

Faster Sample-based Motion Planning using Instance-based Learning

Jia Pan and Sachin Chitta and Dinesh Manocha

Abstract We present a novel approach to improve the performance of sample-based motion planners by learning from prior instances. Our formulation stores the results of prior collision and local planning queries. This information is used to accelerate the performance of planners based on probabilistic collision checking, select new local paths in free space, and compute an efficient order to perform queries along a search path in a graph. We present fast and novel algorithms to perform k -NN (k -nearest neighbor) queries in high dimensional configuration spaces based on locality-sensitive hashing and derive tight bounds on their accuracy. The k -NN queries are used to perform instance-based learning and have a sub-linear time complexity. Our approach is general, makes no assumption about the sampling scheme, and can be used with various sample-based motion planners, including PRM, Lazy-PRM, RRT and RRT*, by making small changes to these planners. We observe up to 100% improvement in the performance of various planners on rigid and articulated robots.

1 Introduction

Motion planning is an important problem in robotics, virtual prototyping and related areas. Most of the practical methods for motion planning of high-DOF (degrees-of-freedom) robots are based on random sampling in configuration spaces, including PRM [13] and RRT [14]. The resulting algorithms avoid explicit computation of obstacle boundaries in the configuration space (\mathcal{C} -space) and use sampling techniques to compute paths in the free space ($\mathcal{C}_{\text{free}}$). The main computations include probing the configuration space for collision-free samples, joining nearby collision-free

Jia Pan and Dinesh Manocha are with the Department of Computer Science, the University of North Carolina, Chapel Hill {panj,dm}@cs.unc.edu · Sachin Chitta is with Willow Garage Inc., Menlo Park, CA 94025, USA sach-inc@willowgarage.com. This research is partially supported by NSF grants 1000579 and 1117127, and willow garage.

samples by local paths, and checking whether the local paths lie in the free space. There is extensive work on different sampling strategies, faster collision checking, or biasing the samples based on local information to handle narrow passages.

The collision detection module is used as an oracle to collect information about the free space and approximate its topology. This module is used to classify a given configuration or a local path as either collision-free (i.e. in $\mathcal{C}_{\text{free}}$) or in-collision (i.e. overlaps with \mathcal{C}_{obs}). Most motion planning algorithms tend to store only the collision-free samples and local paths, and use them to compute a global path from the initial configuration to the goal configuration. However, the in-collision configurations or local paths are typically discarded.

One of our goals is to exploit all prior or historical information related to collision queries and improve the performance of the sample-based planner. Some planners tend to utilize the in-collision configurations or the samples near the boundary of the configuration obstacles (\mathcal{C}_{obs}) to bias the sample generation or improve the performance of planners in narrow passages [4, 9, 18, 23]. However, it can be expensive to perform geometric reasoning based on the outcome of a large number of collision queries in high-dimensional spaces. As a result, most prior planners only use partial or local information about configuration spaces, and can't provide any guarantees in terms of improving the overall performance.

Main Results: We present a novel approach to improve the performance of sample-based planners by learning from prior instances of collision checking, including all in-collision samples. Our formulation uses the historical information generated using collision queries to compute an approximate representation of \mathcal{C} -space as a hash table. Given a new probe or collision query in \mathcal{C} -space, we first perform instance-based learning on the approximate \mathcal{C} -space and compute a collision probability. This probability is used as a similarity result or a prediction of the exact collision query and can improve the efficiency of a planner in the following ways:

- The probability is used to design a collision filter for a local planning query in high-dimensional configuration spaces.
- Explore the \mathcal{C} -space around a given configuration and select a new sample (e.g. for RRT planners [14]).
- Compute an efficient order to perform local planning queries along a path in the graph (e.g. for lazyPRM planners [13]).

The underlying instance-based learning is performed on approximate \mathcal{C} -space using k -NN (k -nearest neighbor) queries. All the prior configurations used by the planning algorithm and their collision outcomes are stored incrementally in a hash table. Given a new configuration or a local path, our algorithm computes the nearest neighbors in the hash table. We use locality-sensitive hashing (LSH) algorithms to perform approximate k -NN computations in high-dimensional configuration spaces. Specifically, we present a line-point k -NN algorithm that can compute the nearest neighbors of a line. We derive bounds on the accuracy and time complexity of k -NN algorithms and show that the collision probability computed using LSH-based k -NN algorithm converges to exact collision detection as the size of dataset increases.

We present improved versions of PRM, lazyPRM, RRT planning algorithms based on instance-based learning. Our approach is general and can be combined with any sampling scheme. Furthermore, it is quite efficient for high-dimensional configuration spaces. We have applied these planners to rigid and articulated robots, and observe up to 100% speedups based on instance-based learning. The additional overheads are in terms of storing the prior instances in a hash table and performing k -NN queries, which take a small fraction of the overall planning time.

The rest of the paper is organized as follows. We survey related work in Section 2. Section 3 gives an overview of sample-based planners, instance-based learning and our approach. We present the learning algorithms and analyze their accuracy and complexity in Section 4. We show the integration of instance-based learning with different motion planning algorithms in Section 5 and highlight the performance of modified planners on various benchmarks in Section 6.

2 Related Work

In this section, we give a brief overview of prior work on the use of machine learning techniques in motion planning and performing efficient collision checking to accelerate sample-based motion planning.

2.1 Machine Learning in Motion Planning

Many techniques have been proposed to improve the performance of sample-based motion planning based on machine learning. Marco et al. [16] combine a set of basic PRM motion planners into a powerful ‘super’ planner by assigning different basic planners to different regions in \mathcal{C} -space, based on offline supervised learning. Some unsupervised learning approaches construct an explicit or implicit representation of $\mathcal{C}_{\text{free}}$ and perform adaptive sampling based on this model. Burns and Brock [5] use entropy to measure each sample’s utility to improve the coverage of PRM roadmap. Hsu et al. [11] adaptively combine multiple sampling strategies to improve the roadmap’s connectivity. Some variants of RRT, which use workspace or taskspace bias (e.g., [10]), can be extended by changing the bias parameters adaptively. Scholz and Stilman [21] combine RRT with reinforcement learning. Given a sufficient number of collision-free samples in narrow passage, learning techniques have been used to estimate a zero-measure subspace to bias the sampling in narrow passages [7]. Our approach is complimentary to all these techniques.

2.2 Collision Detection Oracle in Motion Planning

Some of the previous approaches tend to exploit the knowledge about \mathcal{C} -space gathered using collision checking. Boor et al. [4] use pairs of collision-free and in-collision samples to collect information about \mathcal{C}_{obs} and perform dense sampling near \mathcal{C}_{obs} . The same idea is used in many variants of RRT, such as retraction-based planners [18]. Sun et al. [23] bias sampling near narrow passages by using two in-collision samples and one collision-free sample to identify narrow passages in \mathcal{C} -space. Kavraki et al. [13] use in-collision samples to estimate the visibility of each sample and perform heavier sampling in regions with small visibility. Denny and Amato [9] present a variation of PRM that memorizes in-collision samples and constructs roadmaps in both $\mathcal{C}_{\text{free}}$ and \mathcal{C}_{obs} , in order to generate more samples in narrow passages. All of these methods utilize in-collision samples to provide better sampling strategies for the planners, in the form of different heuristics. Our approach neither makes assumptions about the underlying sampling scheme nor biases the samples. As a result, our algorithm can be combined with these methods.

2.3 k -Nearest Neighbor (k -NN) Search

The problem of finding the k -nearest neighbor within a database of high-dimensional points is well-studied in various areas, including databases, computer vision, and machine learning. Samet’s book [20] provides a good survey of various techniques used to perform k -NN search. In order to handle large and high-dimensional spaces, most practical algorithms are based on approximate k -NN queries [6]. In these formulations, the algorithm is allowed to return a point whose distance from the query point is at most $1 + \epsilon$ times the distance from the query to its k -nearest points; $\epsilon > 1$ is called the *approximation factor*.

3 Overview

In this section, we give an overview of the sample-based planner and provide a brief background on instance-based learning.

3.1 Notations and Symbols

We denote the configuration space as \mathcal{C} -space, and each point within the space as a configuration \mathbf{x} . \mathcal{C} -space is composed of two parts: the collision-free points ($\mathcal{C}_{\text{free}}$) and the in-collision points (\mathcal{C}_{obs}). \mathcal{C} -space may be non-Euclidean, but it is possible to approximately embed a non-Euclidean space into a higher-dimensional Euclidean

space (e.g., using the Linial-London-Robinovich embed [15]) to perform k -nearest neighbor queries. A *local path* in \mathcal{C} -space is a continuous curve that connects two configurations. It is difficult to compute \mathcal{C}_{obs} or $\mathcal{C}_{\text{free}}$ explicitly, therefore sample-based planners use collision checking between the robot and obstacles to probe the \mathcal{C} -space implicitly. These planners perform two kinds of queries: the *point query* and the *local path query*. We use the symbol Q to denote either of these queries.

3.2 Enhance Motion Planner with Instance-based Learning

The goal of a motion planner is to compute a collision-free continuous path between the initial and goal configurations in \mathcal{C} -space. The resulting path should completely lie in $\mathcal{C}_{\text{free}}$ and should not intersect with \mathcal{C}_{obs} . As shown in Figure 1(a), sample-based planners learn about the connectivity of \mathcal{C} -space implicitly based on collision queries. The query results can also be used to bias the sampling scheme of the planner via different heuristics (e.g., retraction rules).

Instance-based learning is a well known family of algorithms in machine learning that learn properties of new problem instances by comparing them with the instances observed earlier that have been stored in memory [19]. In our case, we store all the results of prior collision queries, including collision-free as well as in-collision queries. Our goal is to sufficiently exploit such prior information and accelerate the planner computation. The problem instance in our context is the collision query being performed on a given configuration or a local path in \mathcal{C} -space. In particular, performing exact collision queries for local planning can be expensive. Collision checking for a discrete configuration is relatively cheap, but still can be time consuming if the environment or robot’s geometric representation has a high complexity. We utilize the earlier instances or the stored information by performing k -NN queries and geometric reasoning on query results.

Our new approach to exploit prior information for motion planning is shown in Figure 1(b). When the collision checking routine finishes probing the \mathcal{C} -space for a given query, it stores all the obtained information in a dataset corresponding to historical collision query results. If the query is a point within \mathcal{C} -space, the stored information is its binary collision status. If the query is a local path, the stored information includes the collision status of different configuration points along the path. The resulting dataset of historical collision results constitutes the complete knowledge we have about \mathcal{C} -space, coming from collision checking routines. Therefore, we use it as an approximate description of the underlying \mathcal{C} -space: the in-collision samples are an approximation of \mathcal{C}_{obs} , while the collision-free samples are used to encode $\mathcal{C}_{\text{free}}$. These samples are used by instance-based learning algorithms to estimate the collision status of new queries.

Given a new query Q , either a point or a local path, we first perform k -NN search on the dataset to find its neighbor set S , which provides information about the local \mathcal{C} -space around the query. If S contains sufficient information to infer the collision status of the query, we compute a collision probability for the new query based on

Algorithm 1: learning-based-collision-query(Q)

```

begin
  if  $Q$  is point query then
     $S \leftarrow \text{point-point-}k\text{-NN}(Q)$ 
    if  $S$  provides sufficient information for reasoning then
       $\perp$  approximate-collision-query( $S, Q$ )
    else exact-collision-query( $S, Q$ )
  if  $Q$  is line query then
     $S \leftarrow \text{line-point-}k\text{-NN}(Q)$ 
    if  $S$  provides sufficient information for reasoning then
       $\perp$  approximate-continuous-collision-query( $S, Q$ )
    else exact-continuous-collision-query( $S, Q$ )
end

```

S ; otherwise we perform exact collision checking for this query. The calculated collision probability provides prior information about the collision status of the given query and is useful in many ways. First, it can be used as a culling filter to avoid the exact (and expensive) collision checking for queries that correspond to the configurations or local paths deep inside $\mathcal{C}_{\text{free}}$ or \mathcal{C}_{obs} . Secondly, it can be used to decide an efficient order to perform exact collision checking for a set of queries. For example, many planners like RRT need to select a local path that can best improve the local exploration in $\mathcal{C}_{\text{free}}$, i.e., a local path with a long length but is also collision-free. The collision probability computation can be used to compute an efficient sorting strategy and thereby reduces the number of exact collision tests.

Whether we have sufficient information about S (in Algorithm 1) is related to how much confidence we have in terms of performing inferencing from S . For example, if there exists an in-collision sample very close to the new query, then there is a high probability that the new query is also an in-collision query. The probability decreases when the distance between the sample and the query increases. A description of our prior instance based collision framework is given in Algorithm 1, which is used as an inexpensive routine to perform probabilistic collision detection. More details about this routine and its applications are given in Section 4.

4 Probabilistic Collision Detection based on Instance-based Learning

In this section, we discuss how to avoid the expensive exact collision detection query by estimating the collision probability for a given query. The estimation is implemented by performing efficient k -NN queries on the historical collision query results for the given environment.

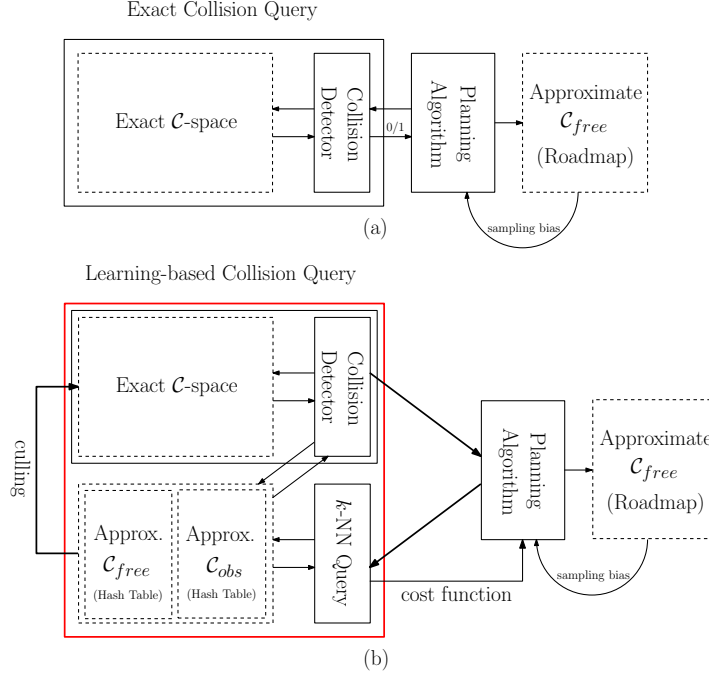


Fig. 1: Comparison between how collision checking is used in prior approaches (a) and our method (b). (a) The collision detection routine is the oracle used by the planner to gather information about \mathcal{C}_{free} and \mathcal{C}_{obs} . The planner performs binary collision queries, either on point configurations or 1-dimensional local paths, and estimates the connectivity of \mathcal{C}_{free} (shown as Approximate \mathcal{C}_{free}). Moreover, some planners utilize the in-collision results to bias sample generation by using different heuristics. (b) Our method also performs collision queries. However, we store all in-collision results (as Approximate \mathcal{C}_{obs}) in addition to collision-free results (as Approximate \mathcal{C}_{free}). Before performing an exact collision query, our algorithm performs a k -NN query on the given configuration or local path to compute a collision probability for each query. The collision probability can be used as a cost function to compute an efficient strategy to perform exact collision queries during motion planning. We use novel LSH-based algorithms to perform k -NN queries efficiently and speed up the overall planner.

4.1 LSH-based Approximate k -NN Query

A key issue in terms of the instance-based learning framework is its computational efficiency. As we generate hypotheses directly from training instances, the complexity of this computation can grow with the size of historical data. If we use exact k -NN computation as the learning method, its complexity can be a linear function of the size of the dataset, especially for high-dimensional spaces. To improve the efficiency of the instance-based learning, we use approximate k -NN algorithms.

Given a dataset $\mathcal{D} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ of N points in \mathcal{R}^d , we consider two types of retrieval queries. One is to retrieve points from \mathcal{D} that are closest to a given point query, i.e. the well-known k -NN query, and we call it the *point-point k -NN* query.

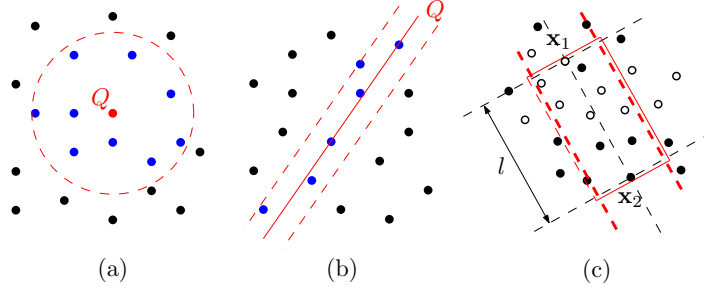


Fig. 2: Two types of k -NN queries used in our method: (a) point-point k -NN; (b) line-point k -NN. Q is the query item and the results of different queries are shown as blue points in each figure. We present novel LSH-based algorithms for fast computation of these queries. (c) Use line-point k -NN query to compute prior instances that can influence the collision status of a local path, which connects \mathbf{x}_1 and \mathbf{x}_2 in \mathcal{C} -space. The query line is the line segment between \mathbf{x}_1 and \mathbf{x}_2 . The white points are prior collision-free samples in the dataset and the black points are prior in-collision samples.

The second query tries to find the points from \mathcal{D} that are closest to a given line in \mathcal{R}^d , whose direction is \mathbf{v} and passes through a point \mathbf{a} , where $\mathbf{v}, \mathbf{a} \in \mathcal{R}^d$. We call the second query the *line-point k -NN* query. The two types of k -NN queries are illustrated in Figure 2.

In order to develop an efficient instance-based learning framework, we use locality-sensitive hashing (LSH) as an approximate method for k -NN queries, which is mainly designed for point-point queries [1]. However, it can be extended to line queries [2] and hyperplane queries [12]. Basri [3] et al. further extend it to perform point/subspace queries.

In this work, we design an efficient LSH-based approximate line-point k -NN query and also derive its error bound similar to the error bound [8] derived for LSH-based point-point k -NN query. The new error bound is important because it enables us to derive the error bound for our approximate collision detection algorithm and eventually prove the completeness of the instance-learning based motion planners. Please refer to [17] for more details of our new line-point k -NN query algorithm, which is omitted in this paper due to limit of space.

4.2 \mathcal{C} -space Reasoning based on k -NN Queries

Our approach stores the outcome of prior instances of exact collision queries, including point queries and local path queries, within a database (shown as Approximate $\mathcal{C}_{\text{free}}$ and Approximate \mathcal{C}_{obs} in Figure 1(b)). In this section, we describe our approach to use those stored instances to perform probabilistic collision queries.

The first case is the query point, i.e., the task is to compute the collision status for a sample \mathbf{x} in \mathcal{C} -space. We first perform point-point k -NN query to compute

the prior collision instances closest to \mathbf{x} . Next, we use a simple k -NN classifier to estimate \mathbf{x} 's *collision probability* as

$$\mathbb{P}[\mathbf{x} \text{ in collision}] = \frac{\sum_{\mathbf{x}_i \in S \setminus \mathcal{C}_{\text{free}}} w_i}{\sum_{\mathbf{x}_i \in S} w_i}, \quad (1)$$

where $w_i = e^{-\lambda \text{dist}(\mathbf{x}_i, \mathbf{x})}$ is the distance-tuned weight for each k -NN neighbor and S is the neighborhood set computed using point-point k -NN query. In this case, the parameter λ controls the magnitude of the weights. λ corresponds to the *obstacle density* of \mathcal{C}_{obs} , if we model \mathcal{C}_{obs} as a point set generated using a Poission point process in the \mathcal{C} -space [22].

The second case is the line query, i.e., the goal is to estimate the collision status of a local path in \mathcal{C} -space. We require the local path to lie within the neighborhood of the line segment l connecting its two endpoints, i.e., the local path should not deviate too much from l . The first step is to perform a line-point k -NN query to find the prior collision query results closest to the infinite line that l lies on. Next, we need to filter out the points in S whose projections are outside the truncated segment of l , as shown in Figure 2(c). Finally, we apply our learning method on the filtered results, denoted as S , to estimate the collision probability of the local path.

We compute the line's collision probability via an optimization scheme. The line l is divided into I segments and we assign each segment, say l_i , a label c_i to indicate its collision status: we assign label $c_i = 0$ if l_i is collision-free and assign label $c_i = 1$ if it is in-collision. Given line l 's neighborhood set S computed using the line-point k -NN query, we now compute the label assignments for these segments. First, we compute the conditional collision probability of one point $\mathbf{x}_j \in S$, given the collision status of one line segment l_i :

$$\mathbb{P}[\mathbf{x}_j \text{ in collision} \mid c_i] = \begin{cases} 1 - e^{-\lambda \text{dist}(l_i, \mathbf{x}_j)}, & c_i = 0, \\ e^{-\lambda \text{dist}(l_i, \mathbf{x}_j)}, & c_i = 1; \end{cases} \quad (2)$$

$$\mathbb{P}[\mathbf{x}_j \text{ collision free} \mid c_i] = \begin{cases} e^{-\lambda \text{dist}(l_i, \mathbf{x}_j)}, & c_i = 0, \\ 1 - e^{-\lambda \text{dist}(l_i, \mathbf{x}_j)}, & c_i = 1, \end{cases} \quad (3)$$

where $\text{dist}(l_i, \mathbf{x}_j)$ is the distance between \mathbf{x}_j and l_i 's midpoint. Given this formalization, we can compute l_i 's posterior collision probability. given l 's neighborhood set S :

$$\mathbb{P}[c_i \mid S] \propto \mathbb{P}[S \mid c_i] \cdot \mathbb{P}[c_i] = \prod_{\mathbf{x}_j \in S} \mathbb{P}[\mathbf{x}_j \text{'s collision status} \mid c_i] \cdot \mathbb{P}[c_i]. \quad (4)$$

The error in this label assignment, i.e., the probability that the computed l_i 's label is not the same as the outcome of the exact collision query, is

$$\mathbb{P}_{\text{error}}[c_i \mid S] = c_i \cdot \mathbb{P}[c_i = 0 \mid S] + (1 - c_i) \cdot \mathbb{P}[c_i = 1 \mid S]. \quad (5)$$

The label assignment algorithm needs to minimize this error probability for each segment l_i . Moreover, the assignment should be coherent, i.e., there is a high prob-

ability that adjacent line segments should have the same label. As a result, we can compute a suitable label assignment $\{c_i^*\}_{i=1}^I$ using a binary integer programming:

$$\{c_i^*\} = \operatorname{argmin}_{\{c_i\} \in \{0,1\}^I} \sum_{i=1}^I \mathbb{P}_{\text{error}}[c_i | S] + \kappa \sum_{i=1}^{I-1} (c_i - c_{i+1})^2, \quad (6)$$

where κ is a weight parameter. The binary integer programming problem can be solved efficiently using dynamic programming. After that, we can estimate the collision probability for the line as

$$\mathbb{P}[l \text{ in collision}] = \max_{i: c_i^*=1} \mathbb{P}[c_i = 1 | S]. \quad (7)$$

As a byproduct, the approximate first time of contact can be given as $\min_i: c_i^*=1 i/I$.

A natural way to use the collision probability formulated as above is to use a specific threshold to justify whether a given query is in-collision or not: if the query's collision probability is larger than the given threshold, we return in-collision; otherwise we return collision-free. We can prove that, for any threshold $0 < t < 1$, the collision status returned by the instance-based learning will converge to the exact collision detection results, when the size of the dataset increases (asymptotically):

Theorem 1. *The collision query performed using LSH-based k -NN will converge to the exact collision detection as the size of the dataset increases, for any threshold between 0 and 1.*

5 Accelerating Sample-based Planners

In this section, we first discuss techniques to accelerate various sample-based planners based on instance-based learning. Next, we analyze the factors that can influence the performance of the learning-based planners. Finally, we prove the completeness and optimality for the instance-learning enhanced planners.

The basic approach to benefit from the learning framework is highlighted in Algorithm 1, i.e., use the computed collision probability as a filter to reduce the number of exact collision queries. If a given configuration or local path query is close to in-collision instances, then it has a high probability of being in-collision. Similarly, if a query has many collision-free instances around it, it is likely to be collision-free. In our implementation, we only cull away queries with high collision probability. For queries with high collision-free probability, we still perform exact collision tests on them in order to guarantee that the overall collision detection algorithm is conservative. In Figure 3(a), we show how our probabilistic culling strategy can be integrated with PRM algorithm by only performing exact collision checking (collide) for queries with collision probability (icollide) larger than a given threshold t . Note that the neighborhood search routine (near) can use LSH-based point-point k -NN query.

<pre> sample($\mathcal{D}^{\text{out}}, n$) $V \leftarrow \mathcal{D} \cap \mathcal{C}_{\text{free}}, E \leftarrow \emptyset$ foreach $v \in V$ do $U \leftarrow \text{near}(G^{V,E}, v, \mathcal{D}^{\text{in}})$ foreach $u \in U$ do if $\text{icollide}(v, u, \mathcal{D}^{\text{in}}) < t$ if $\neg \text{collide}(v, u, \mathcal{D}^{\text{out}})$ $E \leftarrow E \cup (v, u)$ near: nearest neighbor search. icollide: probabilistic collision checking based on k-NN. collide: exact local path collision checking. $\mathcal{D}^{\text{in/out}}$: prior instances as input/output. </pre> <p>(a) I-PRM</p>	<pre> sample($\mathcal{D}^{\text{out}}, n$) $V \leftarrow \mathcal{D} \cap \mathcal{C}_{\text{free}}, E \leftarrow \emptyset$ foreach $v \in V$ do $U \leftarrow \text{near}(G^{V,E}, v, \mathcal{D}^{\text{in}})\{v\}$ foreach $u \in U$ do $w \leftarrow \text{icollide}(v, u, \mathcal{D}^{\text{in}})$ $l \leftarrow \ (v, u)\$ $E \leftarrow E \cup (v, u)^{w,l}$ do search path p on $G(V, E)$ which minimizes $\sum_e l(e) + \lambda \min_e w(e)$. foreach $e \in p$, $\text{collide}(e, \mathcal{D}^{\text{out}})$ while p not valid </pre> <p>(b) I-lazyPRM</p>
<pre> $V, \mathcal{D} \leftarrow x_{\text{init}}, E \leftarrow \emptyset$ while x_{goal} not reach $x_{\text{rnd}} \leftarrow \text{sample-free}(\mathcal{D}^{\text{out}}, 1)$ $x_{\text{nst}} \leftarrow \text{inearst}(G^{V,E}, x_{\text{rnd}}, \mathcal{D}^{\text{in}})$ $x_{\text{new}} \leftarrow \text{isteer}(x_{\text{nst}}, x_{\text{rnd}}, \mathcal{D}^{\text{in,out}})$ if $\text{icollide}(x_{\text{nst}}, x_{\text{new}}) < t$ if $\neg \text{collide}(x_{\text{nst}}, x_{\text{new}})$ $V \leftarrow V \cup x_{\text{new}}, E \leftarrow E \cup (x_{\text{new}}, x_{\text{nst}})$ inearst: find the nearest tree node that is long and has high collision-free probability. isteer: steer from a tree node to a new node, using icollide for validity checking. rewire: RRT* routine used to update the tree topology for optimality guarantee. </pre> <p>(c) I-RRT</p>	<pre> $V, \mathcal{D} \leftarrow x_{\text{init}}, E \leftarrow \emptyset$ while x_{goal} not reach $x_{\text{rnd}} \leftarrow \text{sample-free}(\mathcal{D}^{\text{out}}, 1)$ $x_{\text{nst}} \leftarrow \text{inearst}(G^{V,E}, x_{\text{rnd}}, \mathcal{D}^{\text{in}})$ $x_{\text{new}} \leftarrow \text{isteer}(x_{\text{nst}}, x_{\text{rnd}}, \mathcal{D}^{\text{in,out}})$ if $\text{icollide}(x_{\text{nst}}, x_{\text{new}}) < t$ if $\neg \text{collide}(x_{\text{nst}}, x_{\text{new}})$ $V \leftarrow V \cup x_{\text{new}}$ $U \leftarrow \text{near}(G^{V,E}, x_{\text{new}})$ foreach $x \in U$, compute weight $c(x) =$ $\lambda \ (x, x_{\text{new}})\ + \text{icollide}(x, x_{\text{new}}, \mathcal{D}^{\text{in}})$ sort U according to weight c. Let x_{min} be the first $x \in U$ with $\neg \text{collide}(x, x_{\text{new}})$ $E \leftarrow E \cup (x_{\text{min}}, x_{\text{new}})$ foreach $x \in U$, $\text{rewire}(x)$ </pre> <p>(d) I-RRT*</p>

Fig. 3: Instance-based learning framework can be used to improve different motion planners. Here we present four modified planners.

In Figure 3(b), we show how to use the collision probability as a cost function with the lazyPRM algorithm [13]. In the basic version of lazyPRM algorithm, the expensive local path collision checking is delayed till the search phase. The basic idea is that the algorithm repeatedly searches the roadmap to compute the shortest path between the initial and the goal node, performs collision checking along the edges, and removes the in-collision edges from the roadmap. However, the shortest path usually does not correspond to a collision-free path, especially in complex environments. We improve the lazyPRM planning using instance-based learning. We compute the collision probability for each roadmap edge during roadmap construction, based on Equation 7. The probability (w) as well as the length of the edge (l) are stored as the costs of the edge. During the search step, we try to compute a shortest path with minimum collision probability, i.e., a path that minimizes the cost $\sum_e l(e) + \lambda \min_e w(e)$, where λ is a parameter that controls the relative weightage of path length and collision probability. As the prior knowledge about the obstacles is considered based on collision probability, the resulting path is more likely to be collision-free.

Finally, the collision probability can be used by motion planner to explore $\mathcal{C}_{\text{free}}$ in an efficient manner. We use RRT to illustrate this benefit (Figure 3(c)). Given

a random sample x_{rnd} , RRT computes a node x_{nst} among prior collision-free configurations that are closest to x_{rnd} and expands from x_{nst} towards x_{rnd} . If there is no obstacle in \mathcal{C} -space, this exploration technique is based on Voronoi heuristic that biases planner in the unexplored regions. However, the existence of obstacles affects its performance: the planner may run into \mathcal{C}_{obs} shortly after expansion and the resulting exploration is limited. Based on instance-based learning, we can first estimate the collision probability for local paths connecting x_{rnd} with each of its neighbors and choose x_{nst} to be the one with a long edge length, but a small collision probability, i.e., $x_{\text{nst}} = \arg\max(l(e) - \lambda \cdot w(e))$, where λ is a parameter used to control the relative weight of these two terms. A similar strategy can also be used for RRT*, as shown in Figure 3(d).

The learning-based planners are faster, mainly because we replace part of the expensive, exact collision queries with relatively cheap k -NN queries. Let the timing cost for a single exact collision query be T_C and for a single k -NN query be T_K , where $T_K < T_C$. Suppose the original planner performs C_1 collision queries and the instance-based learning enhanced planners performs C_2 collision queries and $C_1 - C_2$ k -NN queries, where $C_2 < C_1$. We also assume that the two planners spend the same time A on other computations within a planner, such as sample generation, maintaining the roadmap, etc. Then the speedup ratio obtained by instance-based learning is:

$$R = \frac{T_C \cdot C_1 + A}{T_C \cdot C_2 + T_K \cdot (C_2 - C_1) + A}. \quad (8)$$

Therefore, if $T_C \gg T_K$ and $T_C \cdot C_1 \gg A$, we have $R \approx C_1/C_2$, i.e., if a higher number of exact collision queries are culled, we can obtain a higher speedup. The extreme speedup ratio C_1/C_2 may not be reached, because 1) $T_C \cdot C_1 \gg A$ may not hold, such as when the planner is in narrow passages (A is large) or in open spaces ($T_C \cdot C_1$ is small); 2) $T_C \gg T_K$ may not hold, such as when the environment and robot have low geometric complexity (i.e., T_C is small) or the instance dataset is large and the cost of the resulting k -NN query is high (i.e., T_K is large).

Note that R is only an approximation of the actual acceleration ratio. It may overestimate the speedup because a collision-free local path may have collision probability higher than a given threshold and our method will filter it out. If such a collision-free local path is critical for the connectivity of the roadmap, such false positives due to instance-based learning will cause the resulting planner to perform more exploration and thereby increases the planning time. As a result, we need to choose an appropriate threshold that can provide a balance: we need a large threshold to filter out more collision queries and increase R ; at the same time, we need to use a small threshold to reduce the number of false positives. However, the threshold choice is not important in asymptotic sense, because according to Theorem 1, the false positive error converges to 0 when the dataset size increases.

R may also underestimate the actual speedup. The reason is that the timing cost for different collision queries can be different. For configurations near the boundary of \mathcal{C}_{obs} , the collision queries are more expensive. Therefore, the timing cost of checking the collision status for an in-collision local path is usually larger than that of a collision-free local path, because the former always has one configuration

on the boundary of \mathcal{C}_{obs} . As a result, it is possible to obtain a speedup larger than C_1/C_2 .

Finally, as a natural consequence of Theorem 1, we can prove the completeness and optimality of the new planners:

Theorem 2. *I-PRM, I-lazyPRM, I-RRT are probabilistic complete. I-RRT* is probabilistic complete and asymptotic optimal.*

6 Results

In this section, we highlight the performance of our new planners. Figure 4 and Figure 5 show the articulated PR2 and rigid body benchmarks we used to evaluate the performance. We evaluated each planner on different benchmarks, and for each combination of planner and benchmark we ran 50 instances of the planner and computed the average planning time as an estimate of the planner’s performance on this benchmark.

The comparison results are shown in Table 4 and Table 5, corresponding to PR2 benchmarks and rigid body benchmarks, respectively. Based on these benchmarks, we observe:

- The learning-based planners provide more speedup on articulated models. The reason is that exact collision checking on articulated models is more expensive than exact collision checking on rigid models. This makes T_C larger and results in larger speedups.
- The speedup of I-PRM over PRM is relatively large, as exact collision checking takes a significant fraction of overall time within PRM algorithm. I-lazyPRM also provides good speedup as the candidate path nearly collision-free and can greatly reduce the number of exact collision queries in lazy planners. The speedups of I-RRT and I-RRT* are limited or can even be slower than the original planners, especially on simple rigid body benchmarks. The reason is that the original planners are already quite efficient on the simple benchmarks and instance-based learning can result in some overhead.
- On benchmarks with narrow passages, our approach does not increase the probability of finding a solution. However, probabilistic collision checking is useful in culling some of the colliding local paths.

We need to point out that the acceleration results shown in Table 4 and Table 5 show only part of the speedups that can be obtained using learning-based planners. As more collision queries are performed and results stored in the dataset, the resulting planner has more information about \mathcal{C}_{obs} and $\mathcal{C}_{\text{free}}$, and becomes effective in terms of culling. Ideally, we can filter out all in-collision queries and obtain a high speedup. In practice, we don’t achieve ideal speedups due to two reasons: 1) we only have a limited number of samples in the dataset; 2) the overhead of k -NN query increases as the dataset size increases. As a result, when we perform the global planning query repeatedly, the planning time will decrease to a minimum, and then




Fig. 4: PR2 planning benchmarks: robot arms with different colors show the initial and goal configurations. The first three benchmarks are of the same environment, but the robot’s arm is in different places: (a) moves arm from under desk to above desk; (b) moves arm from under desk to another position under desk and (c) moves arm from inside the box to outside the box. In the final benchmark, the robot tries to move arm from under a shelf to above it. The difficulty order of the four benchmarks is (c) > (d) > (b) > (a).




Fig. 5: Rigid body planning benchmarks: from left to right, apartment, cubicles, easy, twistycool, flange and torus. Apartment tries to move the piano to the hallway near the door entrance; cubicles moves the robot in a simple office-like environment and the robot needs to fly through the basement; easy and twistycool are of a similar environment, but twistycool contains a narrow passage. Both flange and torus contain a narrow passage.

increases. This is shown in Figure 6. To further improve the performance, we need to adaptively change the LSH parameters to perform k -NN queries efficiently for datasets of varying sizes.

	PRM	I-PRM	lazyPRM	I-lazyPRM	RRT	I-RRT	RRT*	I-RRT*
(a)	12.78	9.61 (32%)	1.2	0.87 (37%)	0.96	0.75 (28%)	1.12	1.01 (11%)
(b)	23.7	12.1 (96%)	1.7	0.90 (88%)	1.36	0.89 (52%)	2.08	1.55 (34%)
(c)	fail	fail	fail	fail	4.15	2.77 (40%)	3.45	2.87 (20%)
(d)	18.5	13.6 (36%)	2.52	1.06 (37%)	7.72	5.33 (44%)	7.39	5.42 (36%)

Table 1: Performance comparison on different combinations of planners and PR2 benchmarks (in milliseconds). ‘fail’ means all the queries cannot find a collision-free path within 1,000 seconds. The percentage in the brackets shows the speedup obtained using instance-based learning.

7 Conclusion

In this paper, we use instance-based learning to improve the performance of sample-based motion planners. The basic idea is to store the prior collision results as an ap-

	PRM	I-PRM	lazyPRM	I-lazyPRM	RRT	I-RRT	RRT*	I-RRT*
apartment	5.25	2.54 (106%)	2.8	1.9 (47%)	0.09	0.10 (-10%)	0.22	0.23 (5%)
cubicles	3.92	2.44 (60%)	1.62	1.37 (19%)	0.89	0.87(2%)	1.95	1.83 (7%)
easy	7.90	5.19 (52%)	3.03	2.01 (50%)	0.13	0.15(-13%)	0.26	0.27 (-4%)
flange	fail	fail	fail	fail	48.47	25.6 (88%)	46.07	26.9 (73%)
torus	31.52	23.3 (39%)	4.16	2.75 (51%)	3.95	2.7 (46%)	6.01	4.23 (42%)
twistycool	1/50	3/50	2/50	3/50	4/50	3/50	2/50	3/50

Table 2: Performance comparison on different combinations of planners and rigid body benchmarks (in milliseconds). The percentage in the brackets shows the speedup based on instance-based learning. For twistycool, which has a narrow passage, we record the number of successful queries among 50 queries (for a 1000 second budget).

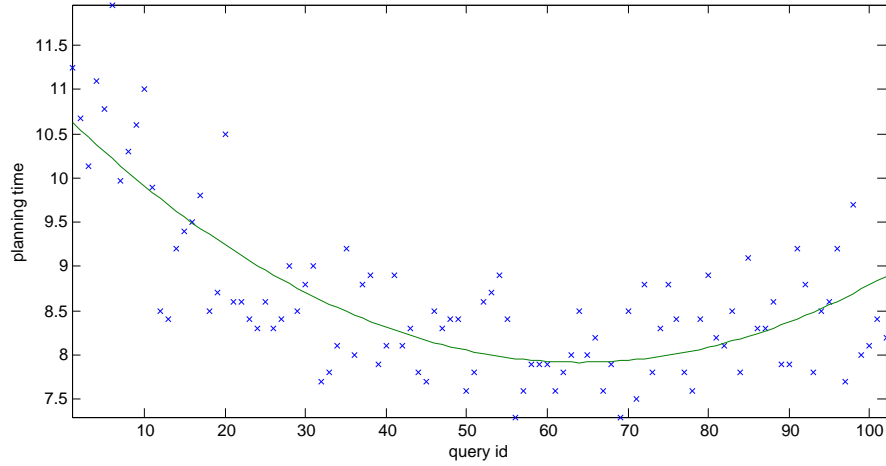


Fig. 6: The time taken by I-PRM when it runs more than 100 times on the benchmark shown in Figure 4 (b). The planning time of a single query first decreases and then increases. The best acceleration acquired is $12.78/7.5 = 70\%$, larger than the 32% in Table 4.

proximate representation of \mathcal{C}_{obs} and $\mathcal{C}_{\text{free}}$ and replace the expensive exact collision detection query by a relatively cheap probabilistic collision query. We integrate approximate collision routine with various sample-based motion planners and observe 30 – 100% speedup on rigid and articulated robots.

References

1. Andoni, A., Indyk, P.: Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. *Communications of the ACM* **51**(1), 117–122 (2008)
2. Andoni, A., Indyk, P., Krauthgamer, R., Nguyen, H.L.: Approximate line nearest neighbor in high dimensions. In: *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 293–301 (2009)

3. Basri, R., Hassner, T., Zelnik-Manor, L.: Approximate nearest subspace search. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **33**(2), 266–278 (2011)
4. Boor, V., Overmars, M., van der Stappen, A.: The gaussian sampling strategy for probabilistic roadmap planners. In: *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 1018–1023 (1999)
5. Burns, B., Brock, O.: Toward optimal configuration space sampling. In: *Proceedings of Robotics: Science and Systems* (2005)
6. Chakrabarti, A., Regev, O.: An optimal randomised cell probe lower bound for approximate nearest neighbour searching. In: *Proceedings of IEEE Symposium on Foundations of Computer Science*, pp. 473–482 (2004)
7. Dalibard, S., Laumond, J.P.: Linear dimensionality reduction in random motion planning. *International Journal of Robotics Research* **30**(12), 1461–1476 (2011)
8. Datar, M., Immorlica, N., Indyk, P., Mirrokni, V.S.: Locality-sensitive hashing scheme based on p-stable distributions. In: *Proceedings of Symposium on Computational Geometry*, pp. 253–262 (2004)
9. Denny, J., Amato, N.M.: The toggle local planner for probabilistic motion planning. In: *Proceedings of IEEE International Conference on Robotics and Automation* (2012)
10. Diankov, R., Ratliff, N., Ferguson, D., Srinivasa, S., Kuffner, J.: Bispaces planning: Concurrent multi-space exploration. In: *Proceedings of Robotics: Science and Systems* (2008)
11. Hsu, D., Sanchez-Ante, G., Sun, Z.: Hybrid PRM sampling with a cost-sensitive adaptive strategy. In: *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 3874–3880 (2005)
12. Jain, P., Vijayanarasimhan, S., Grauman, K.: Hashing hyperplane queries to near points with applications to large-scale active learning. In: *Neural Information Processing Systems* (2010)
13. Kavraki, L., Svestka, P., Latombe, J.C., Overmars, M.: Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation* **12**(4), 566–580 (1996)
14. Kuffner, J., LaValle, S.: RRT-connect: An efficient approach to single-query path planning. In: *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 995–1001 (2000)
15. Linial, N., London, E., Rabinovich, Y.: The geometry of graphs and some of its algorithmic applications. *Combinatorica* **15**(2), 215–245 (1995)
16. Morales, M., Tapia, L., Pearce, R., Rodriguez, S., Amato, N.M.: A machine learning approach for feature-sensitive motion planning. In: *Proceedings of International Workshop on Algorithmic Foundation of Robotics*, pp. 361–376 (2004)
17. Pan, J., Chitta, S., Manocha, D.: Faster sample-based motion planning using instance-based learning. Tech. rep., University of North Carolina, Chapel Hill (2012). Available as <http://cs.unc.edu/~panj/techreport.pdf>
18. Rodriguez, S., Tang, X., Lien, J.M., Amato, N.: An obstacle-based rapidly-exploring random tree. In: *Proceedings of IEEE International Conference on Robotics and Automation*, pp. 895–900 (2006)
19. Russell, S., Norvig, P.: *Artificial Intelligence: A Modern Approach*. Prentice Hall (2003)
20. Samet, H.: *Foundations of Multidimensional and Metric Data Structures*. Morgan Kaufmann Publishers Inc. (2005)
21. Scholz, J., Stilman, M.: Combining motion planning and optimization for flexible robot manipulation. In: *Proceedings of IEEE-RAS International Conference on Humanoid Robots*, pp. 80–85 (2010)
22. Stoyan, D., Kendall, W.S., Mecke, J.: *Stochastic Geometry and Its Applications*. John Wiley & Sons (1996)
23. Sun, Z., Hsu, D., Jiang, T., Kurniawati, H., Reif, J.H.: Narrow passage sampling for probabilistic roadmap planners. *IEEE Transactions on Robotics* **21**(6), 1105–1115 (2005)