Faster Sample-based Motion Planning using Instance-based Learning

Jia Pan, Sachin Chitta*, Dinesh Manocha
Department of Computer Science, University of North Carolina at Chapel Hill
*Willow Garage

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 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.

Motivation

Most robots work in environments with small variations over time

Improve robot's performance by exploiting the knowledge learned from experiences

Three levels of Learning from Experience

low level data	Store, update and learn from sensor data	Low level, need expensive pre- and post- processing
collision queries	Learn an approximate description of obstacles in the configuration space	***Our work: flexible and powerful***
high level planning	Store and reuse the trajectories from planning episodes	High level, need advanced methods for trim and adjustment

Reduce the cost of collision detection

Exact Distance Computation » Exact Collision Computation » Probabilistic Collision Computation

Approximate collision checking with bounded error is enough in world with uncertainty

Overall Pipeline

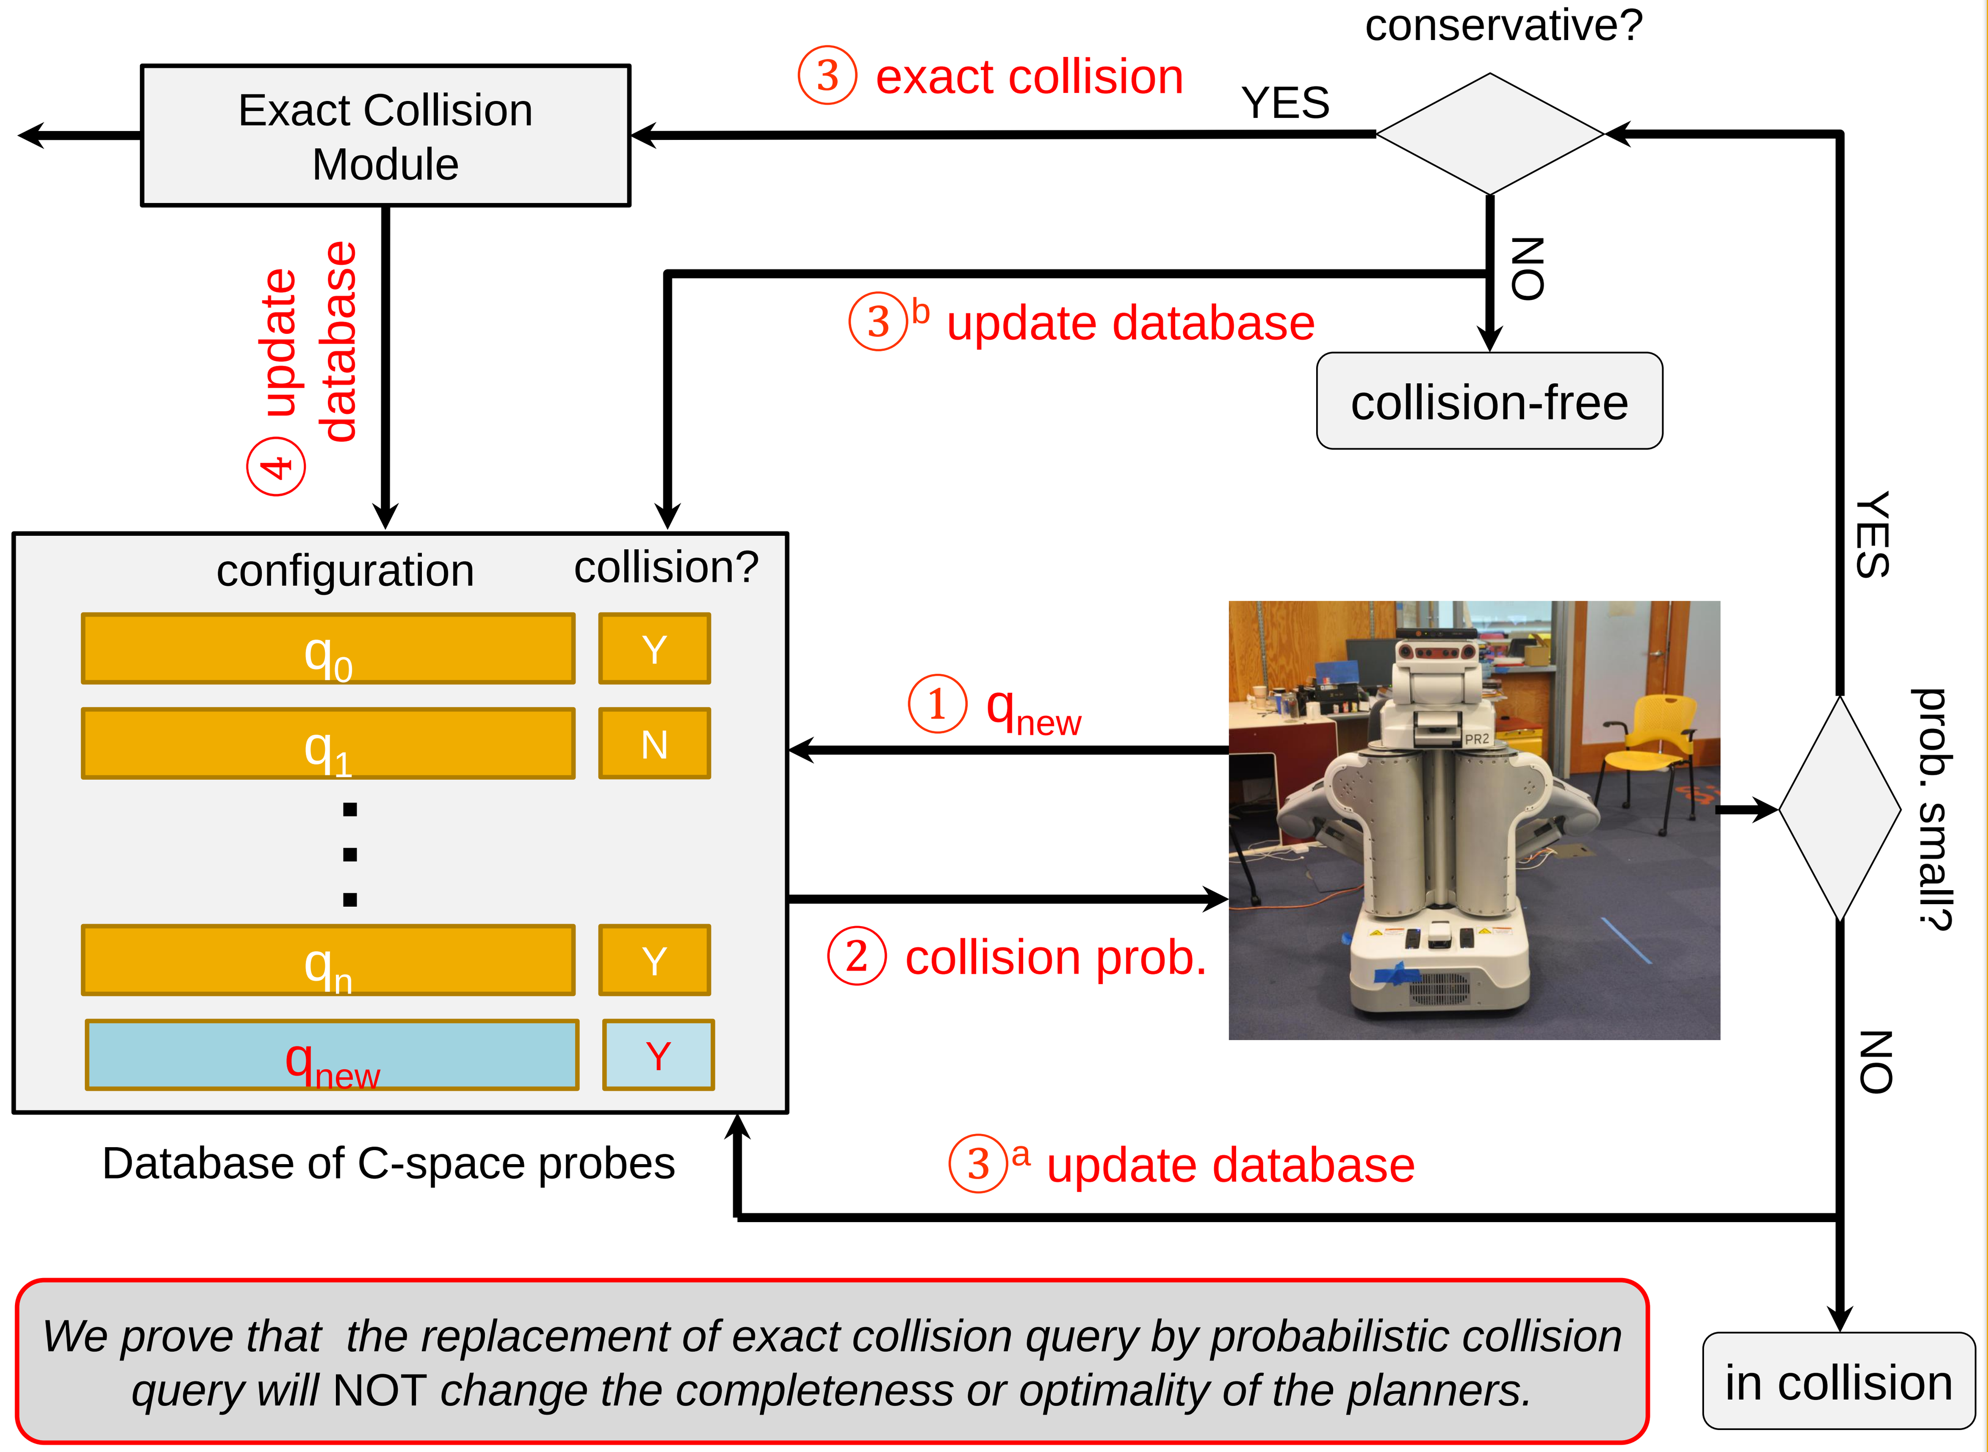

We prove that the replacement of exact collision query by probabilistic collision query will NOT *change the completeness or optimality of the planners.*

Probabilistic Collision Query

The collision probability computation is based on two kinds of k-nearest neighbor algorithms:
1) Find points that are nearest to a given point query;
2) Find points that are nearest to a given line query;
We use locality-sensitive hashing (LSH) based k-NN to provide almost *constant* query cost.

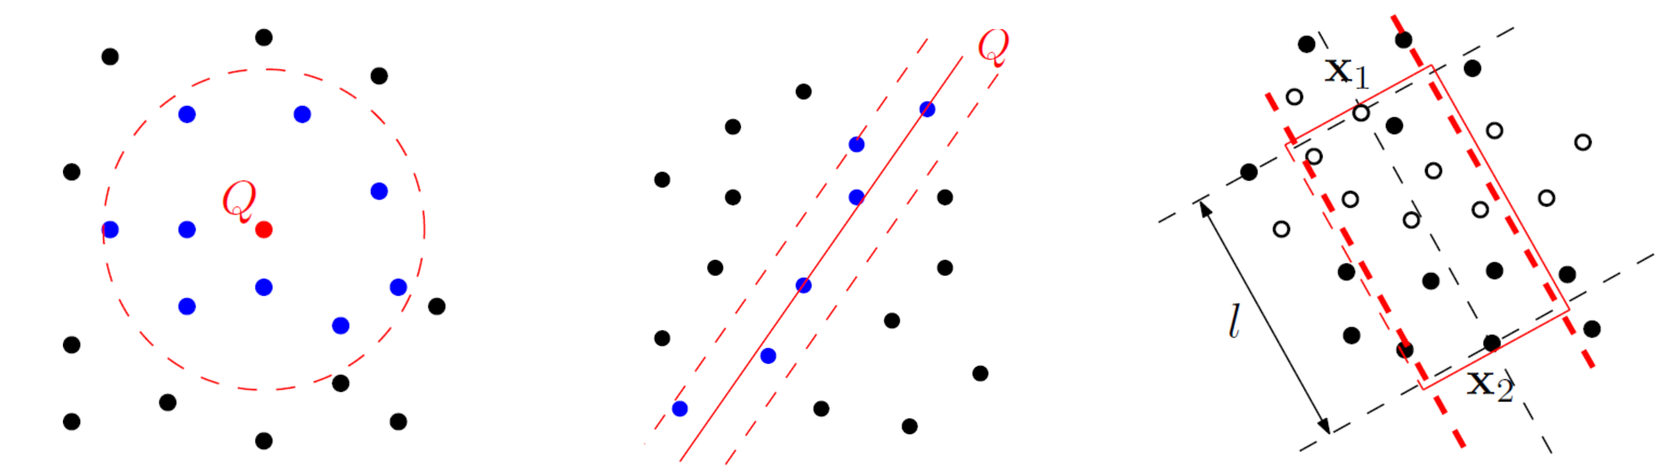

(a) (b) (c)

http://gamma.cs.unc.edu/IBL

Experiments

The instance-based learning algorithm is Integrated into OMPL as an alternative collision detection call. Our method is a general approach and does not change the sampling or planning strategies used by the underlying planners.
All motion planners can benefit it without any programming overhead!

Benchmarks

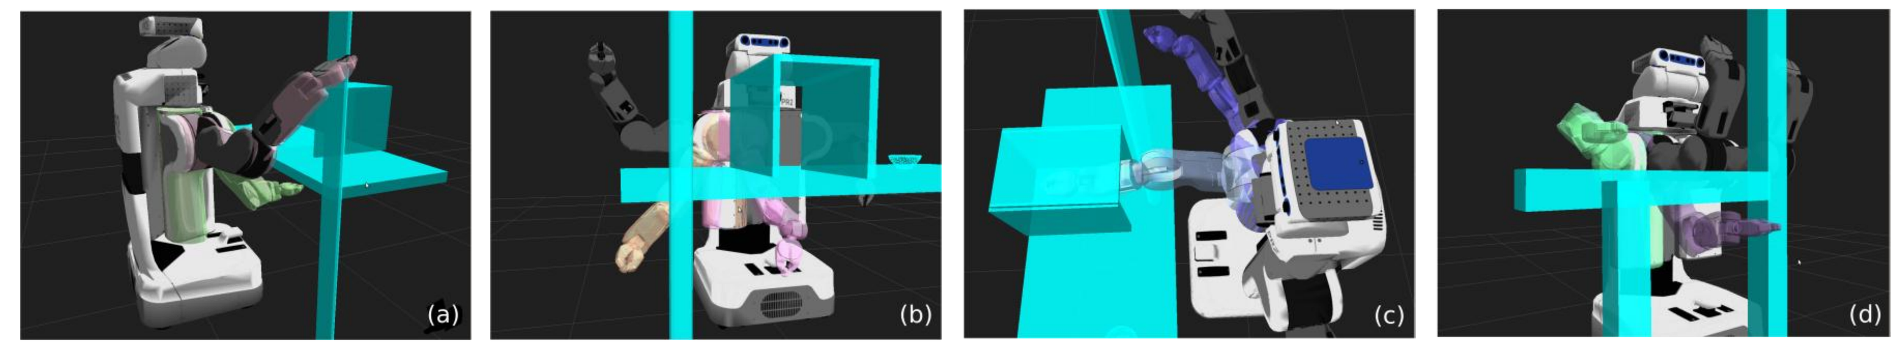

ROS planning wizard benchmarks

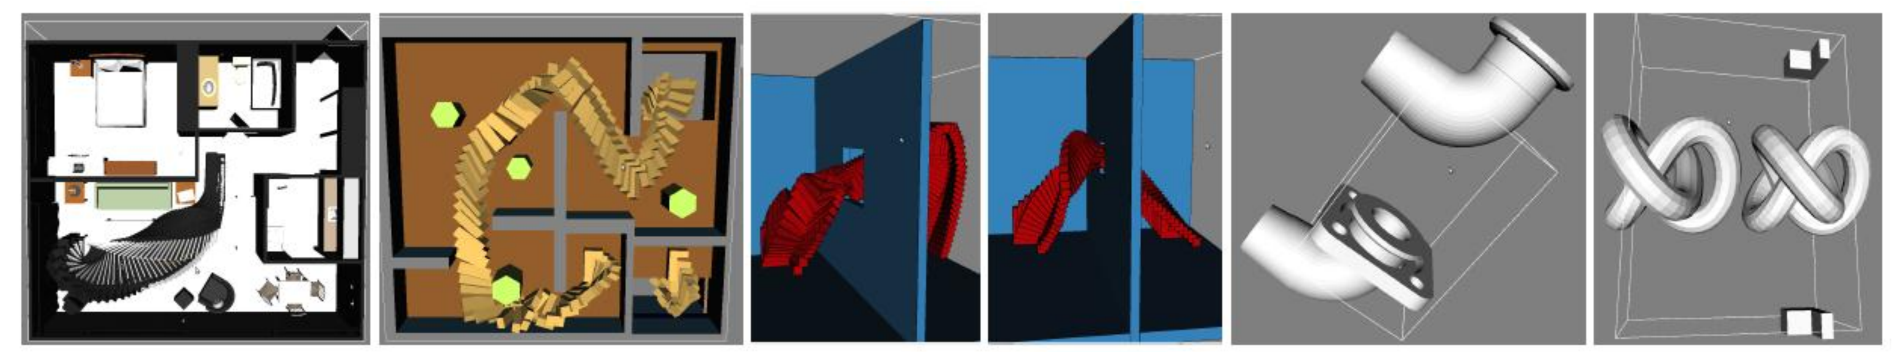

OMPL benchmarks

Results

1) Up to 100% acceleration on various planners.
2) The variation of the planning time is a bit smaller than original planner.
3)The robot is more efficient when it learns more and more:

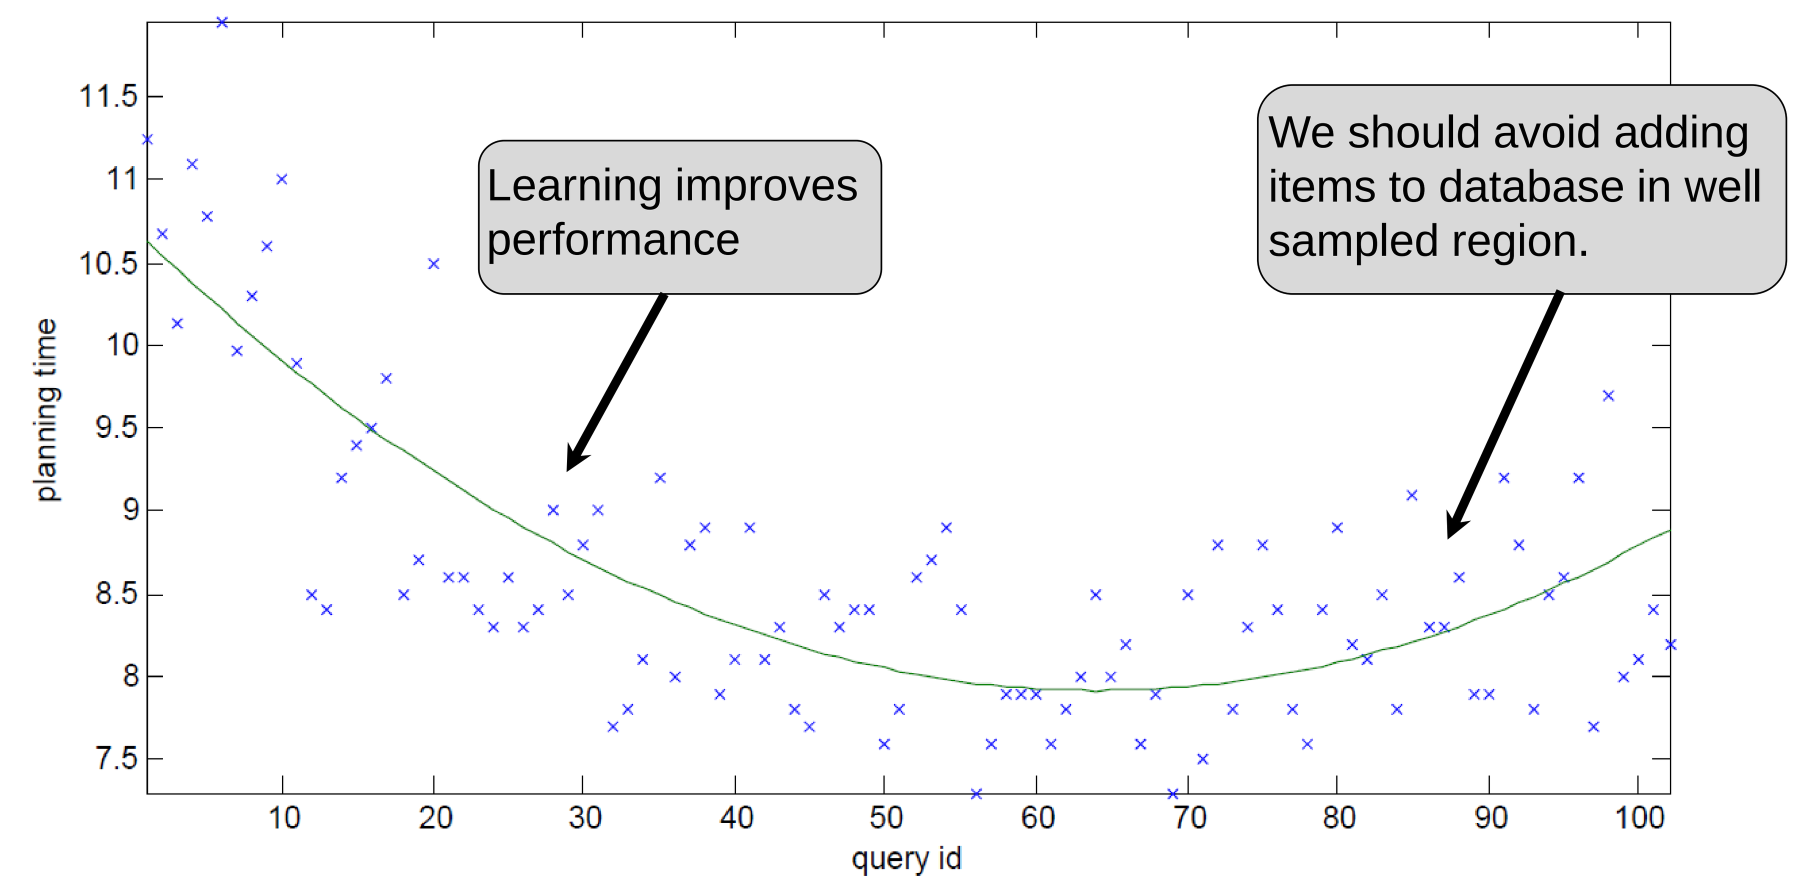

Acknowledgements: NSF and Willow Garage