques

In order to compare our work with entropy-guided PRM to current efforts in deterministic sampling, we



Fig. 4. The two mazes which were compared

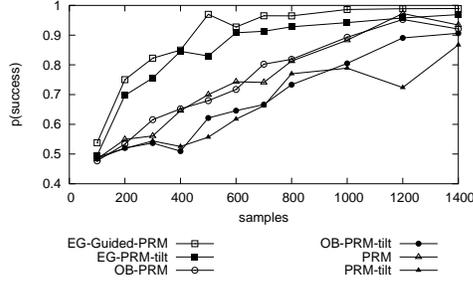


Fig. 5. Results of the maze at different orientations

reproduced the world used in [3] and path-planned in it using our entropy-guided implementation. The results of this are given in table 4. For each algorithm and degree of freedom, the average number of samples necessary to create a roadmap which could path plan through the passageway is listed. As in [3] the numbers are the average of one hundred runs of the algorithms. Entropy-guided PRM significantly outperforms the numbers given for the QRM algorithm.

V. CONCLUSIONS

Decreasing the number of samples necessary to construct successful roadmaps significantly reduces the computational demands of PRM. We have proposed a novel class of sampling strategies for PRM that are based upon information theory. We show that the goal of maximizing information gain suggests a variety of new sampling strategies including one which focuses on connecting the disconnected components in configuration space.

To demonstrate these ideas empirically we have implemented a sampling strategy which uses a simply max-min bounding box to approximate the volume surrounding a connected component. This sampling strategy has been shown to result in significant improvements over existing

DOF	Width	Entropy PRM	QRM
2	0.03	112	195
3	0.05	281	579
6	0.10	3085	4052

TABLE I

THE AVG. NUMBER OF SAMPLES NECESSARY FOR SUCCESS

strategies in multiple environments with increasing benefits in higher dimensional configuration spaces.

ACKNOWLEDGEMENTS

The authors wish to thank our various reviewers for their detailed comments and critiques which have lead us to many improvements.

VI. REFERENCES

- [1] N. Amato, B. Bayazid, L. Dale, C. Jones, and D. Vallejo. OBPRM: An obstacle-based PRM for 3D workspaces. In *Robotics: The Algorithmic Perspective*. AK Peters, 1998.
- [2] M. S. Branicky, S. M. LaValle, K. Olson, and L. Yang. Quasi-randomized path planning. In *Proceedings of the International Conference on Robotics and Automation*, pages 1481–1487, Seoul, Korea, 2001.
- [3] M. S. Branicky, S. M. LaValle, K. Olson, and L. Yang. Deterministic vs. probabilistic roadmaps. Manuscript, 2002.
- [4] D. Hsu, J.-C. Latombe, and R. Motwani. Path planning in expansive configuration spaces. In *Proceedings of the International Conference on Robotics and Automation*, volume 3, pages 2719–2726, 1997.
- [5] L. E. Kavraki, P. Švestka, J.-C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.
- [6] J. H. Reif. Complexity of the mover’s problem and generalizations. In *Proceedings of the Symposium on Foundations of Computer Science*, pages 421–427, 1979.
- [7] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, July 1948.
- [8] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:623–656, October 1948.
- [9] A. G. Sukharev. Optimal strategies of the search for an extremum. *U.S.S.R. Computational Mathematics and Mathematical Physics*, 11(4):119–137, 1971. Translated from Russian, *Zh. Vychisl. Mt. i Mat. Fiz.*, 11(4):910-924.
- [10] P. Wang and K. Gupta. View planning via maximal entropy reduction. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, 2002.
- [11] Y. Yu and K. Gupta. An information theoretic approach to viewpoint planning for motion planning of eye-in-hand systems. In *Proceedings of the International Symposium on Industrial Robotics*, 2000.