

CSE-571
Sampling-Based Motion Planning

Slides based on those from Pieter Abbeel, Zoe McCarthy
Many images from Lavelle, Planning Algorithms

1

Solve by Nonlinear Optimization for Control?


- Could try by, for example, following formulation:

$$\begin{aligned} \min_{u,x} \quad & (x_T - x_G)^T (x_T - x_G) \\ \text{s.t.} \quad & x_{t+1} = f(x_t, u_t) \quad \forall t \\ & u_t \in \mathcal{U}_t \\ & x_t \in \mathcal{X}_t \quad \quad \quad \mathcal{X}_t \text{ can encode obstacles} \\ & x_0 = x_S \end{aligned}$$
- Or, with constraints, (which would require using an infeasible method):

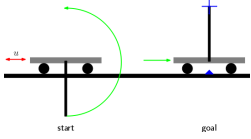
$$\begin{aligned} \min_{u,x} \quad & \|u\| \\ \text{s.t.} \quad & x_{t+1} = f(x_t, u_t) \quad \forall t \\ & u_t \in \mathcal{U}_t \\ & x_t \in \mathcal{X}_t \\ & x_0 = x_S \\ & x_T = x_G \end{aligned}$$
- *Can work surprisingly well, but for more complicated problems can get stuck in infeasible local minima*

2

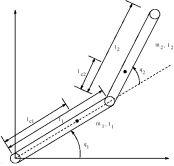
Examples



Helicopter path planning



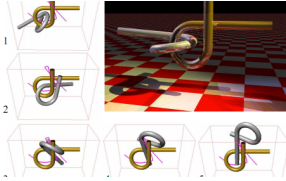
Cartpole swing-up

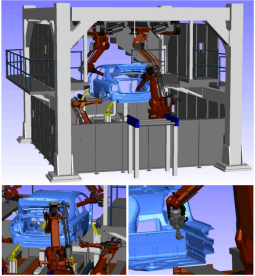


Acrobot


3

Examples

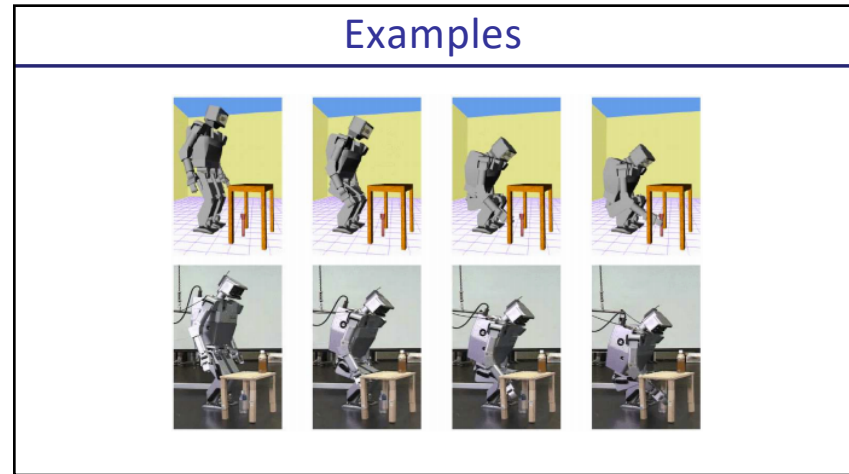




4



5



6

Motion Planning: Outline

- Configuration Space
- Probabilistic Roadmap
- Rapidly-exploring Random Trees (RRTs)
- Extensions
- Smoothing

7

Configuration Space (C-Space)

= { x | x is a pose of the robot }

- obstacles → configuration space obstacles


Workspace

(2 DOF: translation only, no rotation)

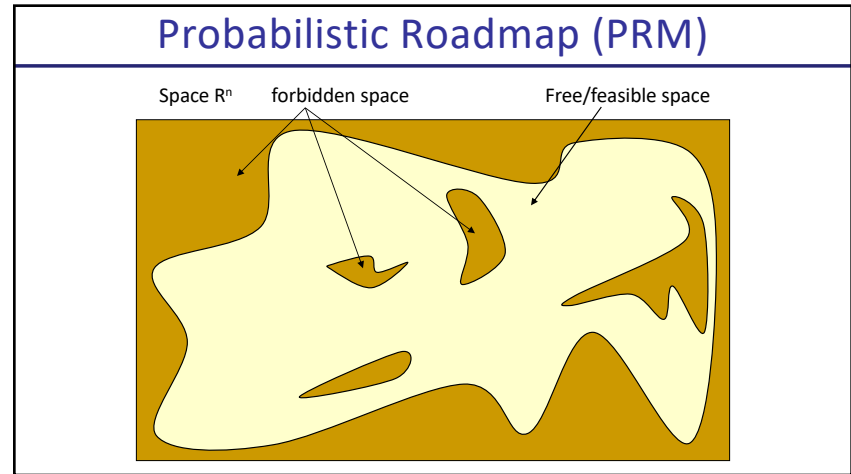
Configuration Space

free space □
obstacles ■

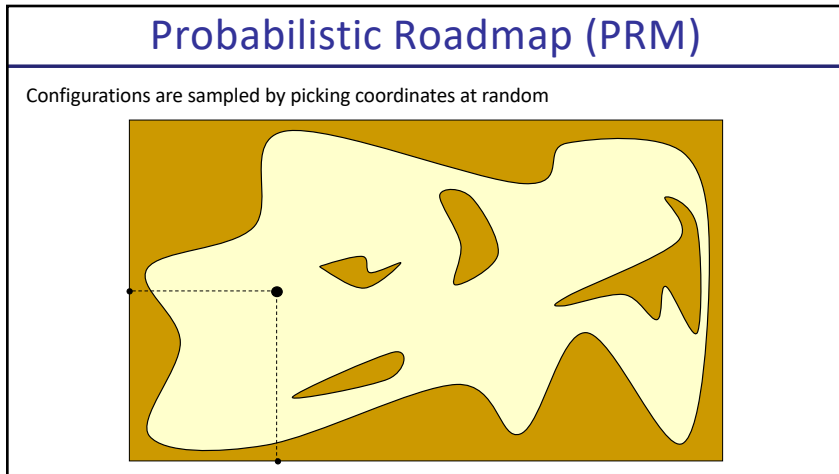
8



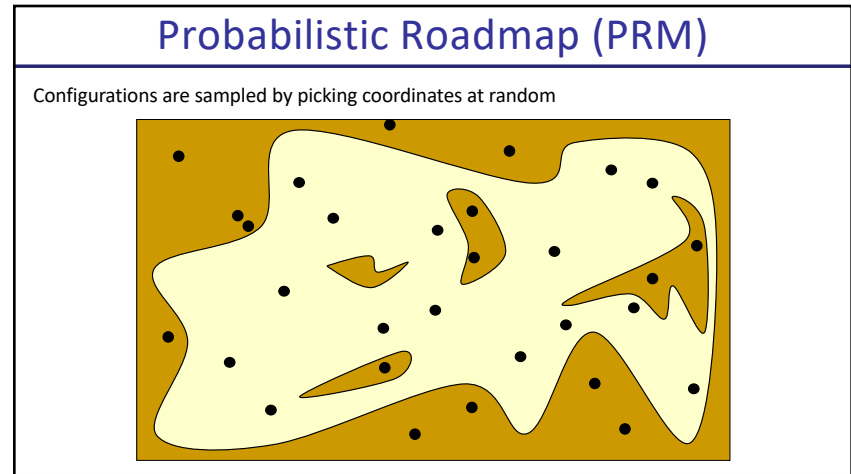
9



10




11



12

Probabilistic Roadmap (PRM)


Sampled configurations are tested for collision



13

Probabilistic Roadmap (PRM)


The collision-free configurations are retained as **milestones**



14

Probabilistic Roadmap (PRM)


Each milestone is linked by straight paths to its nearest neighbors



15

Probabilistic Roadmap (PRM)


Each milestone is linked by straight paths to its nearest neighbors



16

Probabilistic Roadmap (PRM)


The collision-free links are retained as **local paths** to form the PRM



17

Probabilistic Roadmap (PRM)


The start and goal configurations are included as milestones



18

Probabilistic Roadmap (PRM)

The PRM is searched for a path from s to g




19

Probabilistic Roadmap

- Initialize set of points with x_s and x_G
- Randomly sample points in configuration space
- Connect nearby points if they can be reached from each other
- Find path from x_s to x_G in the graph
- Alternatively: keep track of connected components incrementally, and declare success when x_s and x_G are in same connected component

20

PRM Example 1



21

PRM Example 2



22

PRM's Pros and Cons

- Pro:
 - Probabilistically complete: i.e., with probability one, if run for long enough the graph will contain a solution path if one exists.
- Cons:
 - Required to solve 2-point boundary value problem
 - Build graph over state space but no focus on generating a path

23


Rapidly exploring Random Tree (RRT)

Steve LaValle (98)

- Basic idea:
 - Build up a tree through generating "next states" in the tree by executing random controls
 - However: not exactly above to ensure good coverage

24

How to Sample



25

Rapidly exploring Random Tree (RRT)

- Select random point, and expand nearest vertex towards it
 - Biases samples towards largest Voronoi region

26

Rapidly exploring Random Tree (RRT)


```

GENERATE_RRT( $x_{init}, K, \Delta t$ )
1   $\mathcal{T}.init(x_{init});$ 
2  for  $k = 1$  to  $K$  do
3     $x_{rand} \leftarrow \text{RANDOM\_STATE}();$ 
4     $x_{near} \leftarrow \text{NEAREST\_NEIGHBOR}(x_{rand}, \mathcal{T});$ 
5     $u \leftarrow \text{SELECT\_INPUT}(x_{rand}, x_{near});$ 
6     $x_{new} \leftarrow \text{NEW\_STATE}(x_{near}, u, \Delta t);$ 
7     $\mathcal{T}.add\_vertex(x_{new});$ 
8     $\mathcal{T}.add\_edge(x_{near}, x_{new}, u);$ 
9  Return  $\mathcal{T}$ 
  
```

RANDOM_STATE(): often uniformly at random over space with probability 99%, and the goal state with probability 1%, this ensures it attempts to connect to goal semi-regularly

28

Rapidly exploring Random Tree (RRT)



Source: LaValle and Kuffner 01

29

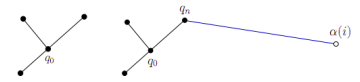
RRT Practicalities

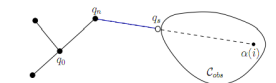
- NEAREST_NEIGHBOR(x_{rand}, T): need to find (approximate) nearest neighbor efficiently
 - KD Trees data structure (upto 20-D) [e.g., FLANN]
 - Locality Sensitive Hashing

- SELECT_INPUT(x_{rand}, x_{near})
 - Two point boundary value problem
 - If too hard to solve, often just select best out of a set of control sequences. This set could be random, or some well chosen set of primitives.

30

RRT Extension


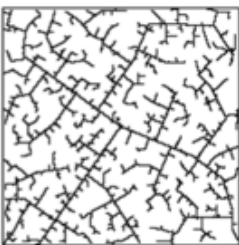
- No obstacles, holonomic:
 

- With obstacles, holonomic:
 

- Non-holonomic: approximately (sometimes as approximate as picking best of a few random control sequences) solve two-point boundary value problem

31

Growing RRT

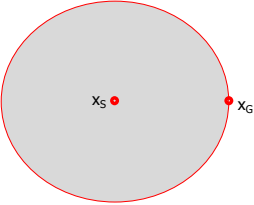



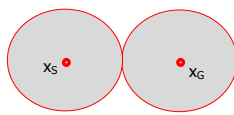
45 iterations 390 iterations

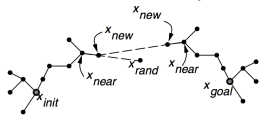
Demo: [http://en.wikipedia.org/wiki/File:Rapidly-exploring_Random_Tree_\(RRT\)_500x373.gif](http://en.wikipedia.org/wiki/File:Rapidly-exploring_Random_Tree_(RRT)_500x373.gif)

32

Bi-directional RRT

- Volume swept out by unidirectional RRT:


- Volume swept out by bi-directional RRT:





- Difference more and more pronounced as dimensionality increases

33

Multi-directional RRT


- Planning around obstacles or through narrow passages can often be easier in one direction than the other



34

Resolution-Complete RRT (RC-RRT)

- Issue: nearest points chosen for expansion are (too) often the ones stuck behind an obstacle



RC-RRT solution:

- Choose a maximum number of times, m , you are willing to try to expand each node
- For each node in the tree, keep track of its Constraint Violation Frequency (CVF)
- Initialize CVF to zero when node is added to tree
- Whenever an expansion from the node is unsuccessful (e.g., per hitting an obstacle):
 - Increase CVF of that node by 1
 - Increase CVF of its parent node by $1/m$, its grandparent $1/m^2, \dots$
- When a node is selected for expansion, skip over it with probability CVF/m

35

RRT*

```

Algorithm 6: RRT*
1  $V \leftarrow \{x_{init}\}; E \leftarrow \emptyset;$ 
2 for  $i = 1, \dots, n$  do
3    $x_{rand} \leftarrow \text{SampleFree};$ 
4    $x_{nearest} \leftarrow \text{Nearest}(G = (V, E), x_{rand});$ 
5    $x_{new} \leftarrow \text{Steer}(x_{nearest}, x_{rand});$ 
6   if  $\text{ObstacleFree}(x_{nearest}, x_{new})$  then
7      $x_{near} \leftarrow \text{Near}(G = (V, E), x_{new}, \min\{\gamma_{RRT^*}(\log(\text{card}(V)) / \text{card}(V))^{1/d}, \eta\});$ 
8      $V \leftarrow V \cup \{x_{new}\};$ 
9      $x_{min} \leftarrow x_{nearest}; c_{min} \leftarrow \text{Cost}(x_{nearest}) + c(\text{Line}(x_{nearest}, x_{new}));$ 
10    foreach  $x_{near} \in X_{near}$  do // Connect along a minimum-cost path
11      if  $\text{CollisionFree}(x_{near}, x_{new}) \wedge \text{Cost}(x_{near}) + c(\text{Line}(x_{near}, x_{new})) < c_{min}$  then
12         $x_{min} \leftarrow x_{near}; c_{min} \leftarrow \text{Cost}(x_{near}) + c(\text{Line}(x_{near}, x_{new}))$ 
13     $E \leftarrow E \cup \{(x_{min}, x_{new})\};$ 
14    foreach  $x_{near} \in X_{near}$  do // Rewire the tree
15      if  $\text{CollisionFree}(x_{new}, x_{near}) \wedge \text{Cost}(x_{new}) + c(\text{Line}(x_{new}, x_{near})) < \text{Cost}(x_{near})$ 
16        then  $x_{parent} \leftarrow \text{Parent}(x_{near});$ 
17         $E \leftarrow (E \setminus \{(x_{parent}, x_{near})\}) \cup \{(x_{new}, x_{near})\};$ 
18 return  $G = (V, E);$ 
    
```

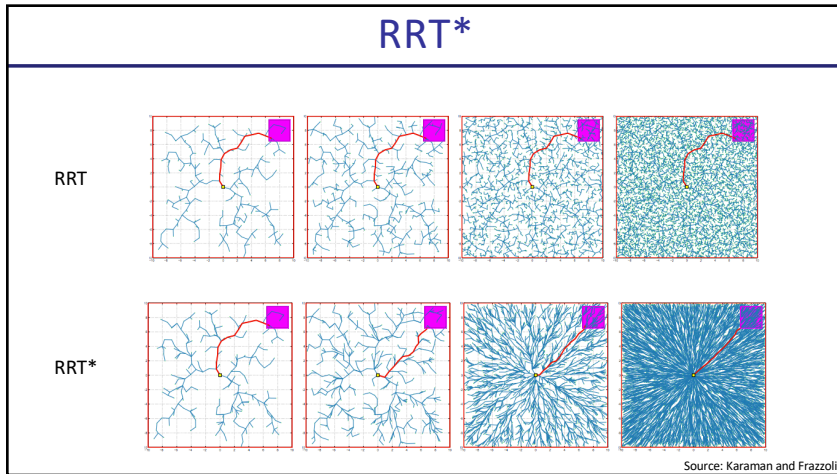
Source: Karaman and Frazzoli

36

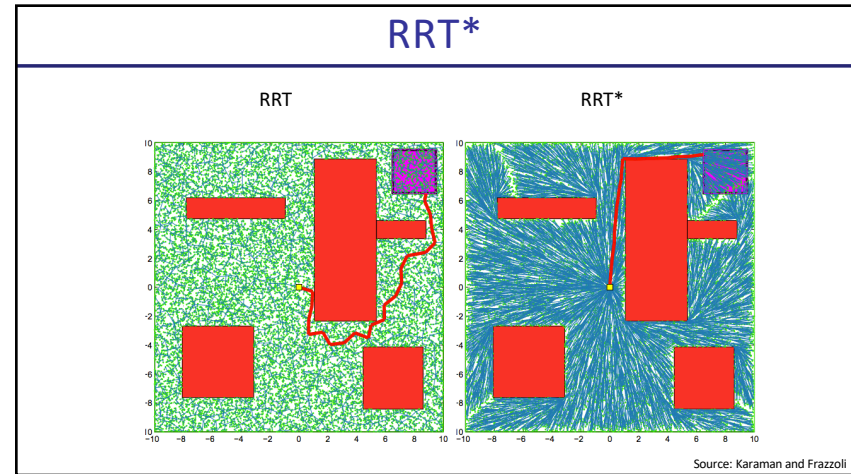
RRT*

- Asymptotically optimal
- Main idea:
 - Swap new point in as parent for nearby vertices who can be reached along shorter path through new point than through their original (current) parent


37



38



39



40

Smoothing

Randomized motion planners tend to find not so great paths for execution: very jagged, often much longer than necessary.

→ In practice: do smoothing before using the path

- **Shortcutting:**
 - along the found path, pick two vertices x_{t1} , x_{t2} and try to connect them directly (skipping over all intermediate vertices)
- **Nonlinear optimization for optimal control**
 - Allows to specify an objective function that includes smoothness in state, control, small control inputs, etc.

41

Additional Resources

- Marco Pavone (<http://asl.stanford.edu/>):
 - Sampling-based motion planning on GPUs: <https://arxiv.org/pdf/1705.02403.pdf>
 - Learning sampling distributions: <https://arxiv.org/pdf/1709.05448.pdf>
- Sidd Srinivasa (<https://personalrobotics.cs.washington.edu/>):
 - Batch informed trees: <https://robotic-esp.com/code/bitstar/>
 - Expensive edge evals: <https://arxiv.org/pdf/2002.11853.pdf>
- Michael Yip (<https://www.ucsdarclab.com/>):
 - Neural Motion Planners: <https://www.ucsdarclab.com/neuralplanning>
- Lydia Kavraki (<http://www.kavrakilab.org/>):
 - Motion in human workspaces: <http://www.kavrakilab.org/nsf-nri-1317849.html>