

CS 287 Advanced Robotics (Fall 2019)

Lecture 9: Motion Planning

Lecture by: Huazhe (Harry) Xu

Slides by: Pieter Abbeel

UC Berkeley EECS

Motion Planning


- **Problem**

- Given start state X_S , goal state X_G
- Asked for: a sequence of control inputs that leads from start to goal


- **Why tricky?**

- Need to avoid obstacles
- For systems with underactuated dynamics: can't simply move along any coordinate at will
 - E.g., car, helicopter, airplane, but also robot manipulator hitting joint limits


Examples



Helicopter path
planning





Cartpole swing-up





Acrobot


Examples



Examples



Examples





Motion Planning: Outline

- Configuration Space
- Optimization-based Motion Planning
- Sampling-based Motion Planning
 - Probabilistic Roadmap
 - Rapidly-exploring Random Trees (RRTs)
 - Smoothing

Motion Planning: Outline

- ***Configuration Space***
- Optimization-based Motion Planning
- Sampling-based Motion Planning
 - Probabilistic Roadmap
 - Rapidly-exploring Random Trees (RRTs)
 - Smoothing

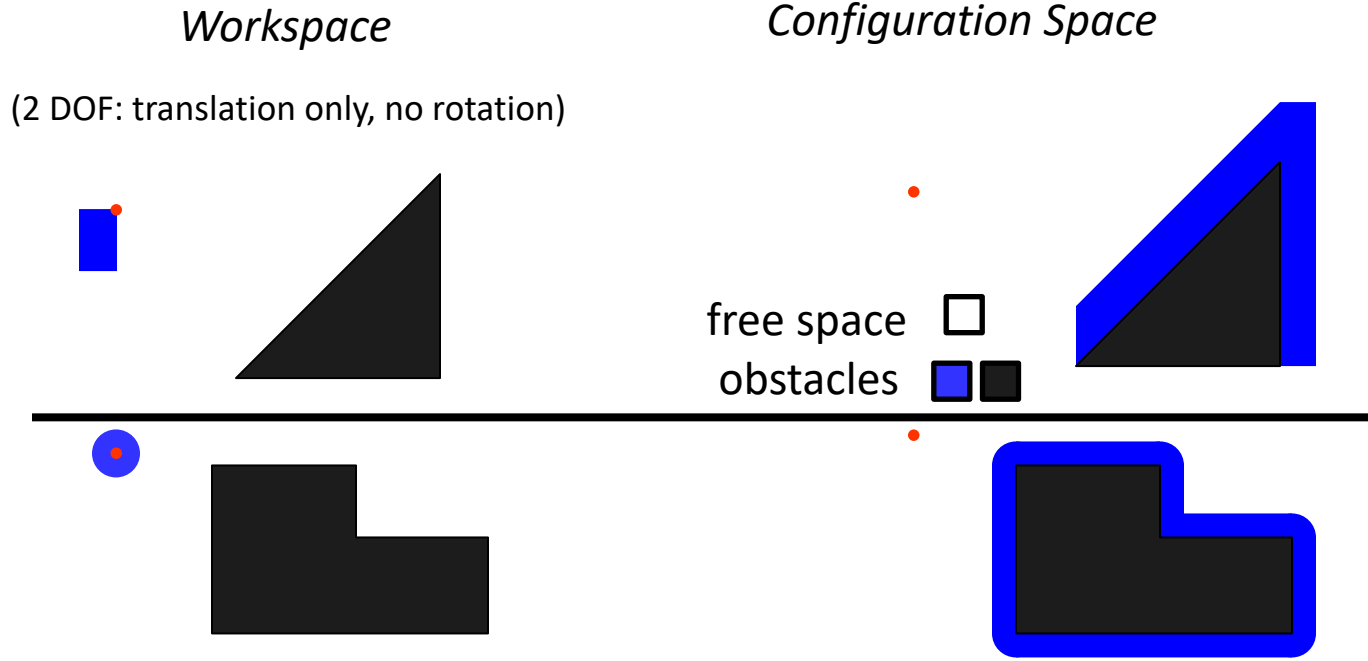
Motion planning



Configuration Space (C-Space)

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

- obstacles \rightarrow configuration space obstacles



Motion Planning: Outline

- Configuration Space
- ***Optimization-based Motion Planning***
- Sampling-based Motion Planning
 - Probabilistic Roadmap
 - Rapidly-exploring Random Trees (RRTs)
 - Smoothing

Optimization-based Motion Planning

- Reactive control
 - Potential-based methods (Khatib '86)
- Optimize over entire trajectory
 - Elastic bands (Quinlan and Khatib '93)
 - CHOMP (Ratliff et al. '09) and variants (STOMP, ITOMP)
 - ***Trajopt (Schulman, et al 2013)***

Solve by Nonlinear Optimization for Control?

- Could try by, for example, following formulation:

$$\begin{aligned} \min_{u,x} \quad & (x_T - x_G)^\top (x_T - x_G) \\ \text{s.t.} \quad & x_{t+1} = f(x_t, u_t) \quad \forall t \\ & u_t \in \mathcal{U}_t \\ & x_t \in \mathcal{X}_t \quad \mathcal{X}_t \text{ can encode obstacles} \\ & x_0 = x_S \end{aligned}$$

- Or, with constraints, (which would require using an infeasible method):

$$\begin{aligned} \min_{u,x} \quad & \|u\| \\ \text{s.t.} \quad & x_{t+1} = f(x_t, u_t) \quad \forall t \\ & u_t \in \mathcal{U}_t \\ & x_t \in \mathcal{X}_t \\ & x_0 = x_S \\ & X_T = x_G \end{aligned}$$

Trajectory Optimization

$$\min_{\theta_{1:T}} \sum_t \|\theta_{t+1} - \theta_t\|^2 + \text{other costs}$$

subject to $\theta_0 = \text{start state}$, θ_T in goal set


joint limits

for all robot parts, for all obstacles:

no collision \longrightarrow ***non-convex***

Solution method: sequential convex optimization

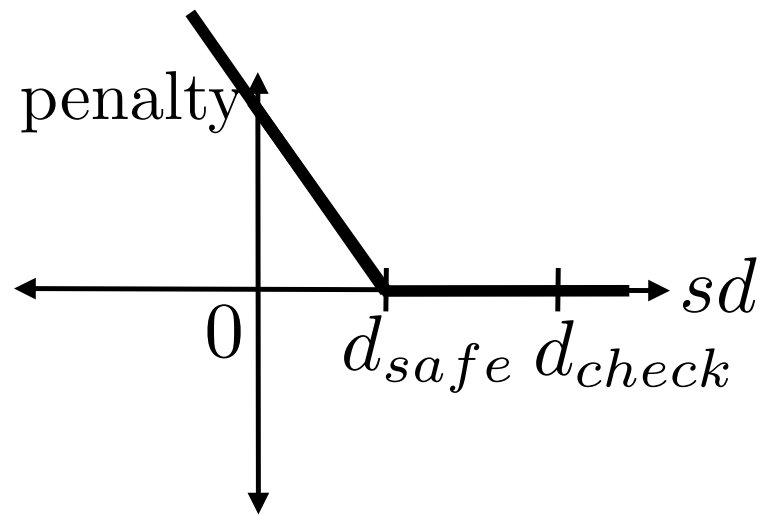
Collision Constraints



$$sd_{AB}(\theta) \approx \hat{n} \cdot (p_B - p_A(\theta))$$

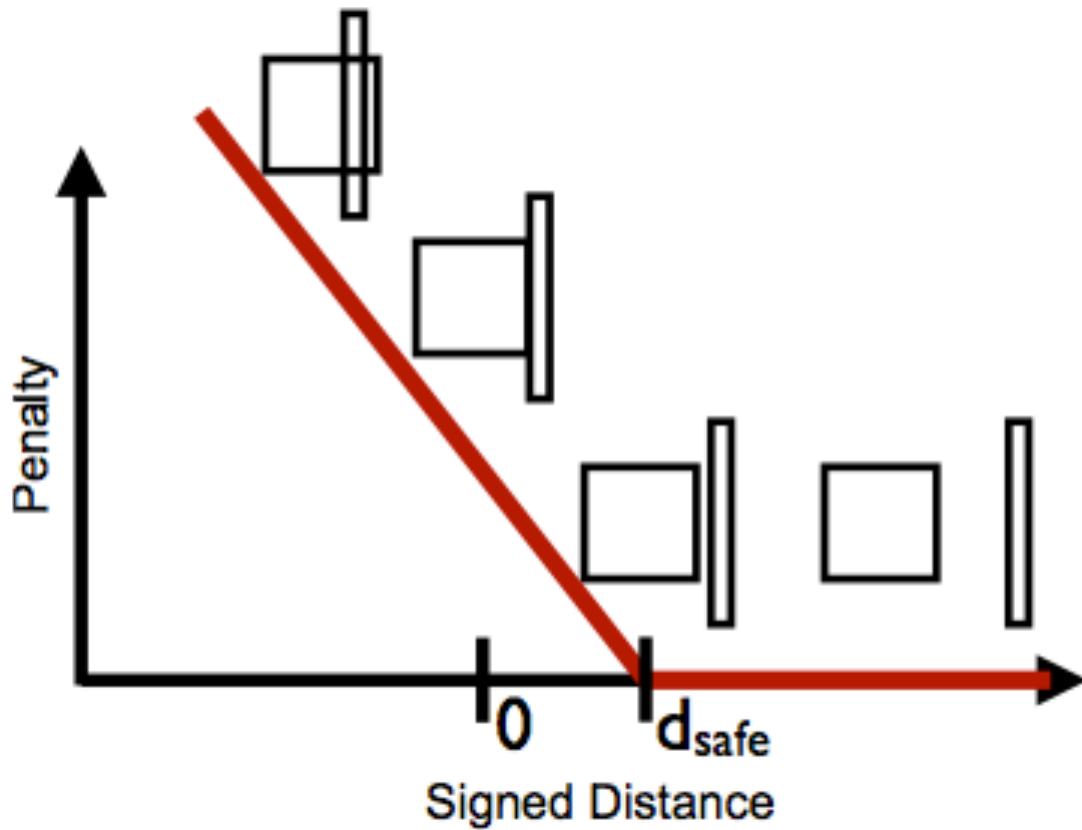
$$\approx sd_{AB}(\theta_0) - \hat{n}^\top J_{P_A}(\theta_0)(\theta - \theta_0)$$

Penalty for Collision Constraints




$$\begin{aligned}sd_{AB}(\theta) &\approx \hat{n} \cdot (p_B - p_A(\theta)) \\ &\approx sd_{AB}(\theta_0) - \hat{n}^\top J_{P_A}(\theta_0)(\theta - \theta_0)\end{aligned}$$


Collision Constraint as L1 Penalty



Collision Constraint as L1 Penalty




Continuous-Time Safety




Collision check against swept-out volume

- Allows coarsely sampling trajectory
 - Overall faster
- Finds better local optima


Collision-free Path for Dubin's Car



Experiments: Industrial Box Picking




Experiments: DRC Robot



Benchmark

7 DOF (one arm)


198 problems



example scene (taken from MoveIt collection)

18 DOF (two arms + base + torso)

96 problems



example scene (imported from Trimble 3d Warehouse / Google Sketchup)

Benchmark Results

Arm planning (7 DOF) 10s limit			
	Trajopt	BiRRT (*)	CHOMP
success	99%	97%	85%
time (s)	0.32	1.2	6.0
path length	1.2	1.6	2.6

Full body (18 DOF) 30s limit			
	Trajopt	BiRRT (*)	CHOMP (**)
success	84%	53%	N/A
time (s)	7.6	18	N/A
path length	1.1	1.6	N/A

(*) Top-performing algorithm from MoveIt/OMPL

(**) Not supported in available implementation

Experiments: PR2




Recycling paper

Step 1: Reach and grab

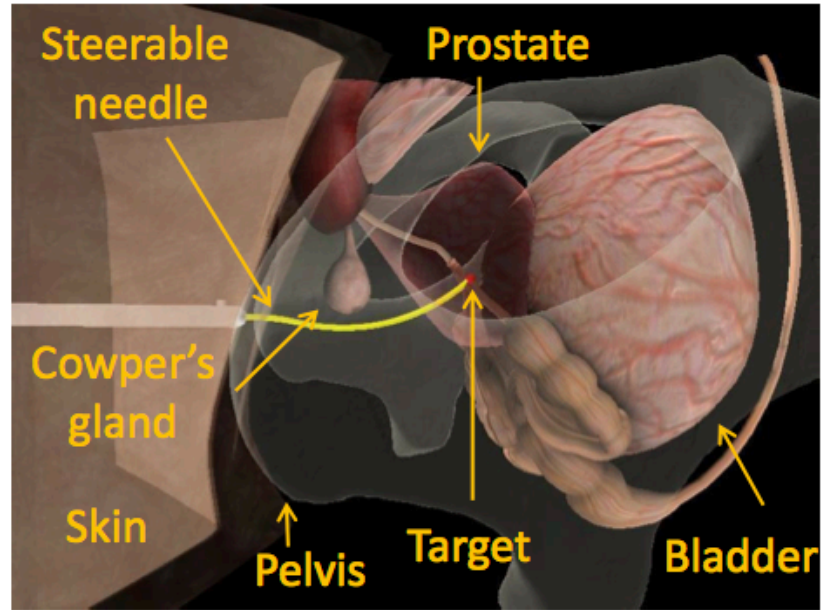
Step 2: Throw away

DOFs: torso, base, right arm

Steerable Needle



Steerable needles inside phantom tissue




Steerable needles navigate around sensitive structures (simulated)


Steerable Needle



Bevel-tip



Highly flexible



State (needle tip) $SE(3): \mathbb{R}^3 \times SO(3)$

- Position: 3D
- Orientation: 3D

Follows constant
curvature paths

Steerable Needle: Opt Formulation

$$\min_{\bar{x}, \mathcal{U}} \alpha_{\Delta} \text{Cost}_{\Delta} + \alpha_{\phi} \text{Cost}_{\phi} + \alpha_{\mathcal{O}} \text{Cost}_{\mathcal{O}},$$



$$\text{s.t. } \log((X_t \cdot \exp(\mathbf{w}_t^{\wedge}) \cdot \exp(\mathbf{v}_t^{\wedge}))^{-1} \cdot X_{t+1})^{\vee} = \mathbf{0}_6,$$

$$\text{sd}(X_t, X_{t+1}, \mathcal{O}_i) \geq d_{\text{safe}} + d_{\text{arc}},$$

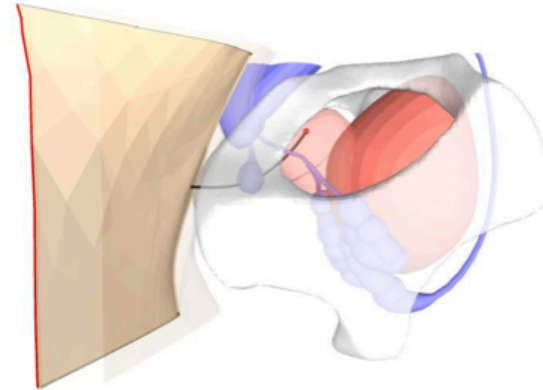

$$X_0 \in Q_{\text{entry}}, X_T \in Q_{\text{target}},$$

$$-\pi \leq \phi_t \leq \pi,$$

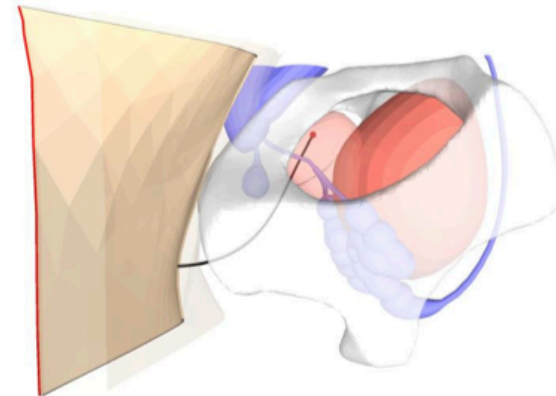
$$\kappa_t = \kappa_{\text{max}}$$



Steerable Needle: Plans



(a) Smaller clearance from obstacles (Cowper's glands) with $\alpha_O = 1$.



(b) Larger clearance from obstacles with $\alpha_O = 10$.

Steerable Needle: Results



	RRT	collocation $\alpha_{\mathcal{O}} = 1$	shooting $\alpha_{\mathcal{O}} = 1$	collocation $\alpha_{\mathcal{O}} = 10$	shooting $\alpha_{\mathcal{O}} = 10$
solved%	67.3%	76.0%	80.3%	79.0%	89.5%
time (s)	9.8 ± 8.1	1.8 ± 1.2	1.6 ± 1.7	1.9 ± 1.3	1.8 ± 1.7
path length	11.1 ± 1.5	11.3 ± 1.4	11.6 ± 1.7	11.9 ± 1.7	13.1 ± 2.3
twist cost	34.9 ± 10.0	1.4 ± 1.4	1.0 ± 1.0	1.6 ± 1.6	1.0 ± 1.0
clearance	0.5 ± 0.4	0.7 ± 0.5	0.5 ± 0.3	1.3 ± 0.4	1.2 ± 0.5

Performance of our approach on the single needle planning case.

Why is minimizing twist important?



Channel Layout (Brachytherapy Implants)



Channel Layout: Opt Formulation

$$\begin{aligned} \min_{\mathcal{X}, \mathcal{U}} \quad & \alpha_{\Delta} \text{Cost}_{\Delta} + \alpha_{\phi} \text{Cost}_{\phi} + \alpha_{\mathcal{O}} \text{Cost}_{\mathcal{O}}, \\ \text{s.t.} \quad & \log((X_t \cdot \exp(\mathbf{w}_t^{\wedge}) \cdot \exp(\mathbf{v}_t^{\wedge}))^{-1} \cdot X_{t+1})^{\vee} = \mathbf{0}_6, \end{aligned}$$


$$\text{sd}(X_t, X_{t+1}, \mathcal{O}_i) \geq d_{\text{safe}} + d_{\text{arc}},$$

$$X_0 \in Q_{\text{entry}}, X_T \in Q_{\text{target}},$$

$$-\pi \leq \phi_t \leq \pi,$$

$$0 \leq \kappa_t \leq \kappa_{\text{max}},$$


$$\Delta \sum_{t=0}^{T-1} \kappa_t \leq c_{\text{max}} \quad \text{for channel planning,}$$



Channel Layout: Results





	RRT	backward shooting
solved%	74.0%	98.0%
time (s)	30.8 ± 17.9	27.7 ± 9.8
path length	41.3 ± 0.3	38.9 ± 0.1
twist cost	65.5 ± 8.4	4.1 ± 1.1

Performance of our approach on the channel layout planning



Try It Yourself

- Code and docs: rl.berkeley.edu/trajopt
- Benchmark: github.com/joschu/planning_benchmark



Experiments: PR2


PR2 Obstacle Course

DOFs: Base, Torso, Arms

Motion Planning: Outline


- Configuration Space
- Optimization-based Motion Planning
- ***Sampling-based Motion Planning***
 - ***Probabilistic Roadmap***
 - Rapidly-exploring Random Trees (RRTs)
 - Smoothing

Probabilistic Roadmap (PRM)




Probabilistic Roadmap (PRM)

Configurations are sampled by picking coordinates at random




Probabilistic Roadmap (PRM)

Configurations are sampled by picking coordinates at random




Probabilistic Roadmap (PRM)

Sampled configurations are tested for collision




Probabilistic Roadmap (PRM)

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




Probabilistic Roadmap (PRM)

Each milestone is linked by straight paths to its nearest neighbors




Probabilistic Roadmap (PRM)

Each milestone is linked by straight paths to its nearest neighbors




Probabilistic Roadmap (PRM)

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




Probabilistic Roadmap (PRM)

The start and goal configurations are included as milestones



Probabilistic Roadmap (PRM)


The PRM is searched for a path from s to g



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

PRM Example 1



PRM Example 2



Sampling

- How to sample uniformly at random from $[0,1]^n$?
 - Sample uniformly at random from $[0,1]$ for each coordinate
- How to sample uniformly at random from the surface of the n-D unit sphere?
 - Sample from n-D Gaussian \rightarrow isotropic; then just normalize
- How to sample uniformly at random for orientations in 3-D?

PRM: Challenges

1. Connecting neighboring points: Only easy for holonomic systems (i.e., for which you can move each degree of freedom at will at any time). Generally requires solving a Boundary Value Problem

$$\begin{aligned} \min_{u,x} \quad & \|u\| \\ \text{s.t.} \quad & x_{t+1} = f(x_t, u_t) \quad \forall t \\ & u_t \in \mathcal{U}_t \\ & x_t \in \mathcal{X}_t \\ & x_0 = x_S \\ & X_T = x_G \end{aligned}$$

Typically solved without collision checking; later verified if valid by collision checking

2. Collision checking:

Often takes majority of time in applications (see Lavelle)

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 entire state space, which might be unnecessarily expensive when what's needed is connecting specific start and goal

Motion Planning: Outline

- Configuration Space
- Optimization-based Motion Planning
- ***Sampling-based Motion Planning***
 - Probabilistic Roadmap
 - ***Rapidly-exploring Random Trees (RRTs)***
 - Smoothing

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

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       $x_{rand} \leftarrow$  RANDOM_STATE();
4       $x_{near} \leftarrow$  NEAREST_NEIGHBOR( $x_{rand}$ ,  $\mathcal{T}$ );
5       $u \leftarrow$  SELECT_INPUT( $x_{rand}$ ,  $x_{near}$ );
6       $x_{new} \leftarrow$  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
SELECT_INPUT(): often a few inputs are sampled, and one that results in x_{new} closest to x_{rand} is retained; sometimes optimization is run to find the best input

Rapidly exploring Random Tree (RRT)

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

Rapidly exploring Random Tree (RRT)




RRT Practicalities


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

RRT Extension

- No obstacles, holonomic:



- With obstacles, holonomic:




- Non-holonomic: approximately solve two-point boundary value problem (often rough approximation: pick best of a few random control sequences)

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)

Bi-directional RRT

- Volume swept out by unidirectional RRT:



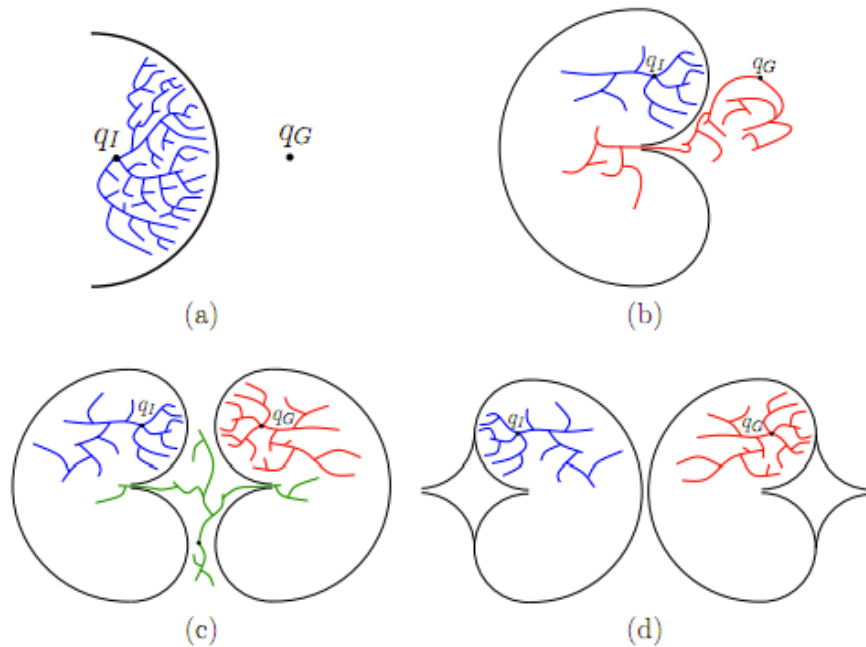
- Volume swept out by bi-directional RRT:



- Difference more and more pronounced as dimensionality increases


Multi-directional RRT

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



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 $1/m^2$, ...
- When a node is selected for expansion, skip over it with probability CVF/m

RRT*

Algorithm 6: RRT*


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

RRT*

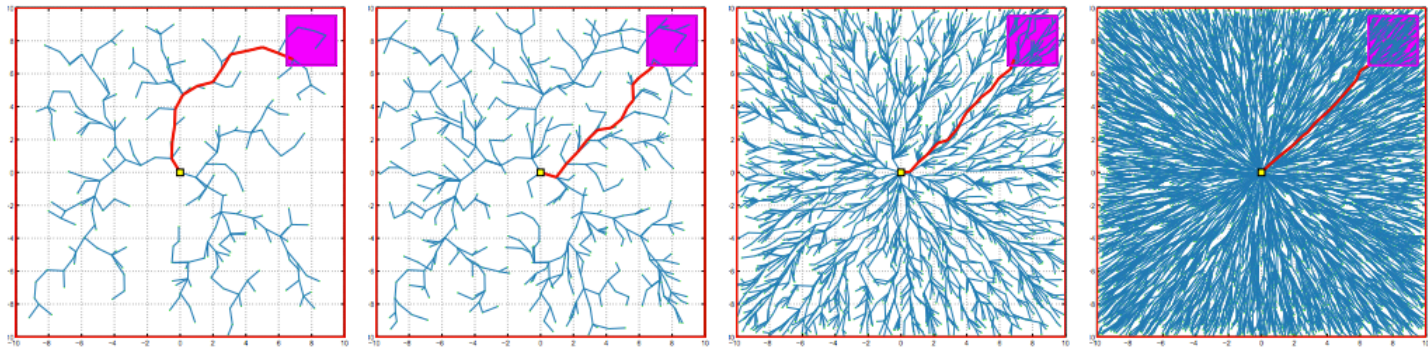
- Asymptotically optimal
- Main idea:
 - Swap new point in as parent for nearby vertices who can be reached along shorter path through new point than through their original (current) parent

RRT*

RRT




RRT*




RRT*

RRT



RRT*



RRT* Kinodynamics

- Requires 2-point boundary value problem solution for optimality
- Li, Littlefield, Bekris 2014 proved that you can get asymptotic optimality from random sampling control trajectories in an RRT like fashion (Naïve Random Tree), without solving a 2-point boundary value problem
- They also show that using pruning can make this efficient in an algorithm called SST*


PRM* Probabilistic Bounds

- Dobson, Moustakides, Bekris 2014
- Gave finite time bounds for the current best path I_n being within a certain threshold of the optimal cost length I_n^* for a fixed delta of the form:

$$\mathbb{P}(|I_n - I_{\epsilon_n}^*| \leq \delta \cdot I_{\epsilon_n}^*) \leq \mathbb{P}_{\text{success}}$$

LQR-trees (Tedrake, IJRR 2010)

- Idea: grow a randomized tree of stabilizing controllers to the goal
 - Like RRT
 - Can discard sample points in already stabilized region




LQR-trees (Tedrake)

Algorithm 1 LQR-tree (\mathbf{f} , \mathbf{x}_G , \mathbf{u}_G , \mathbf{Q} , \mathbf{R})


```
1:  $[\mathbf{A}, \mathbf{B}] \leftarrow$  linearization of  $\mathbf{f}(\mathbf{x}, \mathbf{u})$  around  $(\mathbf{x}_G, \mathbf{u}_G)$ 
2:  $[\mathbf{K}, \mathbf{S}] \leftarrow$  LQR( $\mathbf{A}, \mathbf{B}, \mathbf{Q}, \mathbf{R}$ )
3:  $\rho_c \leftarrow$  level set computed as described in §3.1.1
4:  $T.\text{init}(\{\mathbf{x}_g, \mathbf{u}_g, \mathbf{S}, \mathbf{K}, \rho_c, \text{NULL}\})$ 
5: for  $k = 1$  to  $K$  do
6:    $\mathbf{x}_{\text{rand}} \leftarrow$  random sample
7:   if  $\mathbf{x}_{\text{rand}} \in \mathcal{C}_k$  then
8:     continue
9:   end if
10:   $[t, \mathbf{x}_0(t), \mathbf{u}_0(t)]$  from trajectory optimization with a
    “final tree constraint”
11:  if  $\mathbf{x}_0(t_f) \notin \mathcal{T}_k$  then
12:    continue
13:  end if
14:   $[\mathbf{K}(t), \mathbf{S}(t)]$  from time-varying LQR
15:   $\rho_c \leftarrow$  level set computed as in §3.1.1
16:   $i \leftarrow$  pointer to branch in  $T$  containing  $\mathbf{x}_0(t_f)$ 
17:   $T.\text{add-branch}(\mathbf{x}_0(t), \mathbf{u}_0(t), \mathbf{S}(t), \mathbf{K}(t), \rho_c, i)$ 
18: end for
```

Ck: stabilized
region after
iteration k


LQR-trees (Tedrake)




(a) Goal region



(b) One branch



(c) Six branches



(d) Thirteen branches

Motion Planning: Outline

- Configuration Space
- Optimization-based Motion Planning
- ***Sampling-based Motion Planning***
 - Probabilistic Roadmap
 - Rapidly-exploring Random Trees (RRTs)
 - ***Smoothing***

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 (trajopt)

- Allows to specify an objective function that includes smoothness in state, control, small control inputs, etc.