Algorithmic Robotics and Motion Planning

Sampling-based motion planning I: PRM

Dan Halperin
School of Computer Science
Tel Aviv University

Fall 2019-2020

Overview

- PRM, preliminaries
- the ingredients of PRM
- probabilistic completeness

PRM, the original article

Lydia E. Kavraki, Petr Svestka, Jean-Claude Latombe, Mark H. Overmars: **Probabilistic roadmaps for path planning in high-dimensional configuration spaces.** IEEE Trans. Robotics and Automation 12(4): 566-580 (1996)

PRM = Probabilistic Roadmap

5.6.1 The Basic Method

Once again, let $\mathcal{G}(V, E)$ represent a topological graph in which V is a set of vertices and E is the set of paths that map into $\mathcal{C}_{free}$. Under the multiple-query philosophy, motion planning is divided into two phases of computation:

> **Preprocessing Phase:** During the preprocessing phase, substantial effort is invested to build $\mathcal{G}$ in a way that is useful for quickly answering future queries. For this reason, it is called a *roadmap*, which in some sense should be accessible from every part of $\mathcal{C}_{free}$.
>
> **Query Phase:** During the query phase, a pair, q_I and q_G, is given. Each configuration must be connected easily to $\mathcal{G}$ using a local planner. Following this, a discrete search is performed using any of the algorithms in Section 2.2 to obtain a sequence of edges that forms a path from q_I to q_G.

[LaValle's Planning Algorithms]

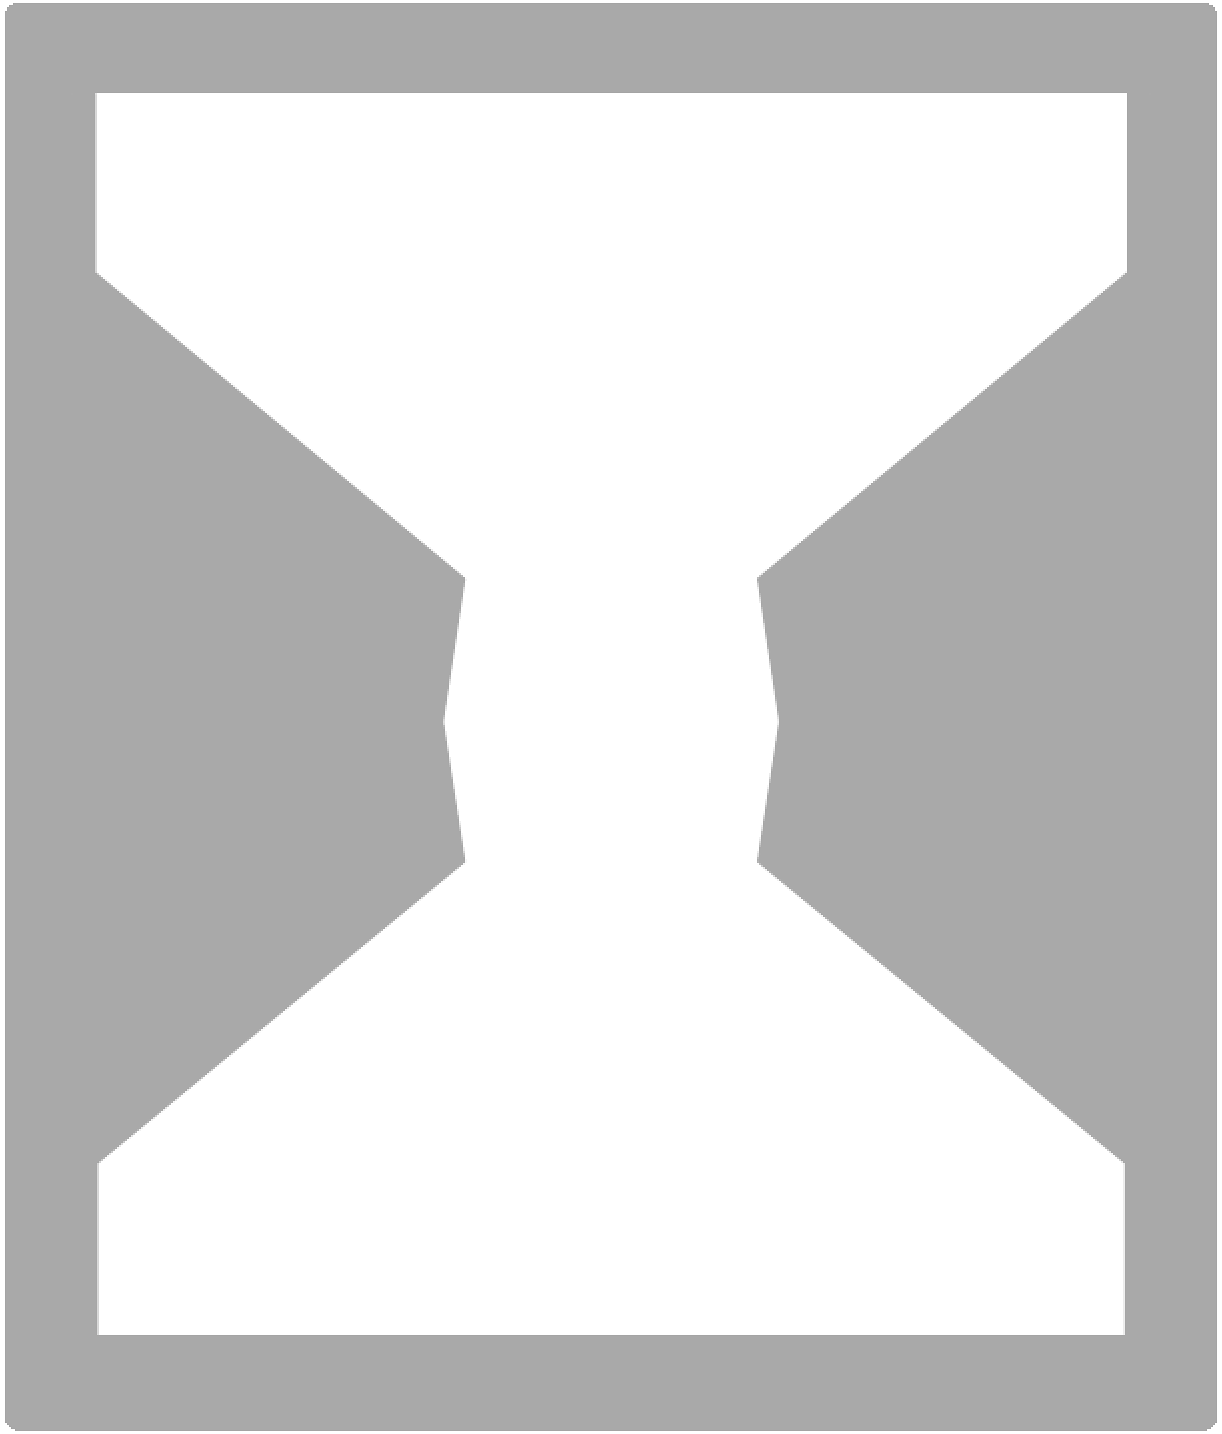

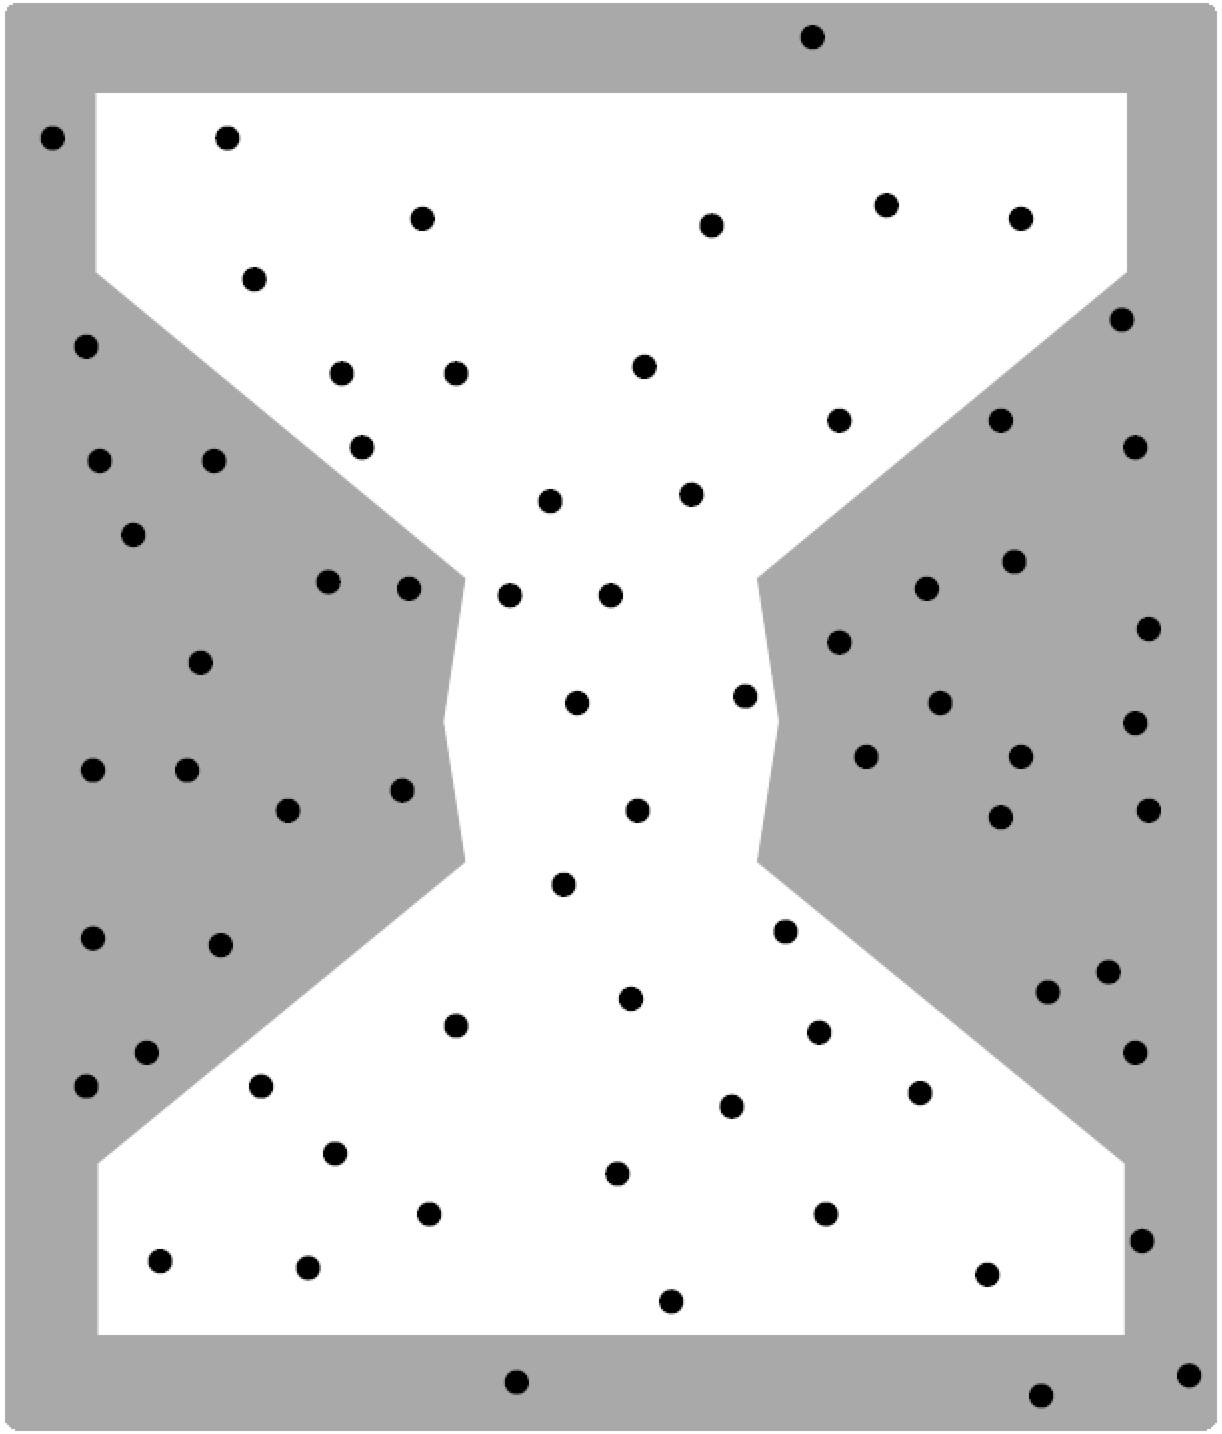

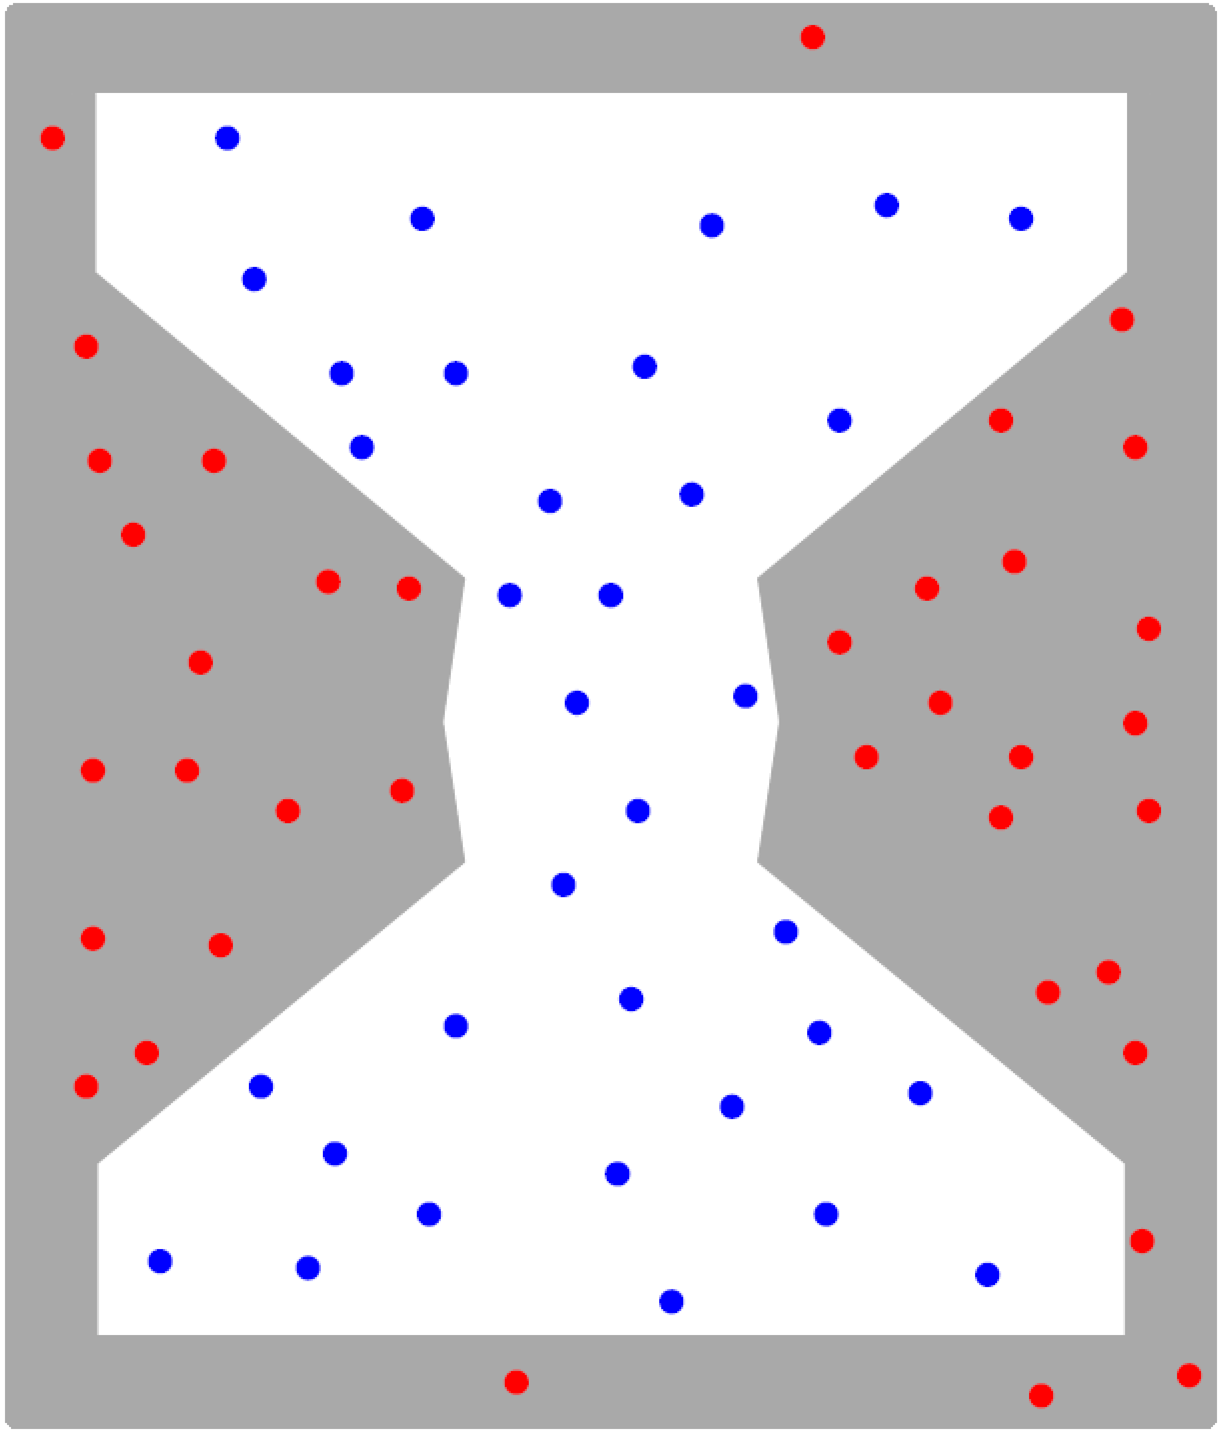

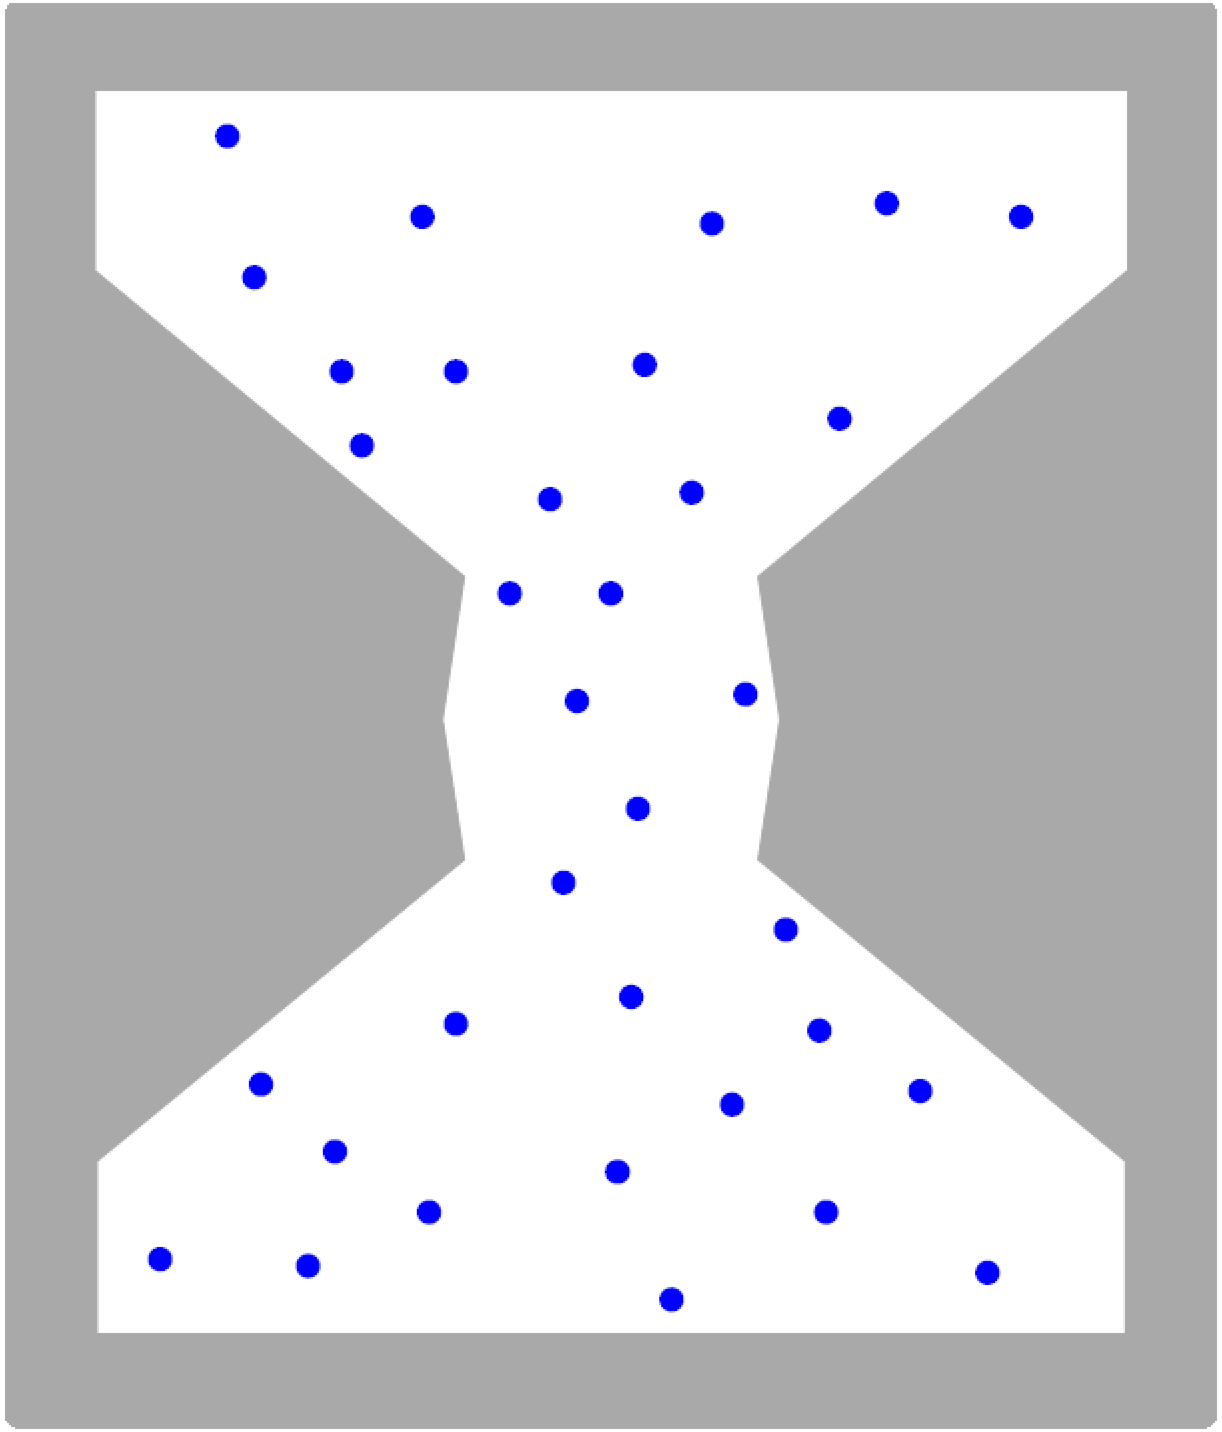

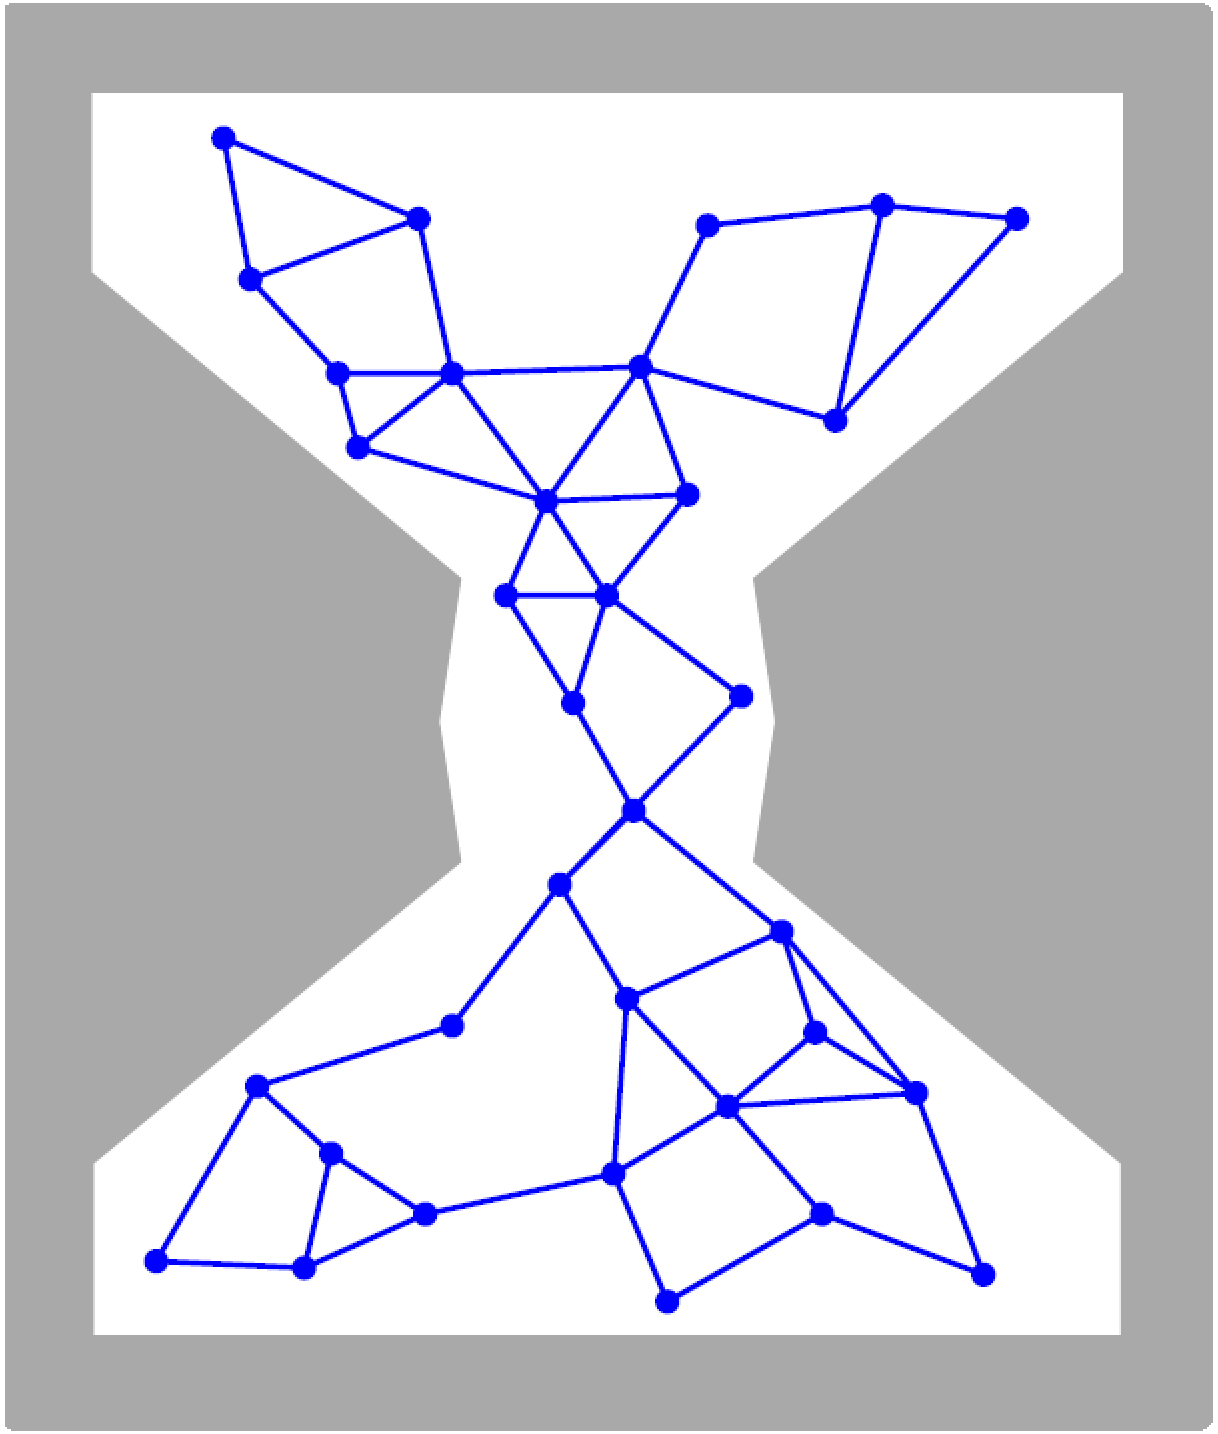

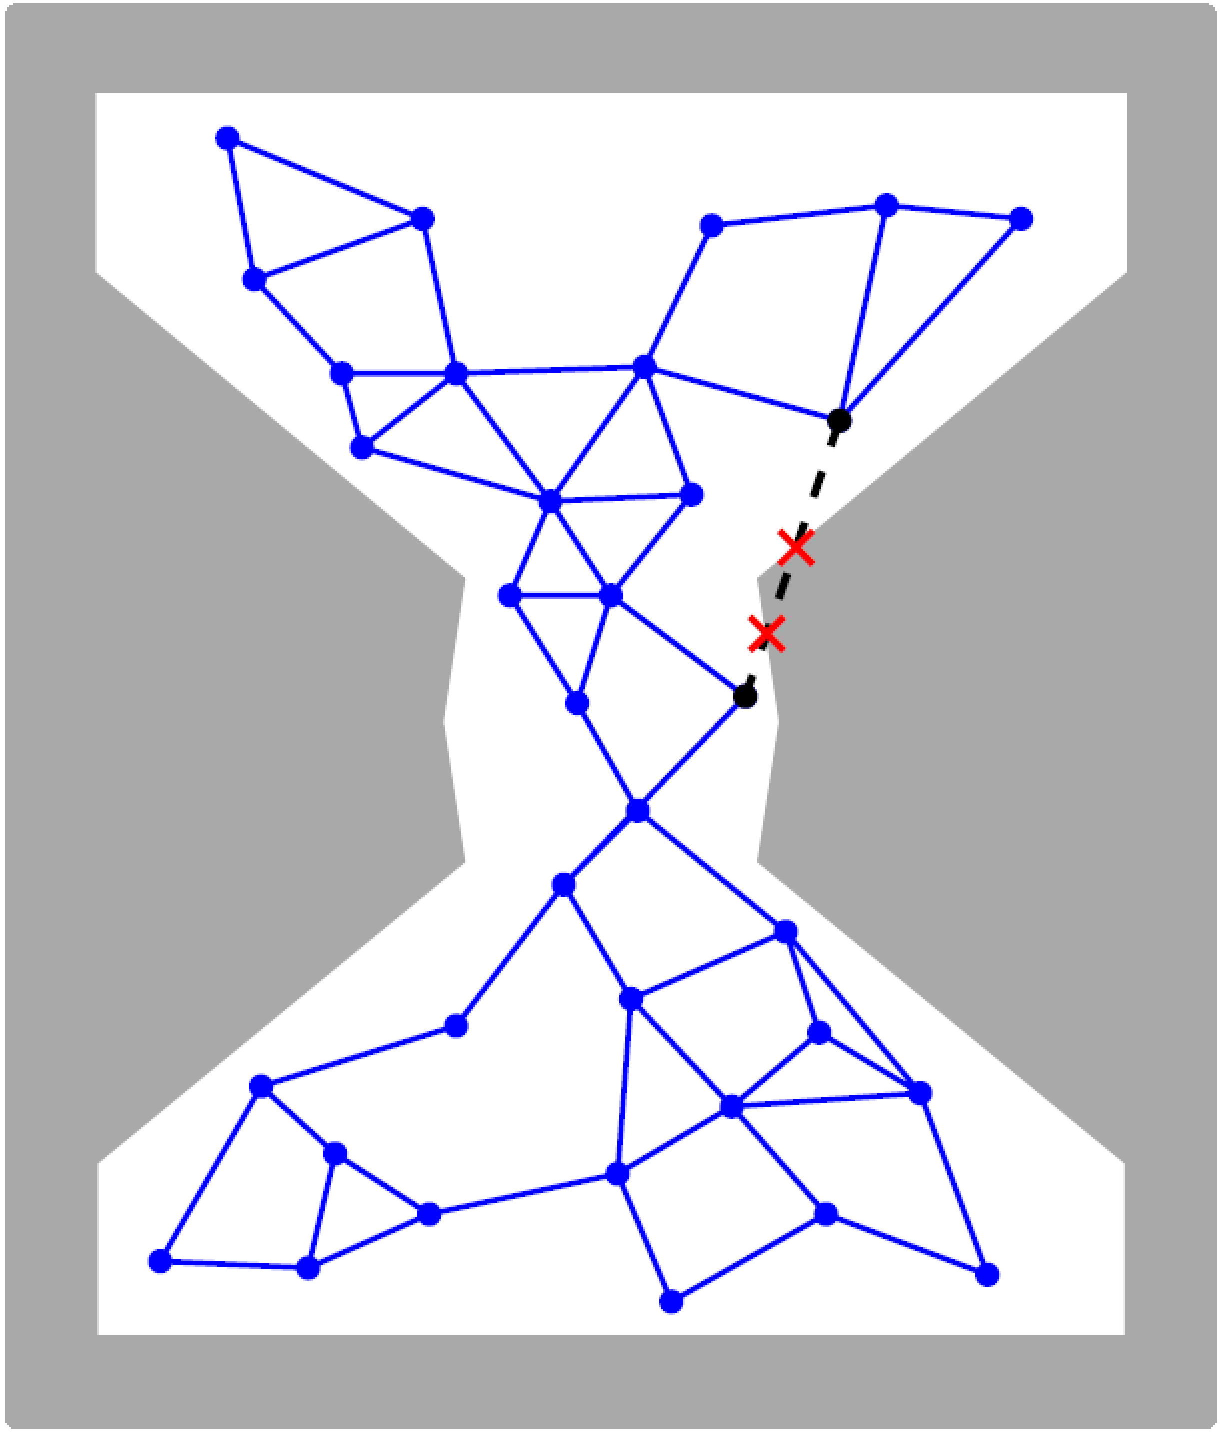

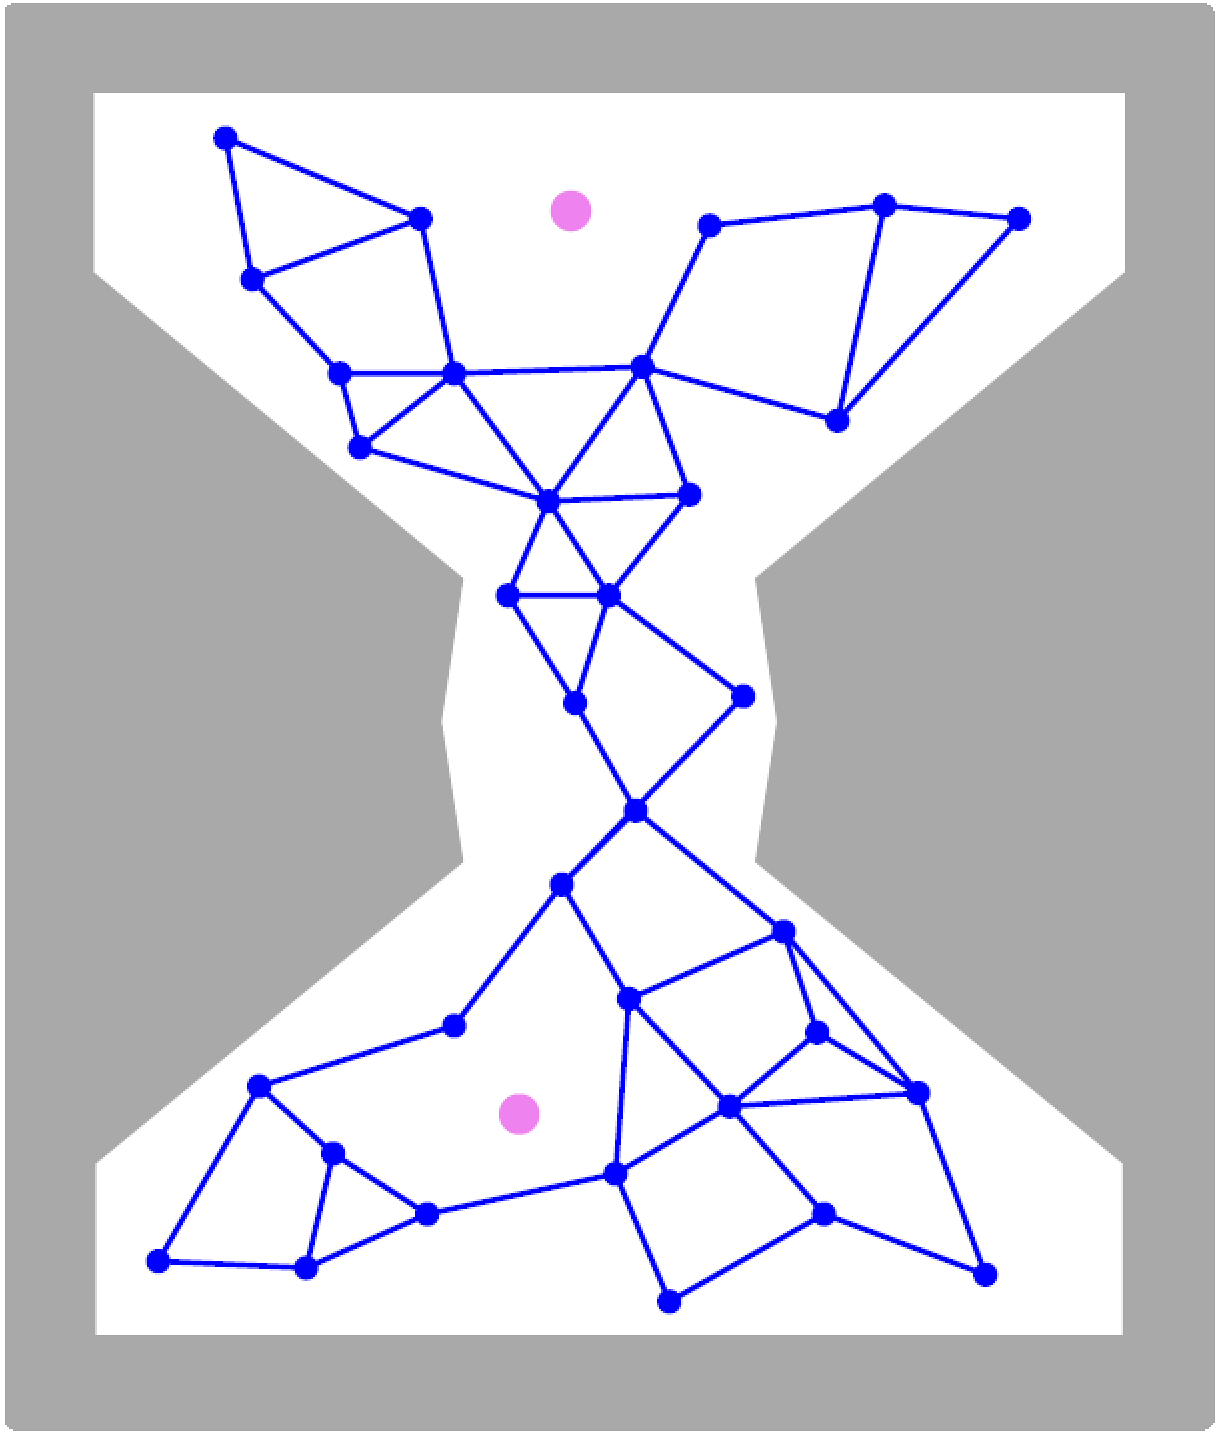

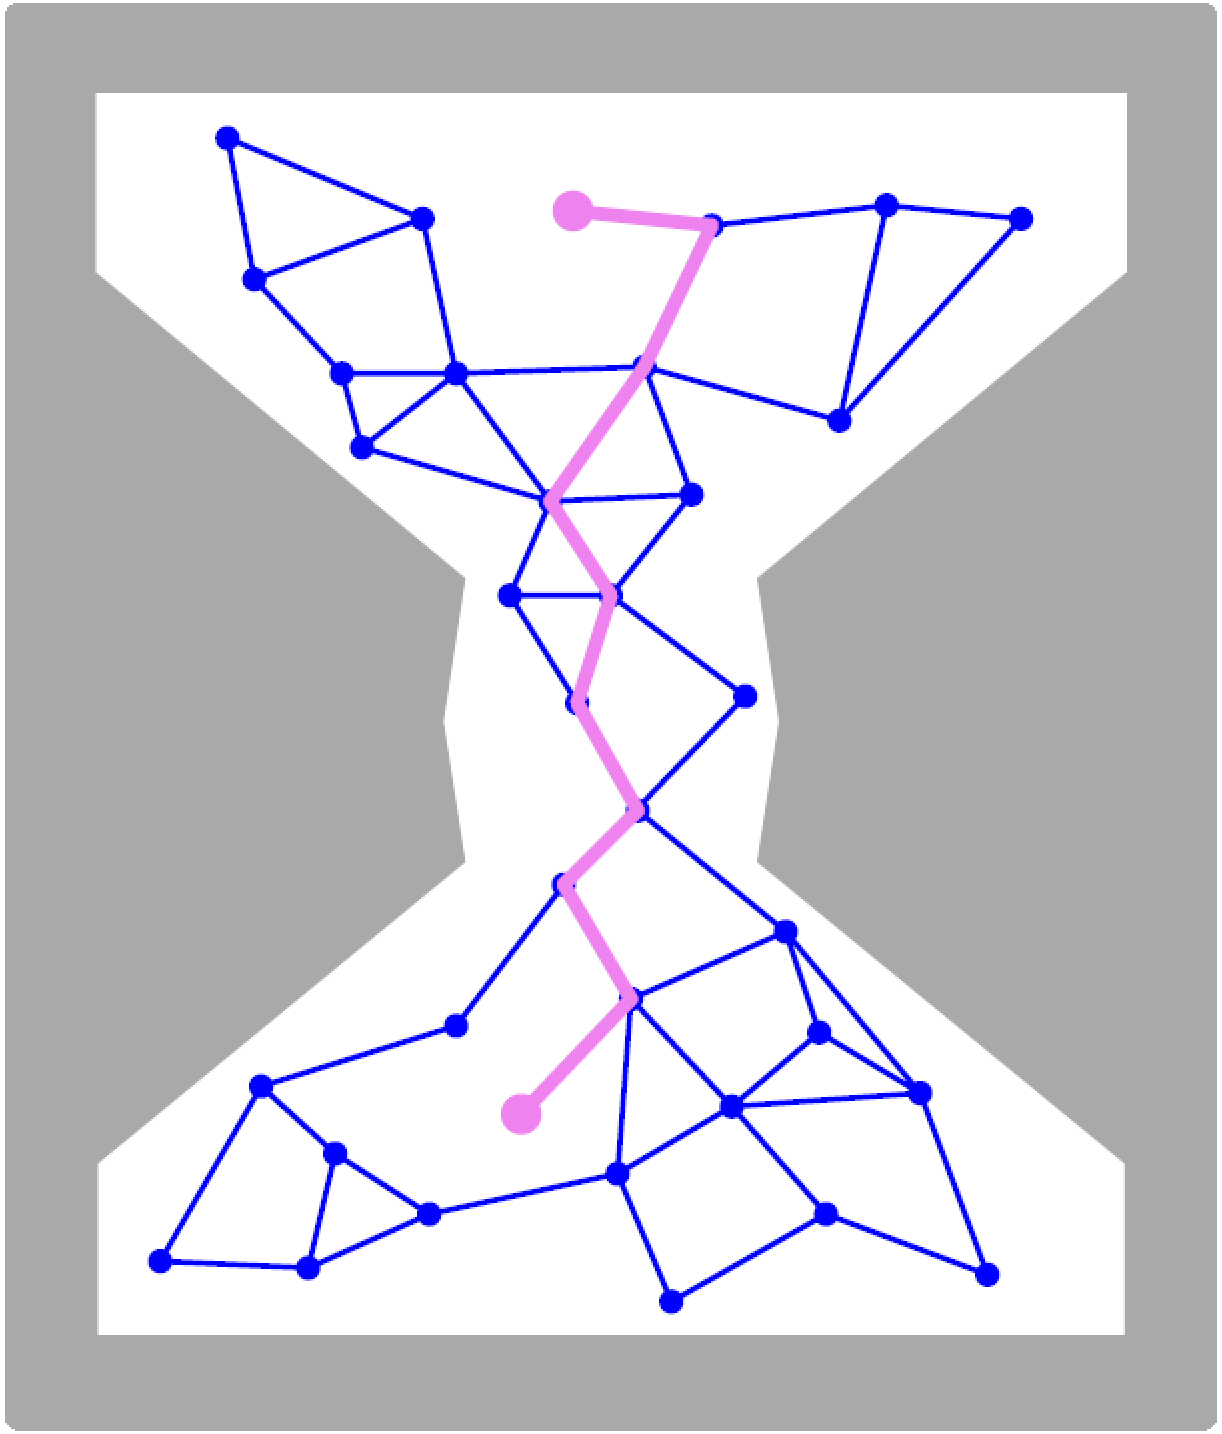

PRM, the basic method

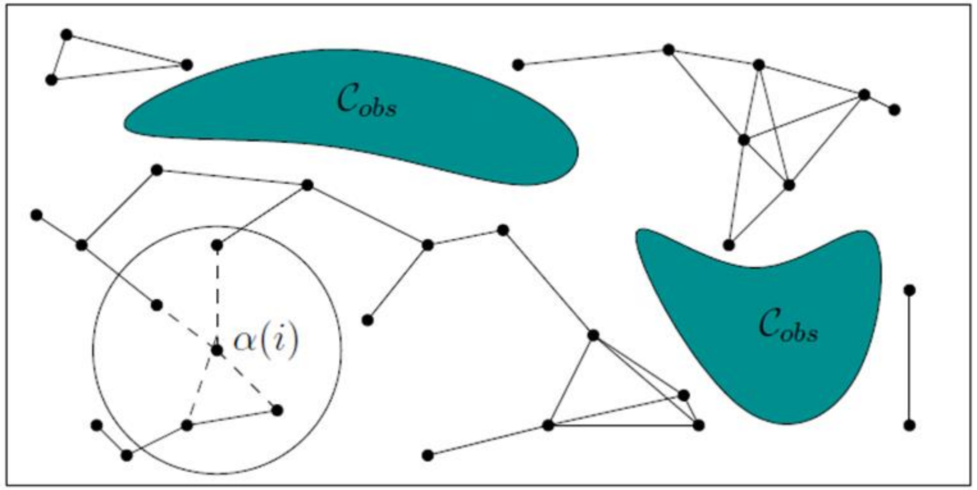

BUILD_ROADMAP

$\mathcal{G}$.init(); $i \leftarrow 0$;
while $i < N$
  if $\alpha(i) \in \mathcal{C}_{free}$ **then**
    $\mathcal{G}$.add_vertex($\alpha(i)$); $i \leftarrow i + 1$;
    for each $q \in$ NEIGHBORHOOD($\alpha(i)$,$\mathcal{G}$)
      if ((**not** $\mathcal{G}$.same_component($\alpha(i), q$)) **and** CONNECT($\alpha(i), q$)) **then**
        $\mathcal{G}$.add_edge($\alpha(i), q$);

[CHECK | DON'T CHECK]

Figure 5.25: The basic construction algorithm for sampling-based roadmaps. Note that i is not incremented if $\alpha(i)$ is in collision. This forces i to correctly count the number of vertices in the roadmap.

[LaValle's Planning Algorithms]

PRM works well in settings with many degrees of freedom

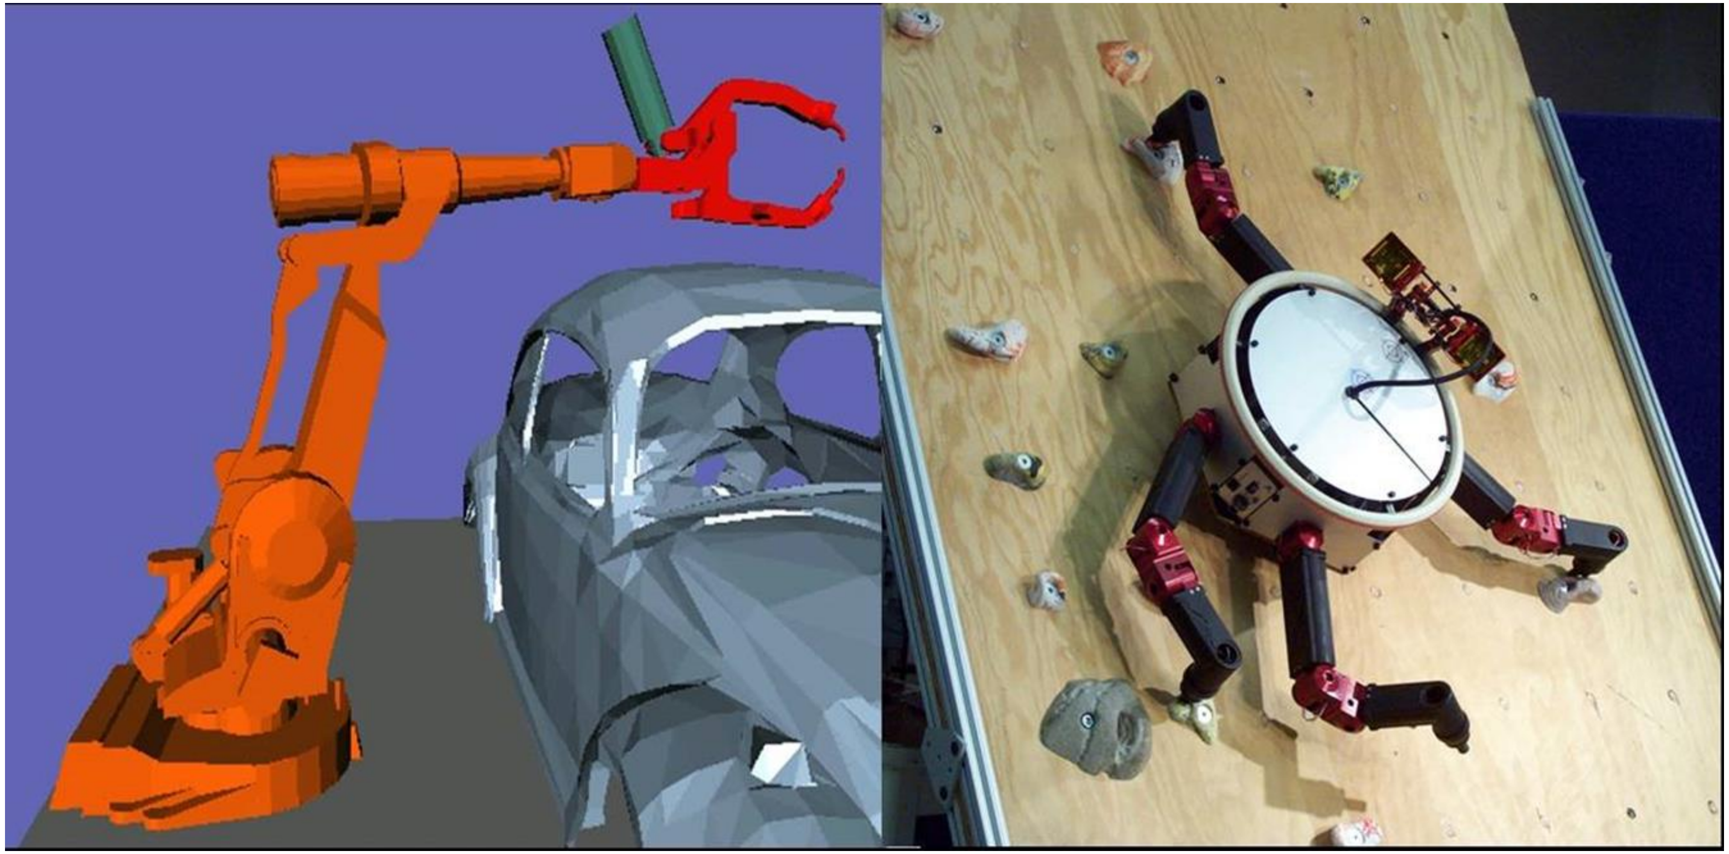

[Latombe]

Preliminary remarks

- PRM is primarily geared toward many start-to-goal queries; later on we will discuss solutions appropriate for one shot (single query) problems
- the success of PRM in practice stems from the rarity of pathological instances; it is easy to cook up examples that will make PRM fail for any variant of PRM

Ingredients

PRM, building blocks

- sampling strategy
- collision detection
- distance metric
- nearest neighbor search
- connection strategy
- local planner
- query phase: extracting path from roadmap

major computational procedures, will be discussed separately

God is in the details

- Sampling-based planners are typically very simple. However, one needs to pay attention to the details. A tiny change to one procedure may result in totally different behavior and results.
- Example I: To which existing nodes (milestones) should we connect a new sample node
- Example II: **if** ((**not** $\mathcal{G}$.same_component($\alpha(i), q$)) **and** CONNECT($\alpha(i), q$)) **then**
  $\mathcal{G}$.add_edge($\alpha(i), q$);

[CHECK | DON'T CHECK]

Sampling strategies

- the basic: uniform
- the narrow passage problem

the following slides are taken from Latombe's

Sampling and Connection Strategies for PRM Planners

Multi-Stage Strategies

Rationale:

One can use intermediate sampling results to identify regions of the free space whose connectivity is more difficult to capture

Two-Stage Sampling

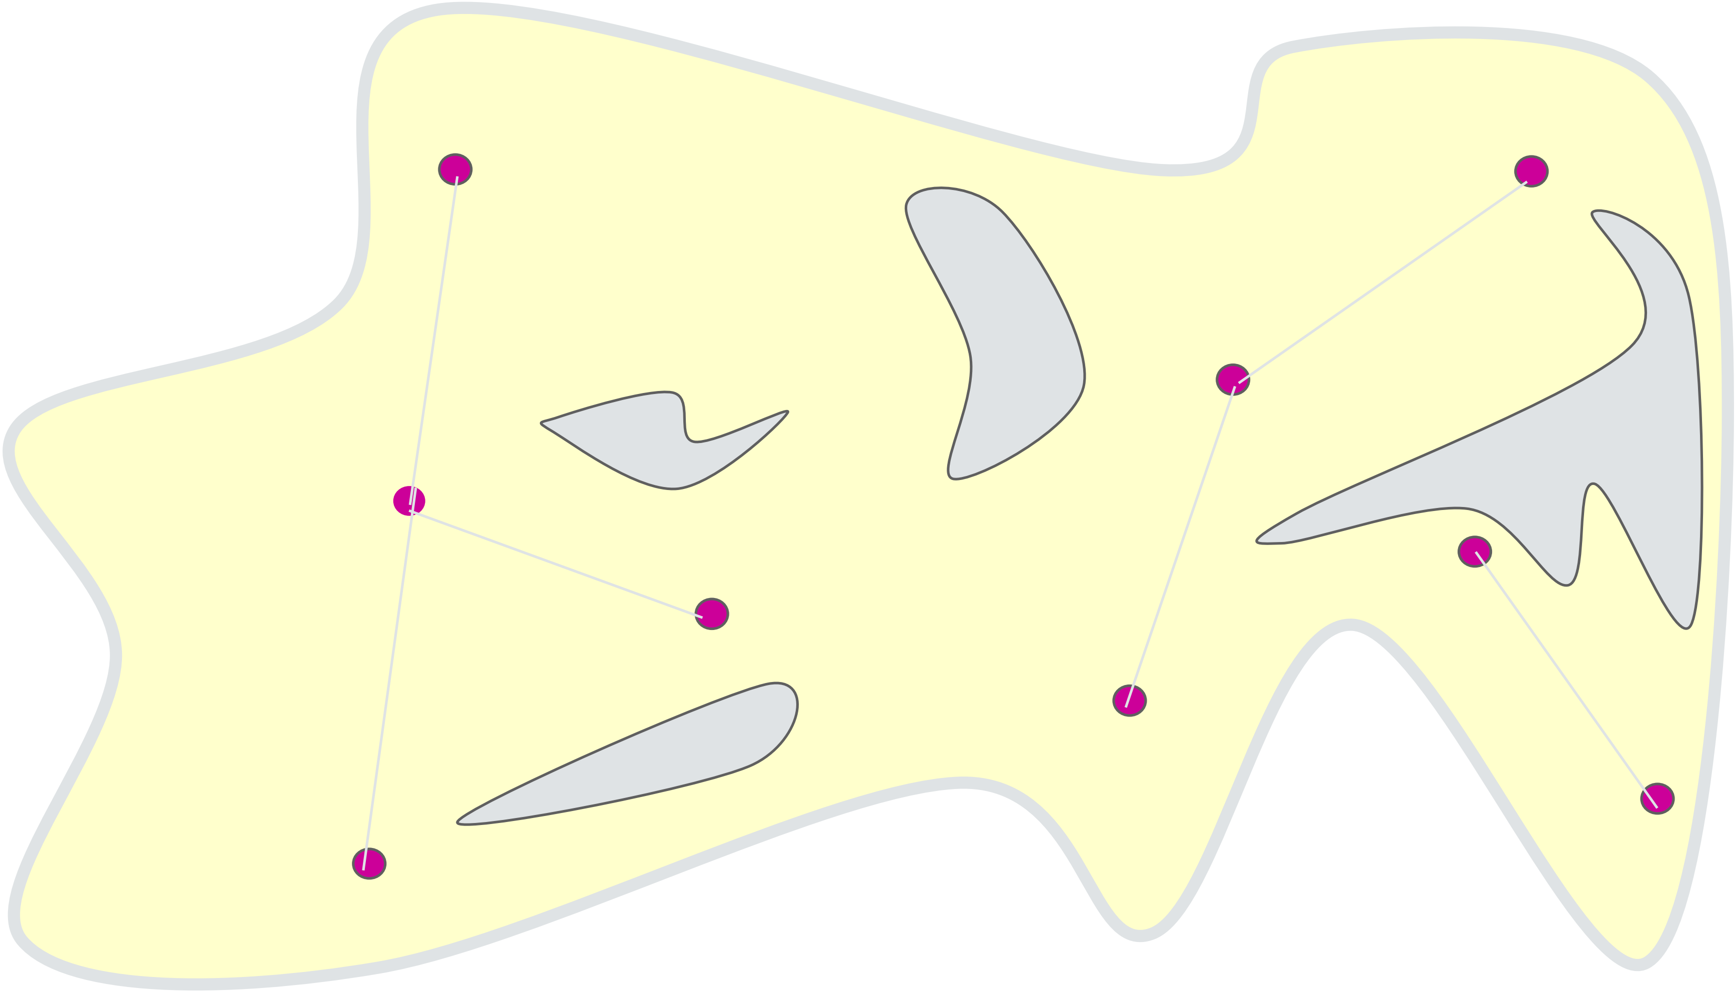

[Kavraki, 94]

Two-Stage Sampling

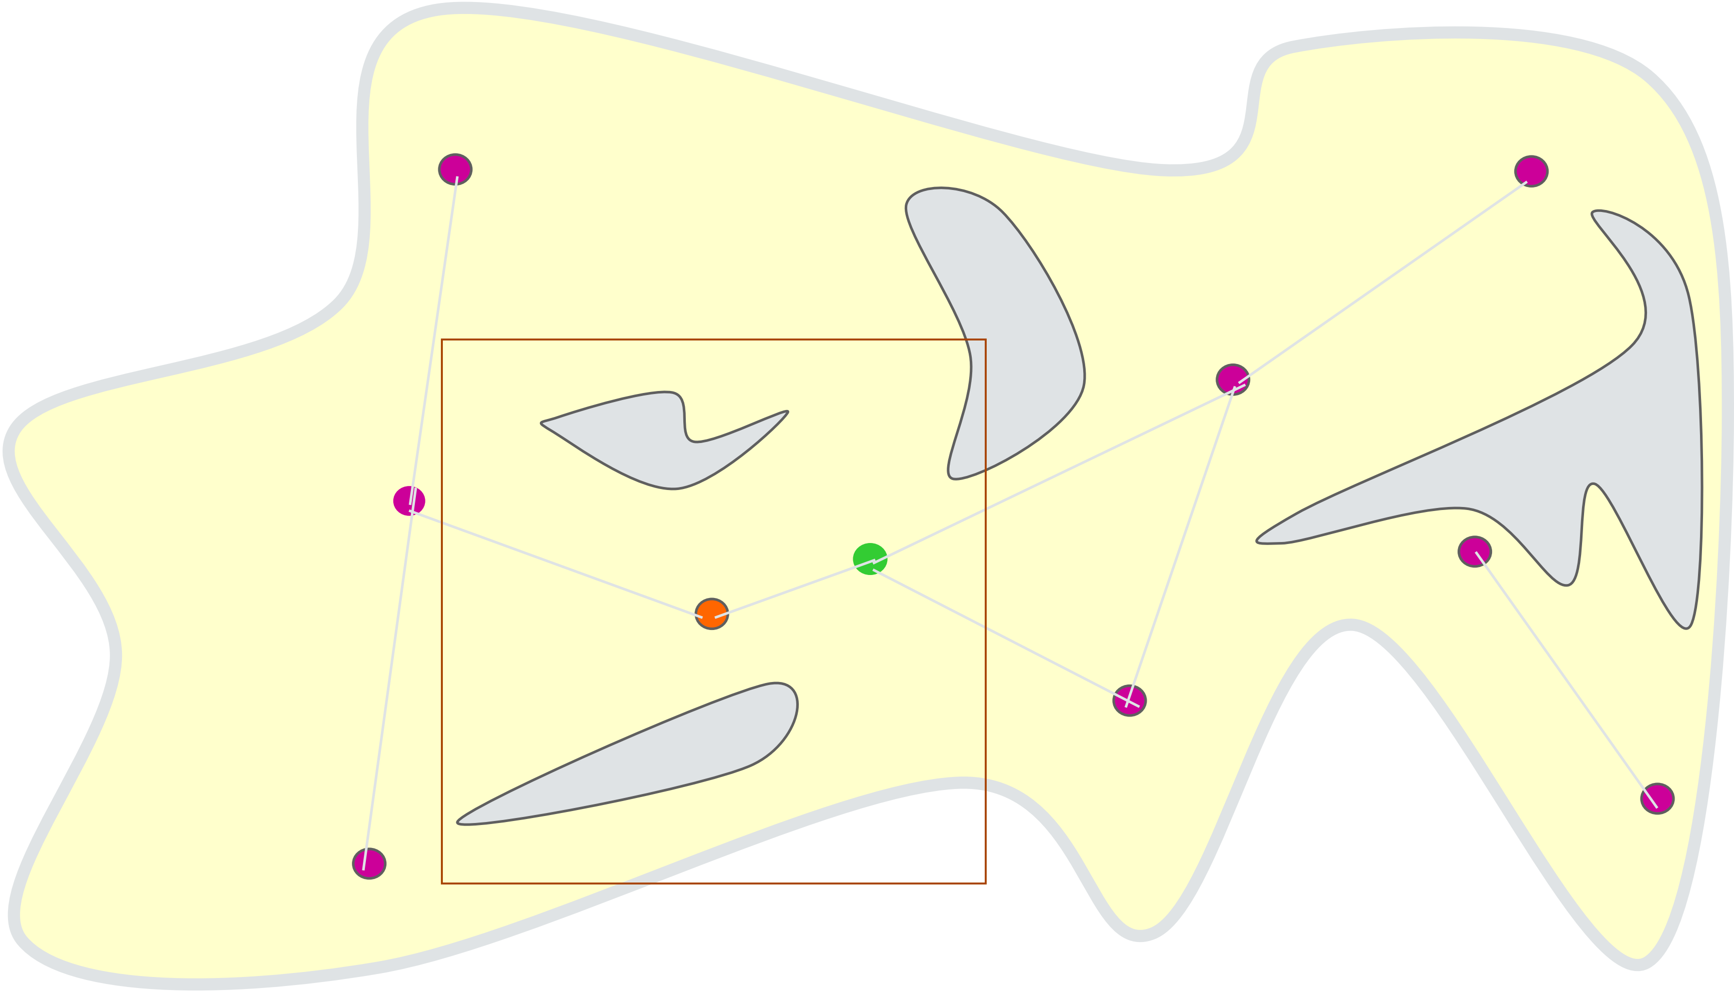

[Kavraki, 94]

Obstacle-Sensitive Strategies

Rationale:

The connectivity of free space is more difficult to capture near its boundary than in wide-open area

Obstacle-Sensitive Strategies

- Ray casting from samples in obstacles

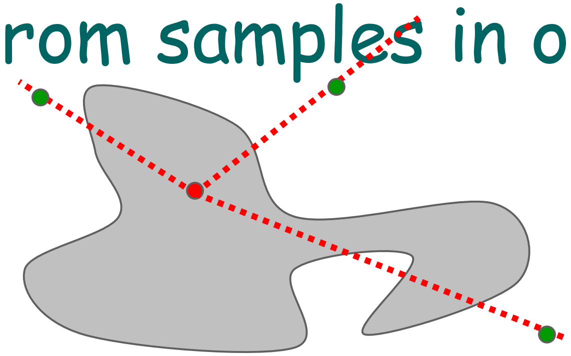

[Amato, Overmars]

- Gaussian sampling

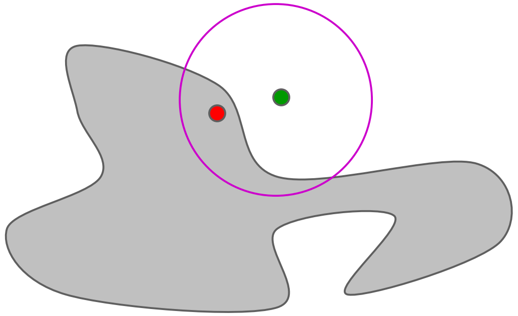

[Boor, Overmars, van der Stappen, 99]

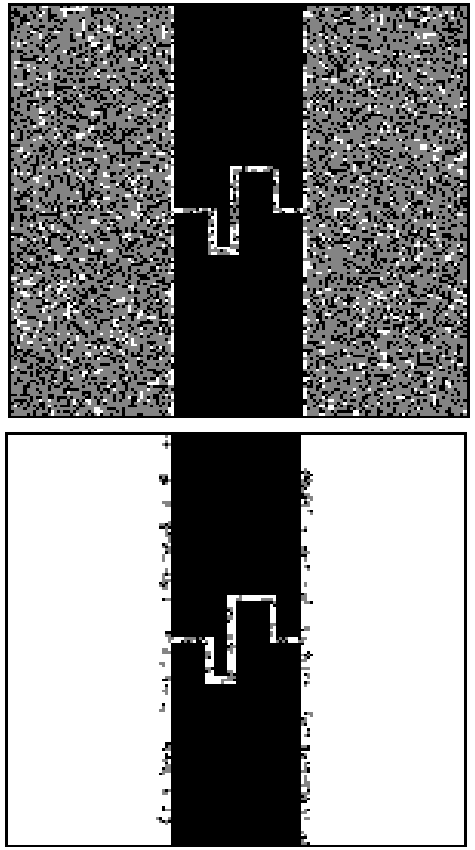

Multi-Query PRM

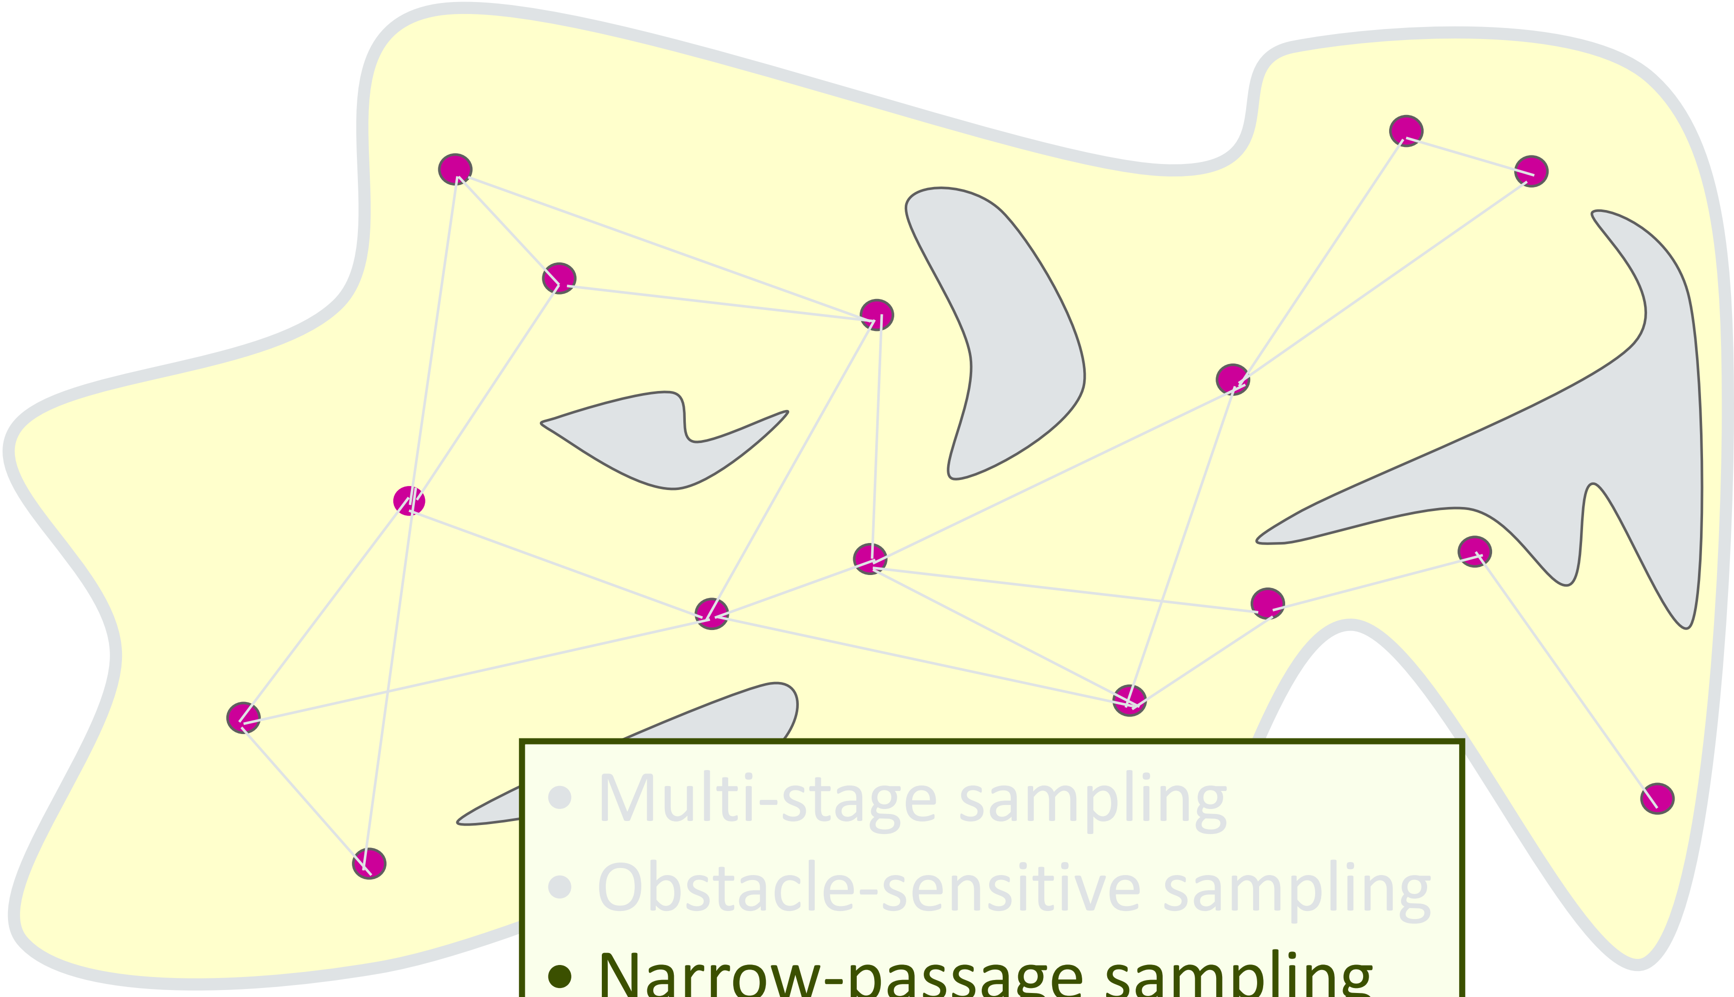

- Multi-stage sampling
- Obstacle-sensitive sampling
- Narrow-passage sampling

Narrow-Passage Strategies

Rationale:

Finding the connectivity of the free space through narrow passage is the only hard problem.

Narrow-Passage Strategies

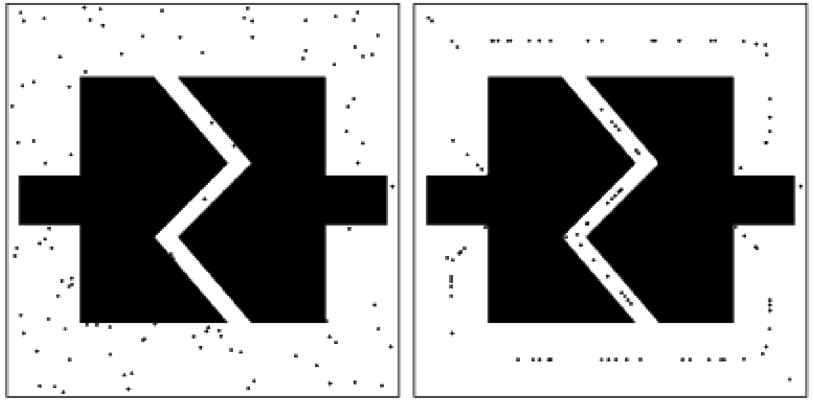

- Medial-Axis Bias

  [Amato, Kavraki]

- Dilatation/contraction of the tree space

  [Baginski, 96; Hsu et al, 98]

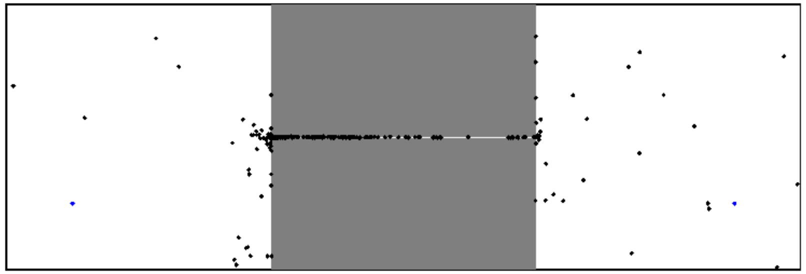

- Bridge test

  [Hsu et al, 02]

Bridge Test

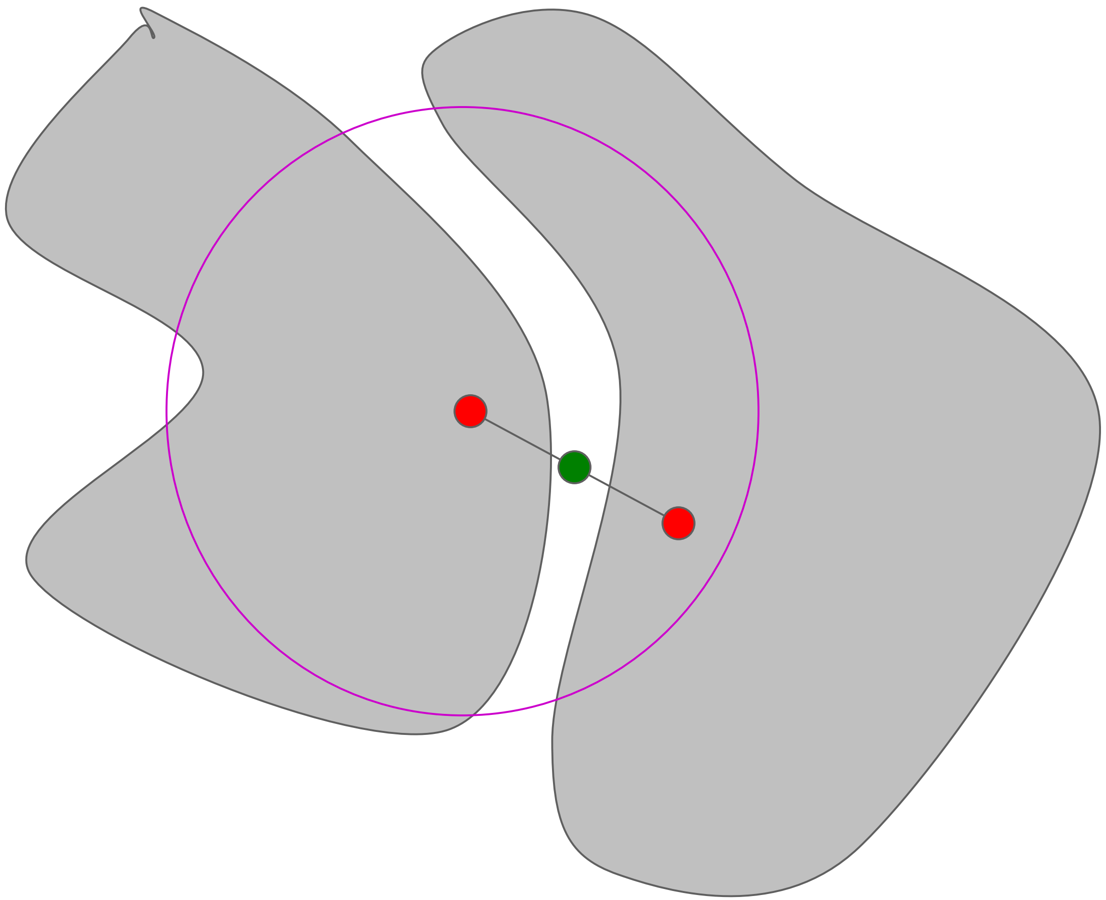

Comparison with Gaussian Strategy

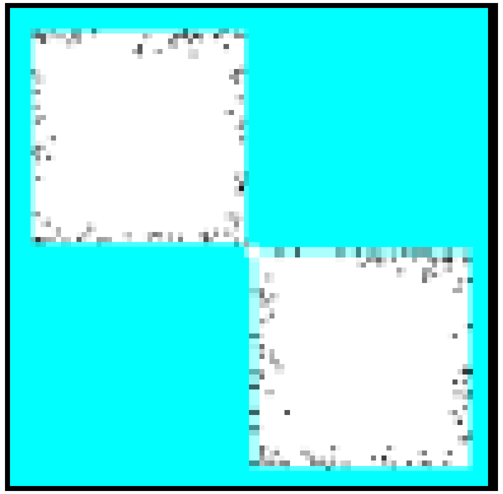

Gaussian

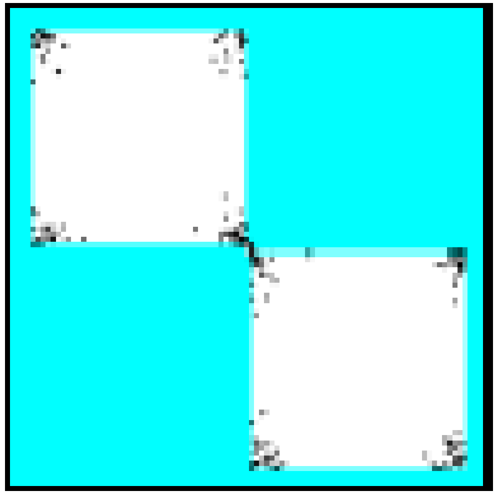

Bridge test

Sampling strategies, deterministic (quasirandom)

- discrepancy
- dispersion
- Van der Corput (1D)
- Halton (dD)
- Hammersley, when the number of samples is known in advance

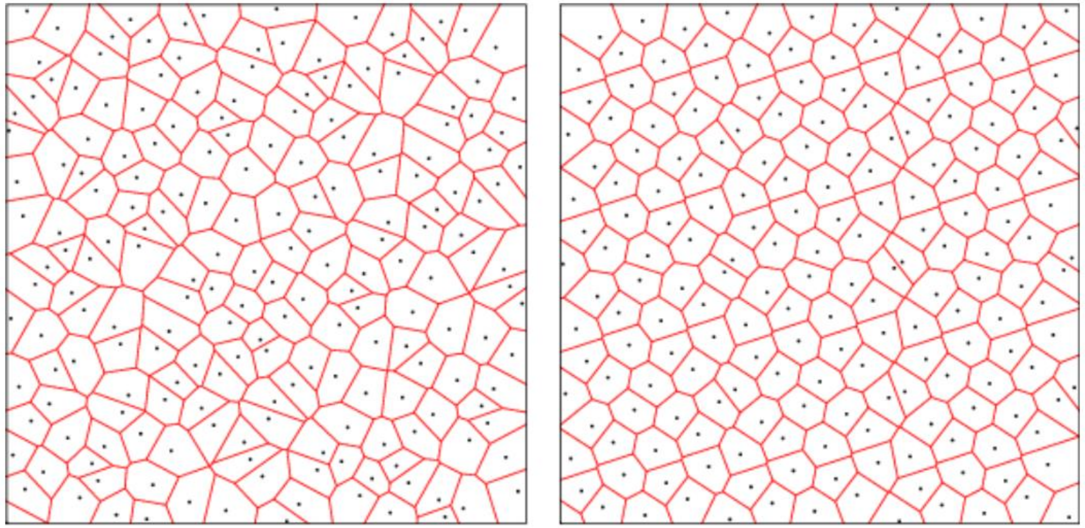

(a) 196 Halton points (b) 196 Hammersley points

[LaValle's Planning Algorithms]

Distance metric

- the distance $d(c_1,c_2)$ between two configurations
- should reflect the likelihood that the local planner will fail to connect c_1 and c_2
- can have drastic effect on the performance of the planner
- embedding to Euclidean space and using the Euclidean distance $d(c_1,c_2)= d_e(emb(c_1),emb(c_2))$
- for rigid motions, separate translation from rotation and combine them with weights w_t, w_r
  - consider Ex 3.2(c)
  - coverage in Kuffner's ICRA 04 paper "Effective sampling and distance metrics for 3D rigid body path planning

Coarse Connnections

[Latombe]

Methods:

1. Connect only pairs of milestones that are not too far apart
2. Connect each milestone to at most **k** other milestones
3. Connect two milestones only if they are in two distinct components of the current roadmap (→ the roadmap is a collection of acyclic graph)
4. Visibility-based roadmap: Keep a new milestone **m** if:
   a) **m** cannot be connected to any previous milestone <u>and</u>
   b) **m** can be connected to 2 previous milestones belonging to distinct components of the roadmap

[Laumond and Simeon, 01]

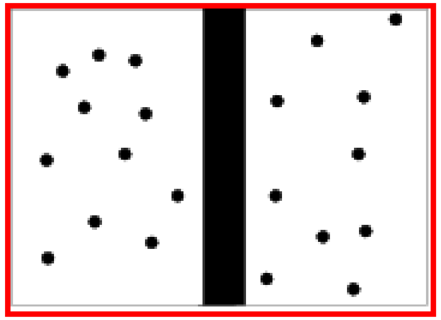

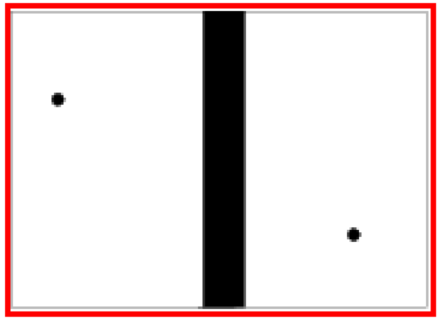

Connection strategies, remarks

- using forests only may hamper path quality (will be discussed in future lesson)
- connecting connected components: as shown earlier in importance sampling, or by employing strong local planners
- lazy evaluation

Local planner

- when simple, easy to apply many times
- if non-deterministic, needs to be saved with the roadmap edge: a storage bottleneck
- typically: dense sampling of C-space line segments between configuration
- two sub-strategies: incremental vs subdivision; none superior in all settings
- ideally conservative: if connection made, the local path is free

Query phase

- connect s and g to the roadmap
- we assume that the edges of the roadmap are given weights according to some optimization criterion
- run Dijkstra, A*, or similar: guarantees quality of path in the graph/roadmap, not necessarily in the free C-space
- perform smoothing/shortcutting
- (we will later see ways to improve path quality)

Analysis

Probabilistic completeness

- A motion planner is said to be complete if the planner in finite time either produces a solution or correctly reports that there is none. Most complete algorithms are geometry-based. The performance of a complete planner is assessed by its computational complexity.
- Probabilistic completeness is the property that as more "work" is performed, the probability that the planner fails to find a path, if one exists, asymptotically approaches zero. Several sample-based methods are probabilistically complete. The performance of a probabilistically complete planner is measured by the rate of convergence.

[Wikipedia]

Probabilistic Completeness of the basic PRM

- The C-space is $[0,1]^d$ in Euclidean space R^d
- F: the free space
- s and t: free start and target configurations
- γ: free path from s to t
- ρ: the clearance of γ
- μ: measure (volume) in R^d
- $B_1(\cdot)$: the unit ball in R^d

Probabilistic Completeness of PRM, cont'd

Theorem:

The probability that s-PRM will find a path between s and t after generating n milestones is given by

Pr[(s,t)SUCCESS] = 1-Pr[(s,t)FAILURE] $\geq$ 1-ceiling$(\frac{2L}{\rho})e^{-\sigma\rho^{d}n}$,

where $\sigma = \frac{\mu(B_1(\cdot))}{2^{d}\,\mu(F)}$.

References

Good starting point

Sampling-based algorithms

Chapter 7 of the book

Principles of robot motion:

theory, algorithms, and implementation

by Choset et al

The MIT Press

2005

comprehensive survey with many references

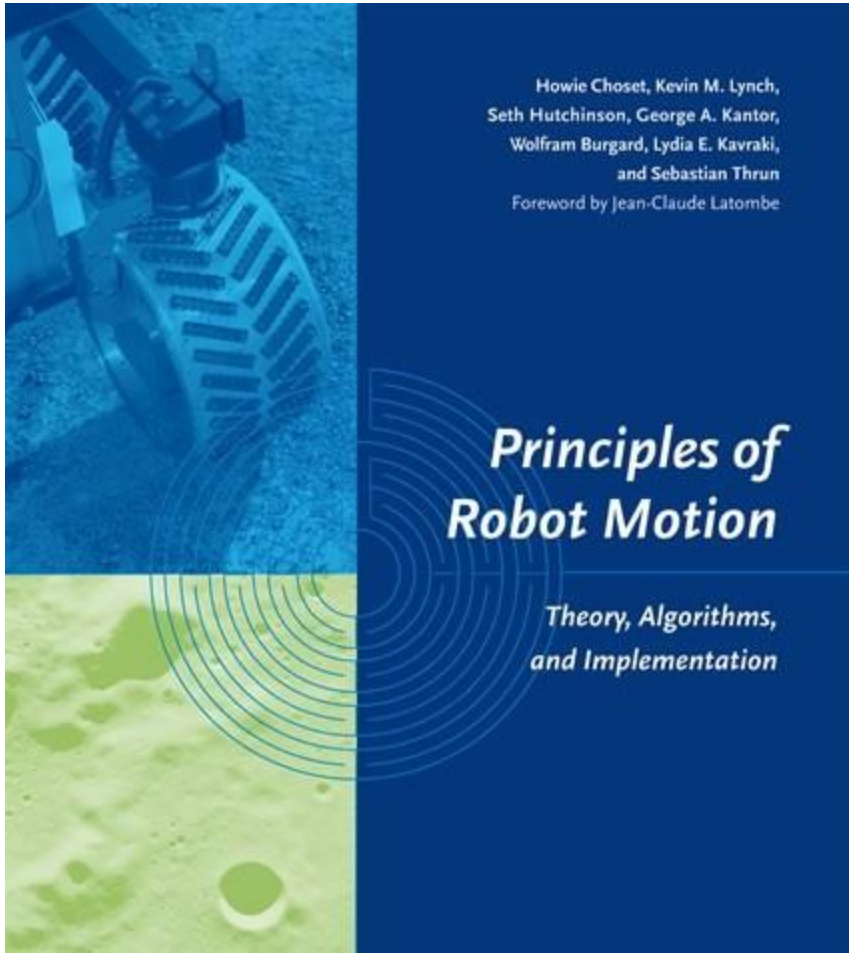

the book

Planning Algorithms

By Steven LaValle

Camrdige University Press, 2006

in-depth coverage of motion planning

available online for free!

http://planning.cs.uiuc.edu/

online bibliography

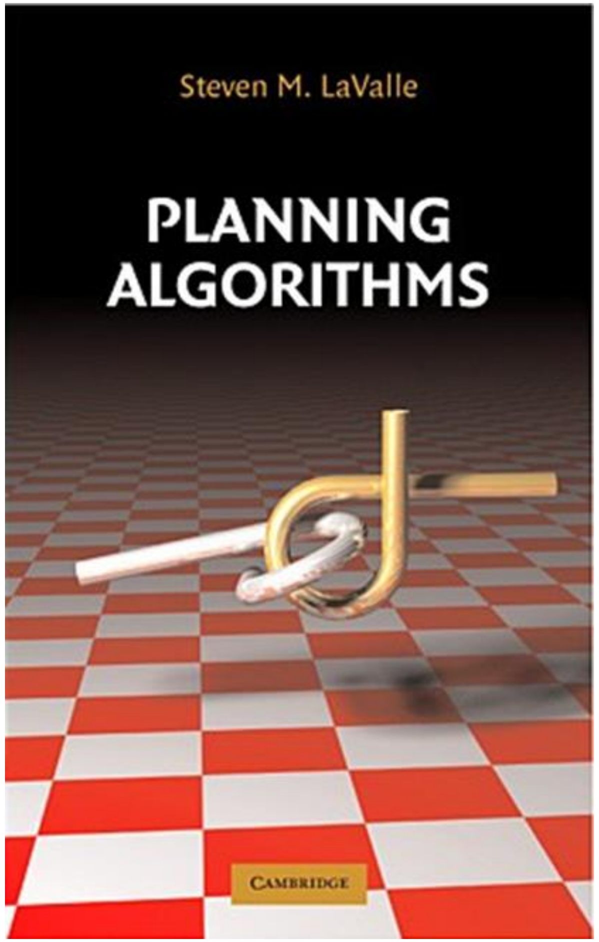

More recent surveys

- Sampling-Based Robot Motion Planning, Oren Salzman, Communications of the ACM, October 2019
- Robotics, Halperin, Kavraki, Solovey, in Handbook of Computational Geometry, 3rd Edition, 2018
- Sampling-Based Robot Motion Planning: A Review, Elbanhawi and Simic, IEEE Access, 2014 (free online)

THE END