Autonomous and Mobile Robotics

Prof. Giuseppe Oriolo

Motion Planning

Probabilistic Methods

Dipartimento di Ingegneria Informatica
Automatica e Gestionale Antonio Ruberti

Sapienza
Università di Roma

sampling-based methods

- build a roadmap of the configuration space $\mathcal{C}$ by repeating this basic iteration:
  - extract a sample $\boldsymbol{q}$ of $\mathcal{C}$
  - use forward kinematics to compute the volume $\mathcal{B}(\boldsymbol{q})$ occupied by the robot $\mathcal{B}$ at $\boldsymbol{q}$
  - check collision between $\mathcal{B}(\boldsymbol{q})$ and obstacles $\mathcal{O}_1,\ldots,\mathcal{O}_p$
  - if $\boldsymbol{q} \in \mathcal{C}_{\text{free}}$, add $\boldsymbol{q}$ to the roadmap; else, discard it
- preliminary computation of $\mathcal{CO}$ is completely avoided: an approximate representation of $\mathcal{C}_{\text{free}}$ is directly built as a collection of connected configurations (roadmap)
- different criteria for sampling lead to different methods: in general, randomized outperforms deterministic

PRM (Probabilistic Roadmap)

- basic iteration to build the PRM:
  - extract a sample $\boldsymbol{q}$ of $\mathcal{C}$ with uniform probability distribution
  - compute $\mathcal{B}(\boldsymbol{q})$ and check for collision
  - if $\boldsymbol{q} \in \mathcal{C}_{\text{free}}$, add $\boldsymbol{q}$ to the PRM; else, discard it
  - search the PRM for "sufficiently near" configurations $\boldsymbol{q}_{\text{near}}$
  - if possible, connect $\boldsymbol{q}$ to $\boldsymbol{q}_{\text{near}}$ with a free local path
- the generation of a free path between $\boldsymbol{q}$ and $\boldsymbol{q}_{\text{near}}$ is delegated to a procedure called local planner: e.g., throw a linear path and check it for collision
- the chosen metric in $\mathcal{C}$ plays a role in identifying $\boldsymbol{q}_{\text{near}}$

disconnected components

narrow passages are scarcely sampled

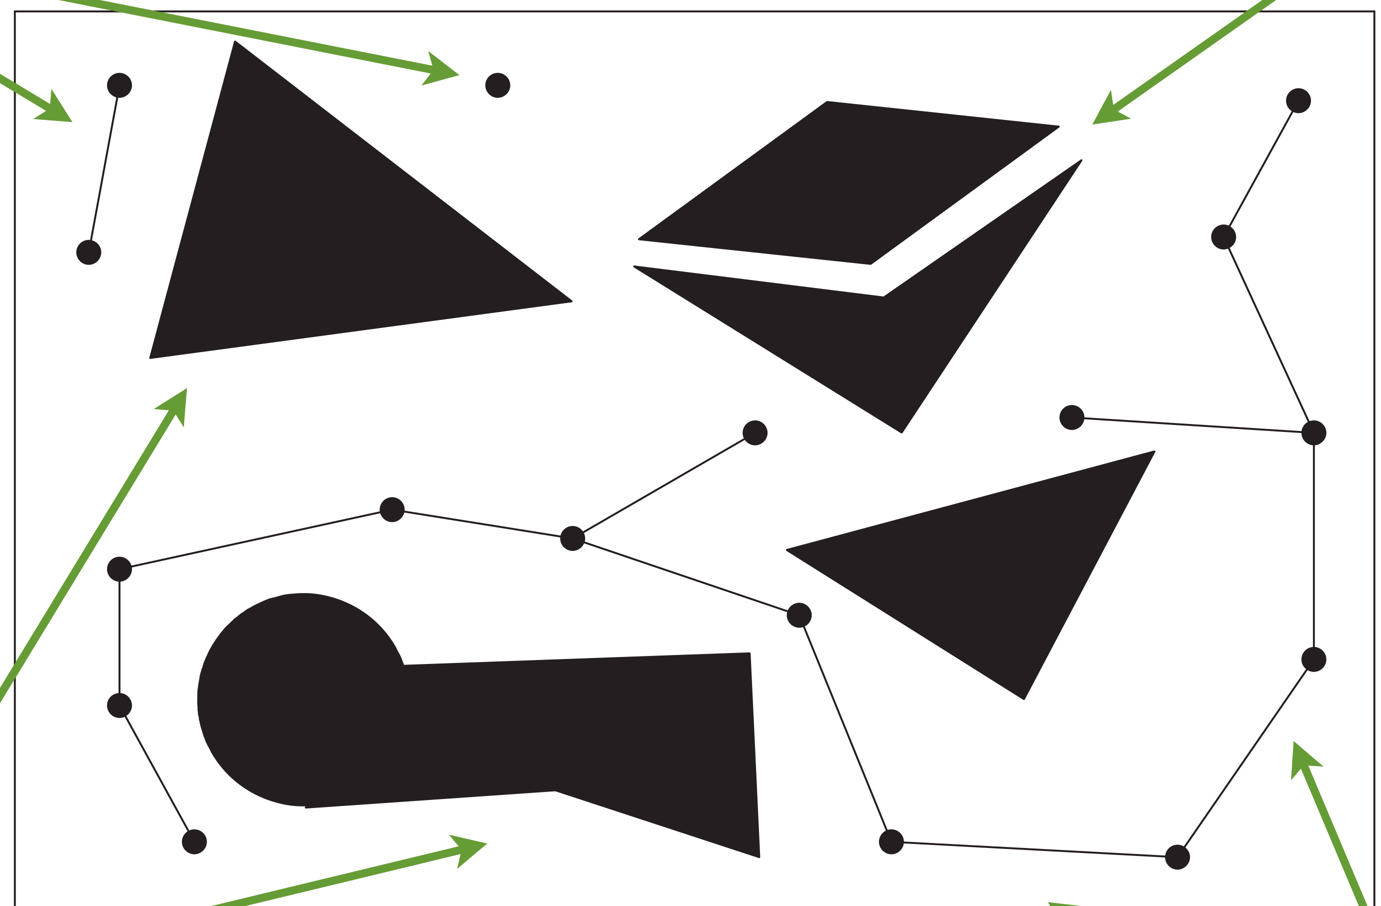

$\mathcal{C}$-obstacles are never computed

local paths

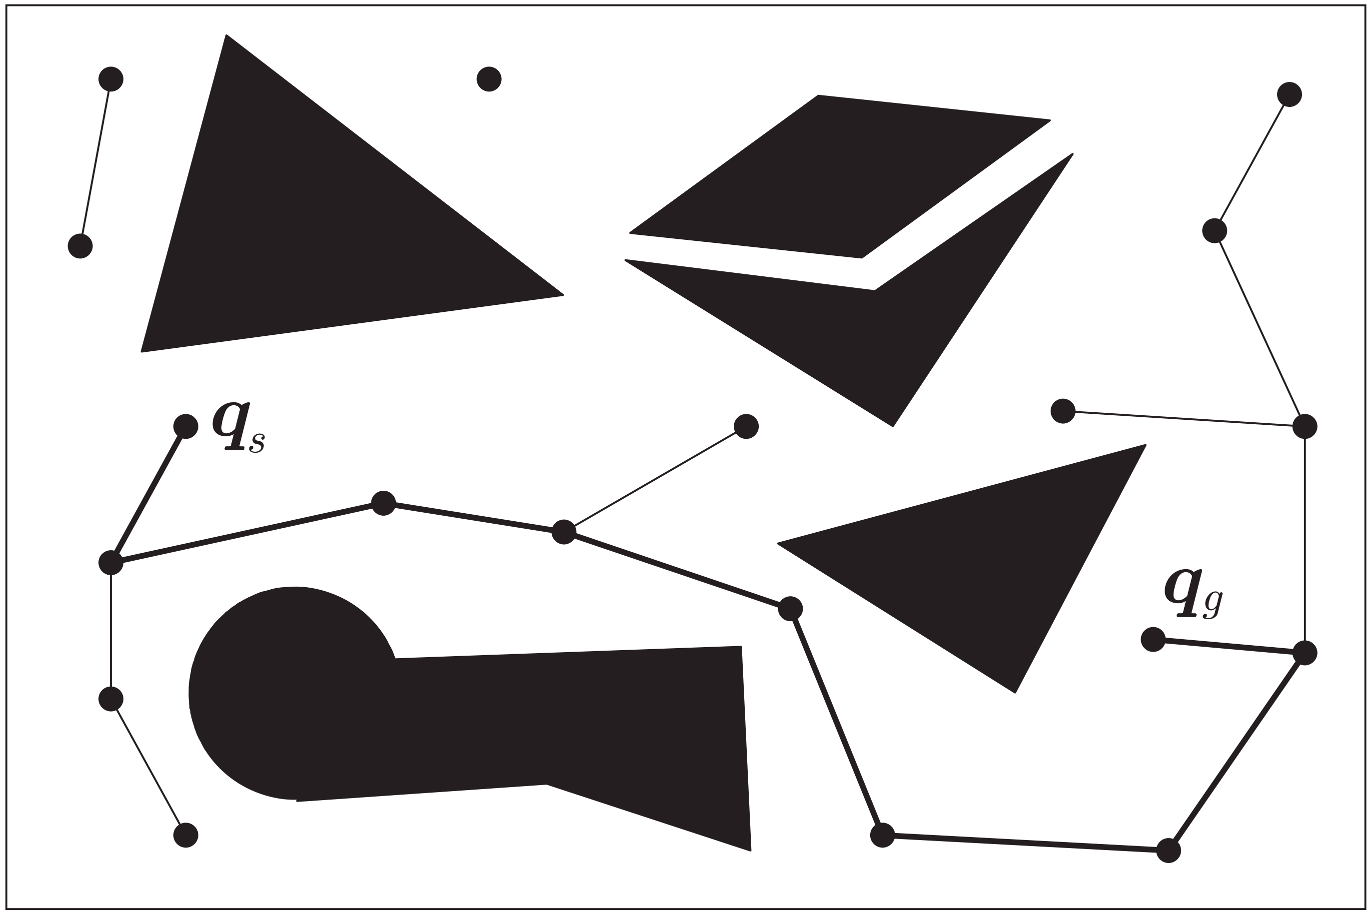

- construction of the PRM is arrested when
  1. disconnected components become less than a threshold, or
  2. a maximum number of iterations is reached
- if q_s and q_g can be connected to the same component, a solution can be found by graph search; else, enhance the PRM by performing more iterations

- the PRM method is probabilistically complete, i.e., the probability of finding a solution whenever one exists tends to 1 as the execution time tends to ∞; and is multiple-query (new queries enhance the PRM)

- the main advantage is speed; the time PRM needs to find a solution in high-dimensional spaces can be orders of magnitude smaller than previous planners

- narrow passages are critical; heuristics may be used to design biased (non-uniform) probability distributions aimed at increasing sampling in such areas

RRT (Rapidly-exploring Random Tree)

- basic iteration to build the tree T_s rooted at $\boldsymbol{q}_s$:
  - generate $\boldsymbol{q}_{\rm rand}$ in $\mathcal{C}$ with uniform probability distribution
  - search the tree for the nearest configuration $\boldsymbol{q}_{\rm near}$
  - choose $\boldsymbol{q}_{\rm new}$ at a distance δ from $\boldsymbol{q}_{\rm near}$ in the direction of $\boldsymbol{q}_{\rm rand}$
  - check for collision $\boldsymbol{q}_{\rm new}$ and the segment from $\boldsymbol{q}_{\rm near}$ to $\boldsymbol{q}_{\rm new}$
  - if check is negative, add $\boldsymbol{q}_{\rm new}$ to T_s (expansion)

- the chosen metric in $\mathcal{C}$ plays a role in identifying $\boldsymbol{q}_{\rm near}$

- T_s rapidly covers $\mathcal{C}_{\rm free}$ because the expansion is biased towards unexplored areas (actually, towards larger Voronoi regions)

tree is
rooted at $\boldsymbol{q}_s$

$\mathcal{C}$-obstacles are
never computed

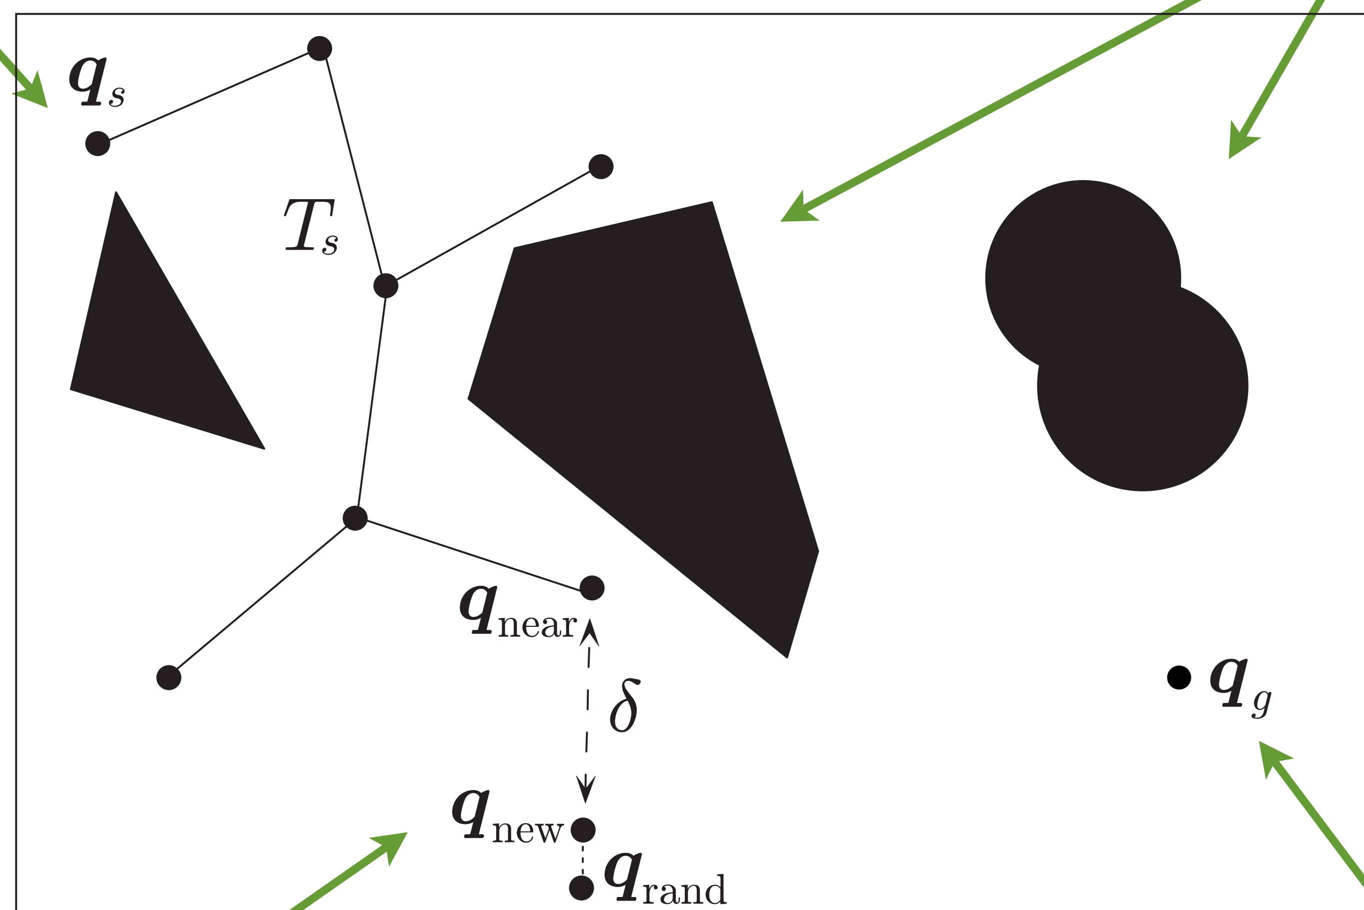

tree
expansion

no bias
towards $\boldsymbol{q}_g$

RRT in empty 2D space

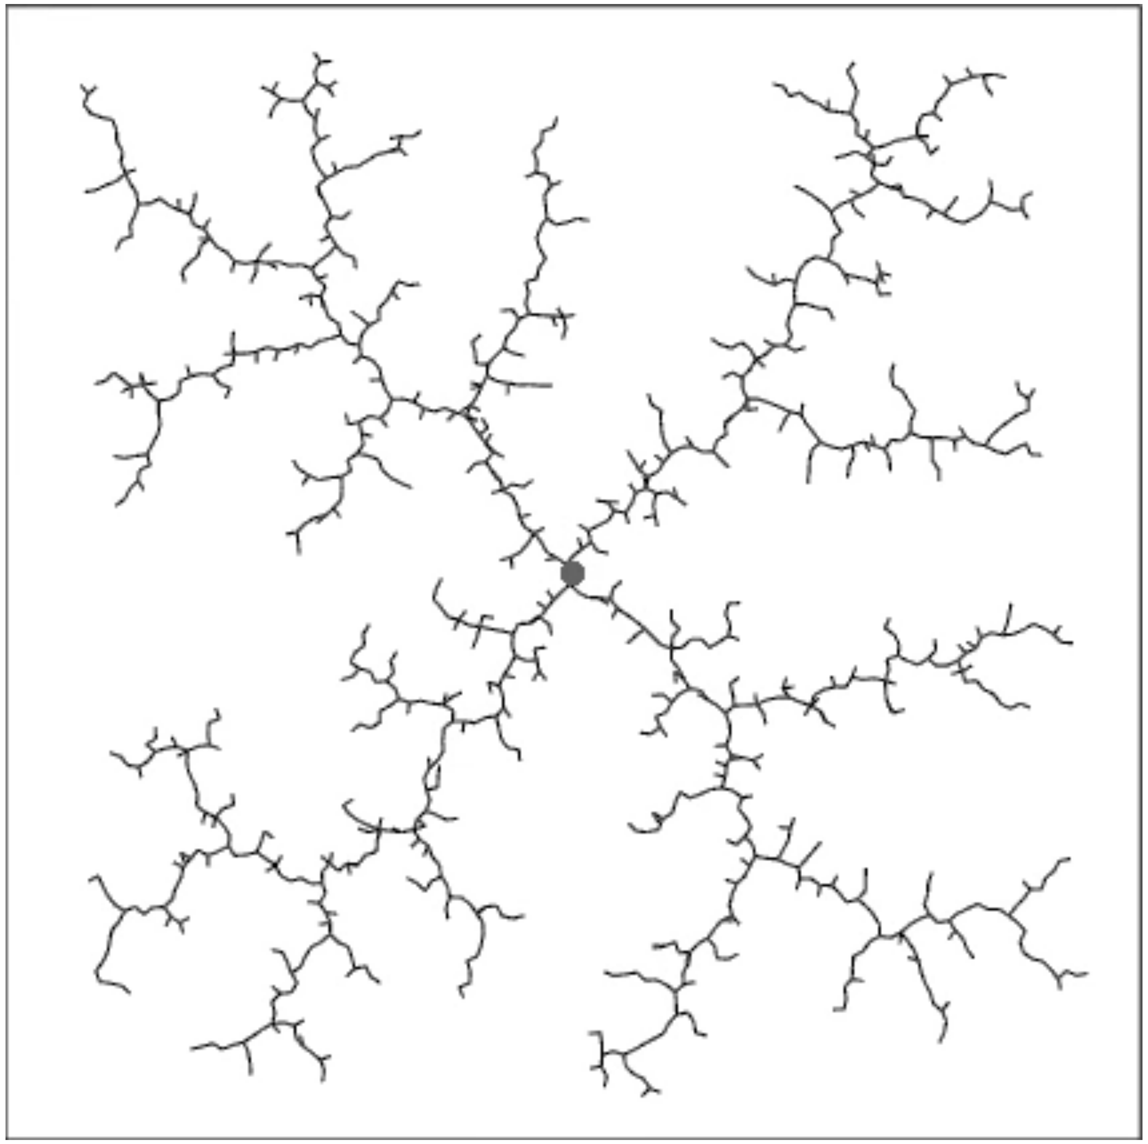

quickly explores all areas, much more efficiently than other simple strategies, e.g., random walks

- to introduce a bias towards $\boldsymbol{q}_g$, one may grow two trees T_s and T_g, respectively rooted at $\boldsymbol{q}_s$ and $\boldsymbol{q}_g$ (bidirectional RRT)
- alternate expansion and connection phases: use the last generated $\boldsymbol{q}_{\text{new}}$ of T_s as a $\boldsymbol{q}_{\text{rand}}$ for T_g, and then repeat switching the roles of T_s and T_g

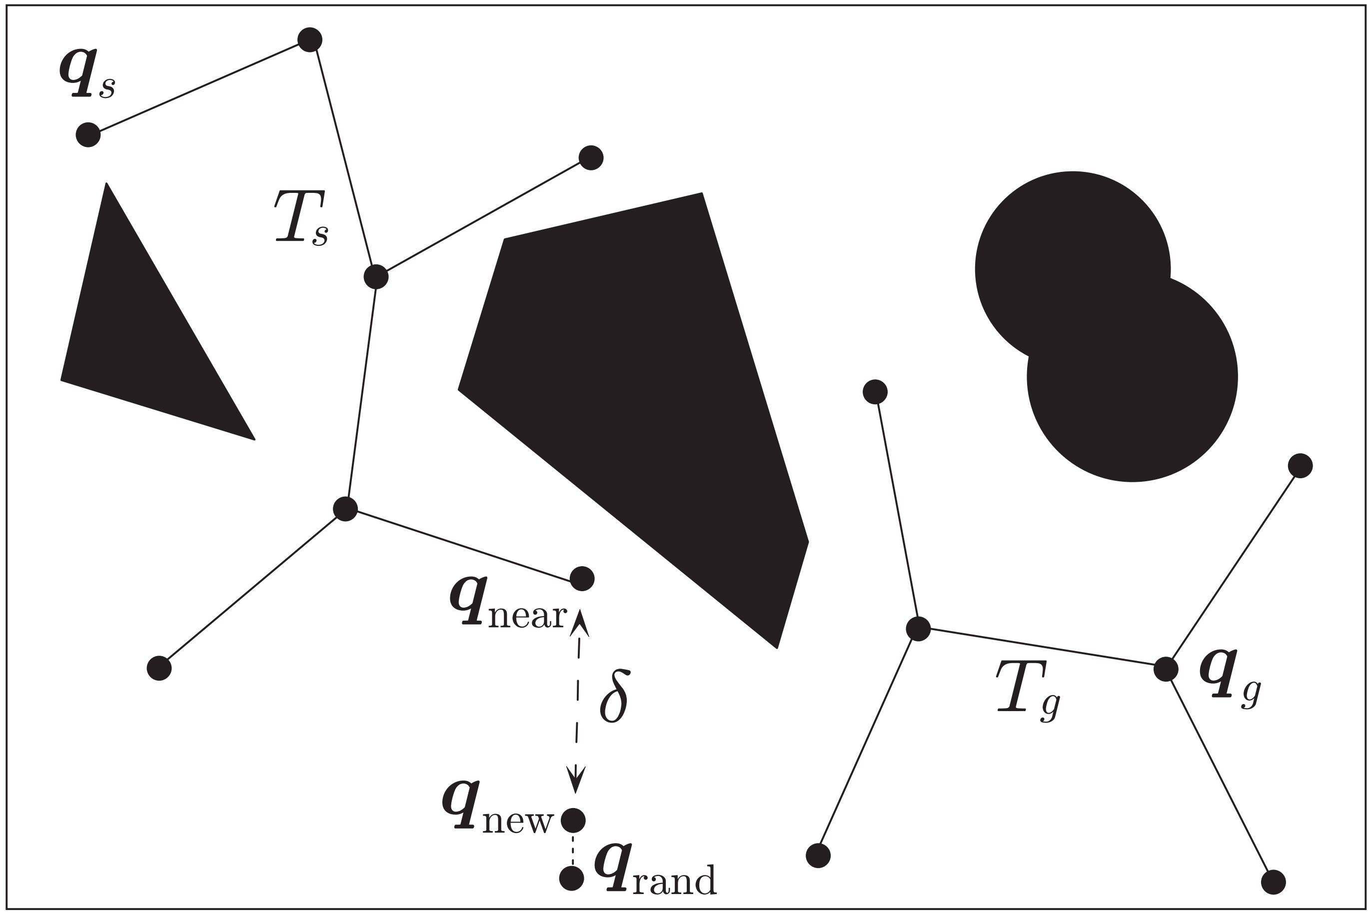

- bidirectional RRT is probabilistically complete and single-query (trees are rooted at $\boldsymbol{q}_s$ and $\boldsymbol{q}_g$, and in any case new queries may require significant work)

- many variations are possible: e.g., one may use an adaptive stepsize δ to speed up motion in wide open areas (greedy exploration)

- can be modified to address many extensions of the canonical planning problem, e.g., moving obstacles, nonholonomic constraints, manipulation planning

a benchmark problem: the Alpha Puzzle

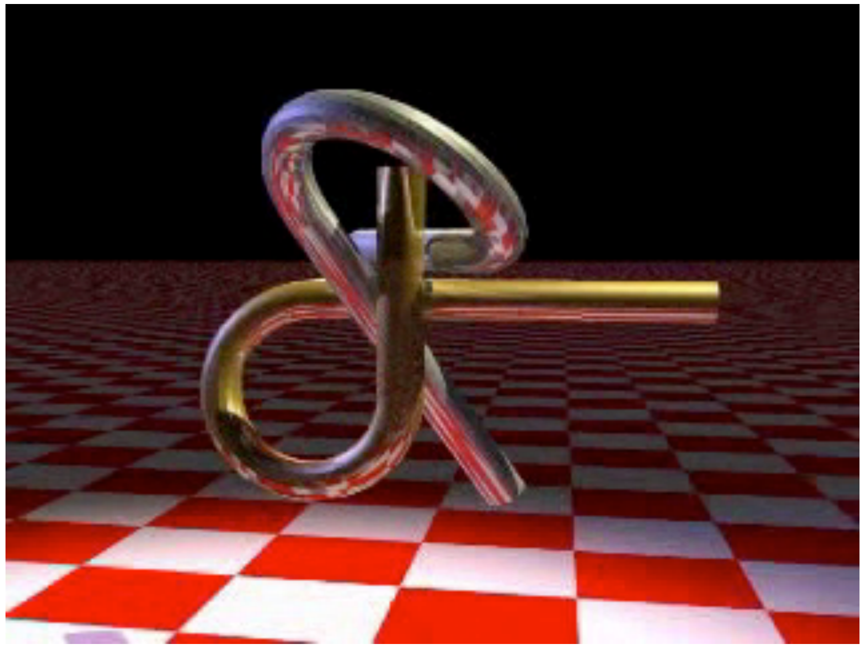

- 6-dof configuration space + narrow passages
- solved by bidirectional RRT in few mins (average)
- in practice, this problem is not solvable by classical methods such as retraction or cell decomposition

RRT: extension to nonholonomic robots

- motion planning for a unicycle in $\mathcal{C} = \mathrm{R}^2 \times SO(2)$

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{pmatrix} = \begin{pmatrix} \cos\theta \\ \sin\theta \\ 0 \end{pmatrix} v + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \omega$$

- linear paths in $\mathcal{C}$ such as those used to connect $\boldsymbol{q}_{\rm near}$ to $\boldsymbol{q}_{\rm rand}$ are not admissible in general

- one possibility is to use motion primitives, i.e., a finite set of admissible local paths, produced by a specific choice of the velocity inputs

- for example, one may use (Dubins car)

$$v = \bar{v} \qquad \omega = \{-\bar{\omega}, 0, \bar{\omega}\} \qquad t \in [0, \Delta]$$

resulting in 3 possible paths in forward motion

- the algorithm is the same with the only difference that $\boldsymbol{q}_{\rm new}$ is generated from $\boldsymbol{q}_{\rm near}$ selecting one of the possible paths (either randomly or as the one that leads the unicycle closer to $\boldsymbol{q}_{\rm rand}$)

- if $\boldsymbol{q}_g$ can be reached from $\boldsymbol{q}_s$ with a collision-free concatenation of primitives, the probability that a solution is found tends to 1 as the time tends to ∞

solution path made by concatenation

primitives