

ECE276B: Planning & Learning in Robotics

Lecture 8: Sampling-based Planning

Lecturer:

Nikolay Atanasov: natanasov@ucsd.edu

Teaching Assistants:

Tianyu Wang: tiw161@eng.ucsd.edu

Yongxi Lu: yol070@eng.ucsd.edu



UC San Diego

JACOBS SCHOOL OF ENGINEERING

Electrical and Computer Engineering



Search-based vs Sampling-based Planning

- ▶ Search-based planning:
 - ▶ Generates a systematic discrete representation (graph) of C_{free}
 - ▶ Searches the representation for a path guaranteeing to find one if it exists (resolution complete)
 - ▶ Can interleave the representation construction with the search, i.e., adds nodes only when necessary
 - ▶ Provides suboptimality bounds on the solution
 - ▶ Can get computationally expensive in high dimensions



Search-based vs. Sampling-based Planning

- ▶ Sampling-based planning:
 - ▶ Generates a sparse sample-based representation (graph) of C_{free}
 - ▶ Searches the representation for a path guaranteeing that the probability of finding one if it exists approaches 1 as the number of iterations $\rightarrow \infty$ (probabilistically complete)
 - ▶ Can interleave the representation construction with the search, i.e., adds samples only when necessary
 - ▶ Provides asymptotic suboptimality bounds on the solution
 - ▶ Well-suited for high-dimensional planning as it is faster and requires less memory than search-based planning in many domains



Probabilistic Roadmap (PRM)

Step 1. Preprocessing Phase: Build a roadmap (graph) \mathcal{G} which, hopefully, should be accessible from any point in C_{free}

- ▶ **Nodes:** randomly sampled valid configurations $x_i \in C_{free}$
- ▶ **Edges:** added between samples that are easy to connect with a simple local controller (e.g., follow straight line)



Step 2. Query Phase: Given a start configuration x_s and goal configuration x_T , connect them to the roadmap \mathcal{G} using a local planner, then search the augmented roadmap for a shortest path from x_s to x_T


▶ **Pros and Cons:**

- ▶ Simple and highly effective in high dimensions
- ▶ Can result in suboptimal paths, no guarantees on suboptimality
- ▶ Difficulty with narrow passages
- ▶ Useful for multiple queries with different start and goal in the same environment

Step 1: Preprocessing Phase

Algorithm 1 Build Roadmap

```
1:  $\mathcal{G}.init()$ 
2: for  $i = 1, \dots, N$  do
3:   Sample  $x_{rand}$ 
4:   if  $x_{rand} \in \mathcal{C}_{free}$  then ▷ New sample
5:      $\mathcal{G}.add\_vertex(x_{rand})$ 
6:     for  $x \in \text{NEIGHBORHOOD}(x_{rand}, \mathcal{G})$  do ▷ Region around sample
7:       if (not  $\mathcal{G}.same\_component(x_{rand}, x)$ ) and  $\text{CONNECT}(x_{rand}, x)$  then
8:          $\mathcal{G}.add\_edge(x_{rand}, x)$  ▷ Can be connected by local planner
```



Step 1: Preprocessing Phase


- ▶ Efficient implementation of $x \in \text{NEIGHBORHOOD}(x_{rand}, \mathcal{G})$:
 - ▶ select all nodes within a fixed radius from x_{rand}
 - ▶ select K nodes closest to x_{rand}
 - ▶ select K (often just 1) closest points from each of the components in \mathcal{G}
- ▶ $\mathcal{G}.\text{same_component}(x_{rand}, x)$ may be replaced by “ $|\text{Children}(x)| < K$ ”
- ▶ Sampling strategies:
 - ▶ Sample x_{rand} uniformly from C_{free}
 - ▶ Select an existing nodes with probability inversely proportional to how well connected it is and generate a random motion from it to get x_{rand}
 - ▶ Bias sampling towards obstacle boundaries
 - ▶ Bias sampling away from obstacles

PRM vs RRT


- ▶ **Rapidly Exploring Random Tree (RRT):**
 - ▶ One of the most popular planning techniques
 - ▶ Introduced by Steven LaValle in 1998
 - ▶ Many, many, many extensions and variants (articulated robots, kinematics, dynamics, differential constraints)
- ▶ **PRM:** a graph constructed from random samples. It can be search for a path whenever a start node x_s and goal node x_T are specified. PRMs are well-suited for repeated planning between different pairs of x_s and x_T (*multiple queries*)
- ▶ **RRT:** a tree is constructed from random samples with root x_s . The tree is grown until it contains a path to x_T . RRTs are well-suited for single-shot planning between a single pair of x_s and x_T (*single query*)
- ▶ There exist extensions of RRTs that try to reuse a previously constructed tree when replanning in response to map updates

Rapidly Exploring Random Tree (RRT)


- ▶ Sample a new configuration x_{rand} , find the nearest neighbor x_{near} in \mathcal{G} and connect them:



- ▶ If the nearest point x_{near} lies on an existing edge, then split the edge:



- ▶ If there is an obstacle, the edge travels up to the obstacle boundary, as far as allowed by a collision detection algorithm



Rapidly Exploring Random Tree (RRT)


- ▶ What about the goal? Occasionally (e.g., every 100 iterations) add the goal configuration x_T and see if it gets connected to the tree
- ▶ RRT can be implemented in the original workspace (need to do collision checking) or in configuration space
- ▶ Challenges with a C-Space implementation:
 - ▶ What distance function do we use to find the nearest configuration?
 - ▶ e.g., distance along the surface of a torus for a 2 link manipulator
 - ▶ An edge represents a path in C-Space. How do we construct a collision-free path between two configurations?
 - ▶ We do not have to connect the configurations all the way. Instead, use a small step size ϵ and a local steering function to get closer to the second configuration.

Rapidly Exploring Random Tree (RRT)

- ▶ **No preprocessing:** starting with an initial configuration x_s build a graph (actually, tree) until the goal configuration x_T is part of it

Algorithm 2 Build_RRT(x_s)

- 1: $\mathcal{T}.\text{init}(x_s)$
 - 2: **for** $i = 1 \dots N$ **do**
 - 3: Sample x_{rand}
 - 4: EXTEND(\mathcal{T}, x_{rand})
 - 5: **return** \mathcal{T}
-



Algorithm 3 Extend(\mathcal{T}, x_{rand})

- 1: $x_{near} \leftarrow \text{NEARESTNEIGHBOR}(\mathcal{T}, x_{rand})$
- 2: $x_{new} \leftarrow \text{STEER}(x_{near}, x_{rand})$
- 3: **if** $\text{OBSTACLEFREE}(x_{near}, x_{new})$ **then**
- 4: $\mathcal{T}.\text{add_vertex}(x_{new})$
- 5: $\mathcal{T}.\text{add_edge}(x_{near}, x_{new})$
- 6: **if** $x_{new} = x_{rand}$ **then**
- 7: **return** Reached
- 8: **else**
- 9: **return** Advanced
- 10: **return** Trapped


- ▶ closest node in the tree
- ▶ moves by at most ϵ from x_{near} towards x_{rand}

Rapidly Exploring Random Tree (RRT)

- ▶ RRT without ϵ (called Rapidly Exploring Dense Tree (RDT)):






45 iterations




2345 iterations

- ▶ RRT with ϵ




Example: RRT Algorithm

- ▶ Start node x_s
- ▶ Goal node x_t
- ▶ Gray obstacles




Example: RRT Algorithm

- ▶ Sample x_{rand} in the workspace
- ▶ Steer from x_s towards x_{rand} by a fixed distance ϵ to get x_1
- ▶ If the segment from x_s to x_1 is collision-free, insert x_1 into the tree



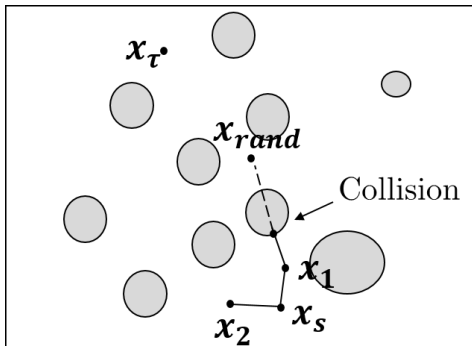
Example: RRT Algorithm

- ▶ Sample x_{rand} in the workspace
- ▶ Find the closest node x_{near} to x_{rand}
- ▶ Steer from x_{near} towards x_{rand} by a fixed distance ϵ to get x_2
- ▶ If the segment from x_{near} to x_2 is collision-free, insert x_2 into the tree




Example: RRT Algorithm

- ▶ Sample x_{rand} in the workspace
- ▶ Find the closest node x_{near} to x_{rand}
- ▶ Steer from x_{near} towards x_{rand} by a fixed distance ϵ to get x_3
- ▶ If the segment from x_{near} to x_3 is collision-free, insert x_3 into the tree




Example: RRT Algorithm

- ▶ Sample x_{rand} in the workspace
- ▶ Find the closest node x_{near} to x_{rand}
- ▶ Steer from x_{near} towards x_{rand} by a fixed distance ϵ to get x_3
- ▶ If the segment from x_{near} to x_3 is collision-free, insert x_3 into the tree






Example: RRT Algorithm

- ▶ Continue until a node that is a distance ϵ from the goal is generated
- ▶ Either terminate the algorithm or search for additional feasible paths






Sampling in RRTs

- ▶ The vanilla RRT algorithm provides uniform coverage of space




- ▶ Alternatively, the growth may be biased by the largest Voronoi region




Sampling in RRTs


- ▶ Goal-biased sampling: with probability $(1 - p_g)$, x_{rand} is chosen as a uniform sample in C_{free} and with probability p_g , $x_{rand} = x_T$



(a) $p_g = 0$



(b) $p_g = 0.1$







(c) $p_g = 0.5$

Handling Robot Dynamics with Steer()

- ▶ Steer() extends the tree towards a given random sample x_{rand}
- ▶ Consider a car-like robot with non-holonomic constraints (can't slide sideways) in $SE(2)$. Obtaining a feasible path from $x_{rand} = (0, 0, 90^\circ)$ to $x_{near} = (1, 0, 90^\circ)$ is as hard as the original problem
- ▶ Steer() resolves this by not requiring the motion to get all the way to x_{rand} . We just apply the best control input for a fixed duration to obtain x_{new} and a dynamically feasible trajectory to it

Bug Traps

- ▶ Growing two trees, one from start and one for goal, often has better performance in practice.



Bi-directional RRT

Algorithm 4 BALANCED_BIDIRECTIONAL_RRT(x_s, x_T)

```
1:  $\mathcal{T}_a.\text{init}(x_s); \mathcal{T}_b.\text{init}(x_T);$ 
2: for  $i = 1 \dots N$  do
3:   Sample  $x_{rand}$ 
4:    $x_{near} \leftarrow \text{NEARESTNEIGHBOR}(\mathcal{T}_a, x_{rand})$ 
5:    $x_c \leftarrow \text{STEER}(x_{near}, x_{rand})$ 
6:   if  $x_c \neq x_{near}$  then
7:      $\mathcal{T}_a.\text{add\_vertex}(x_c)$ 
8:      $\mathcal{T}_a.\text{add\_edge}(x_{near}, x_c)$ 
9:      $x'_{near} \leftarrow \text{NEARESTNEIGHBOR}(\mathcal{T}_b, x_c)$ 
10:     $x'_c \leftarrow \text{STEER}(x'_{near}, x_c)$ 
11:    if  $x'_c \neq x'_{near}$  then
12:       $\mathcal{T}_b.\text{add\_vertex}(x'_c)$ 
13:       $\mathcal{T}_b.\text{add\_edge}(x'_{near}, x'_c)$ 
14:    if  $x'_c = x_c$  then return SOLUTION
15:    if  $|\mathcal{T}_b| < |\mathcal{T}_a|$  then SWAP( $\mathcal{T}_a, \mathcal{T}_b$ )
16: FAILURE
```

RRT-Connect



- ▶ J. Kuffner and S. LaValle, “RRT-Connect: An Efficient Approach to Single-Query Path Planning,” ICRA’00
- ▶ Bi-directional tree + relax the ϵ constraint on tree growth

Algorithm 5 RRT_CONNECT(x_s, x_t)

```
1:  $\mathcal{T}_a$ .init( $x_s$ );  $\mathcal{T}_b$ .init( $x_t$ );
2: for  $k = 1 \dots K$  do
3:   Sample  $x_{rand}$ 
4:   if not EXTEND( $\mathcal{T}_a, x_{rand}$ ) = Trapped then
5:     if CONNECT( $\mathcal{T}_b, x_{new}$ ) = Reached then
6:       return PATH( $\mathcal{T}_a, \mathcal{T}_b$ )
7:   SWAP( $\mathcal{T}_a, \mathcal{T}_b$ )
8: return Failure
9:
10: function CONNECT( $\mathcal{T}, x$ )
11:   repeat
12:      $S \leftarrow$  EXTEND( $\mathcal{T}, x$ )
13:   until not ( $S =$  Advanced)
14:   return  $S$ 
```



▷ x_{new} was just added to \mathcal{T}_a

Example: Single RRT-Connect Iteration





Example: Single RRT-Connect Iteration

- ▶ One tree is grown to a random target




Example: Single RRT-Connect Iteration

- ▶ The new node becomes a target for the other tree




Example: Single RRT-Connect Iteration

- ▶ Determine the nearest node to the target




Example: Single RRT-Connect Iteration

- ▶ Try to add a new collision-free branch




Example: Single RRT-Connect Iteration

- ▶ If successful, keep extending the branch




Example: Single RRT-Connect Iteration

- ▶ If successful, keep extending the branch




Example: Single RRT-Connect Iteration

- ▶ If successful, keep extending the branch




Example: Single RRT-Connect Iteration

- ▶ If the branch reaches all the way to the target, a feasible path is found!





Example: Single RRT-Connect Iteration



- ▶ If the branch reaches all the way to the target, a feasible path is found!





Example: RRT-Connect



Example: RRT-Connect



Example: RRT-Connect




Why are RRTs so popular?

- ▶ The algorithm is very simple once the following subroutines are implemented:
 - ▶ Random sample generator
 - ▶ Nearest neighbor
 - ▶ Collision checker
 - ▶ Steer
- ▶ Pros:
 - ▶ Sparse exploration requires little memory and computation
 - ▶ RRTs find feasible paths quickly in practice
 - ▶ Can add heuristics on top, e.g., bias the sampling towards the goal
- ▶ Cons:
 - ▶ Solutions can be highly sub-optimal and require path smoothing as a post-processing step
 - ▶ The smoothed path is still restricted to the same homotopy class



Path Smoothing

- ▶ Start with the initial point (1)
- ▶ Make connections to subsequent points in the path (2), (3), (4), ...
- ▶ When a connection collides with obstacles, add the previous waypoint to the smoothed path
- ▶ Continue smoothing from this point on



Search-based vs Sampling-based Planning

- ▶ RRT:
 - ▶ Sparse exploration requires little memory and computation
 - ▶ Solutions can be highly sub-optimal and require post-processing (path smoothing) which may be difficult
- ▶ Weighted A*:
 - ▶ Systematic exploration may require a lot of memory and computation
 - ▶ Returns a path with (sub-)optimality guarantees



RRT Guarantees

- ▶ RRT and RRT-Connect are **probabilistically complete**: the probability that a feasible path will be found if one exists, approaches 1 exponentially as the number of samples approaches infinity
- ▶ Assuming C_{free} is connected, bounded, and open, for any $x \in C_{free}$,
 $\lim_{N \rightarrow \infty} \mathbb{P}(\|x - x_{near}\| < \epsilon) = 1$, where x_{near} is the closest node to x in \mathcal{T}
- ▶ RRT is **not optimal**: the probability that RRT converges to an optimal solution, as the number of samples approaches infinity, is zero under reasonable technical assumptions (S. Karaman, E. Frazzoli, RSS'10)
- ▶ **Problem**: once we build an RRT we never modify it

RRT*

- ▶ S. Karaman, E. Frazzoli, “Incremental Sampling-based Algorithms for Optimal Motion Planning,” RSS’10
- ▶ RRT*: RRT + rewiring of the tree to ensure asymptotic optimality

Algorithm 1: Body of RRT and RRG Algorithms

```
1  $V \leftarrow \{x_{init}\}; E \leftarrow \emptyset; i \leftarrow 0;$   
2 while  $i < N$  do  
3    $G \leftarrow (V, E);$   
4    $x_{rand} \leftarrow \text{Sample}(i); i \leftarrow i + 1;$   
5    $(V, E) \leftarrow \text{Extend}(G, x_{rand});$ 
```

Algorithm 2: Extend_{RRT}


```
1  $V' \leftarrow V; E' \leftarrow E;$   
2  $x_{nearest} \leftarrow \text{Nearest}(G, x);$   
3  $x_{new} \leftarrow \text{Steer}(x_{nearest}, x);$   
4 if  $\text{ObstacleFree}(x_{nearest}, x_{new})$  then  
5    $V' \leftarrow V' \cup \{x_{new}\};$   
6    $E' \leftarrow E' \cup \{(x_{nearest}, x_{new})\};$   
7 return  $G' = (V', E')$ 
```

Algorithm 4: Extend_{RRT*}

```
1  $V' \leftarrow V; E' \leftarrow E;$   
2  $x_{nearest} \leftarrow \text{Nearest}(G, x);$   
3  $x_{new} \leftarrow \text{Steer}(x_{nearest}, x);$   
4 if  $\text{ObstacleFree}(x_{nearest}, x_{new})$  then  
5    $V' \leftarrow V' \cup \{x_{new}\};$   
6    $x_{min} \leftarrow x_{nearest};$   
7    $X_{near} \leftarrow \text{Near}(G, x_{new}, |V|);$   
8   for all  $x_{near} \in X_{near}$  do  
9     if  $\text{ObstacleFree}(x_{near}, x_{new})$  then  
10       $c' \leftarrow \text{Cost}(x_{near}) + c(\text{Line}(x_{near}, x_{new}));$   
11      if  $c' < \text{Cost}(x_{new})$  then  
12         $x_{min} \leftarrow x_{near};$   
13    $E' \leftarrow E' \cup \{(x_{min}, x_{new})\};$   
14   for all  $x_{near} \in X_{near} \setminus \{x_{min}\}$  do  
15     if  $\text{ObstacleFree}(x_{new}, x_{near})$  and  
16      $\text{Cost}(x_{near}) >$   
17      $\text{Cost}(x_{new}) + c(\text{Line}(x_{new}, x_{near}))$  then  
18      $x_{parent} \leftarrow \text{Parent}(x_{near});$   
19      $E' \leftarrow E' \setminus \{(x_{parent}, x_{near})\};$   
20      $E' \leftarrow E' \cup \{(x_{new}, x_{near})\};$   
21 return  $G' = (V', E')$ 
```



RRT*: Extend Step

- ▶ Generate a new potential node x_{new} identically to RRT
- ▶ Instead of finding the closest node in the tree, find all nodes within a neighborhood \mathcal{N}
- ▶ Let $x_{nearest} = \arg \min_{x_{near} \in \mathcal{N}} g_{x_{near}} + c_{x_{near}, x_{new}}$, i.e., the node in \mathcal{N} that lies on the currently known shortest path from x_s to x_{new}
- ▶ Add node: $\mathcal{V} \leftarrow \mathcal{V} \cup \{x_{new}\}$
- ▶ Add edge: $\mathcal{E} \leftarrow \mathcal{E} \cup \{(x_{nearest}, x_{new})\}$
- ▶ Set the label of x_{new} to $g_{x_{new}} = g_{x_{nearest}} + c_{x_{nearest}, x_{new}}$




RRT*: Rewire Step


- ▶ Check all nodes $x_{near} \in \mathcal{N}$ to see if re-routing through x_{new} reduces the path length (**label correcting!**):
- ▶ If $g_{x_{new}} + c_{x_{new}, x_{near}} < g_{x_{near}}$, then remove the edge between x_{near} and its parent and add a new edge between x_{near} and x_{new}



RRT vs RRT*




(a) RRT




(b) RRT*

- ▶ Same nodes in the tree, only the edge connections are different. Notice how the RRT* edges are almost straight lines (optimal paths).
- ▶ S. Karaman and E. Frazzoli, "Incremental Sampling-based Algorithms for Optimal Motion Planning," International Journal of Robotics Research, 2010.

RRT vs RRT*



(a) RRT



(b) RRT*

- ▶ Same nodes in the tree, only the edge connections are different. Notice how the RRT* edges are almost straight lines (optimal paths).
- ▶ S. Karaman and E. Frazzoli, "Incremental Sampling-based Algorithms for Optimal Motion Planning," International Journal of Robotics Research, 2010.