

Motion planning: sampling-based planners III

Vojtěch Vonásek

Department of Cybernetics
Faculty of Electrical Engineering
Czech Technical University in Prague

- One may consider sampling-based planning as a “magic” tool
... but that’s not true at all!

Sampling-based planners have many issues

- Narrow passage problem
 - Difficulty of sampling small region in $\mathcal{C}_{\text{free}}$ surrounded by \mathcal{C}_{obs}
 - Problematic if (all) solutions have to pass that region
- Sensitivity to metric & parameters
 - How to measure distance in \mathcal{C} ?
 - Selecting a good metric is as difficult as motion planning!
 - Many methods have “too many” parameters
 - Some parameters are hidden (or not well described)
 - How to tune the parameters?
- Supporting functions
 - Collision detection & nearest-neighbor search
 - Fast and reliable implementation


How do we recognize the issue? → performance measurement!

Narrow passage (NP)


- A region $\mathcal{R} \subseteq \mathcal{C}_{\text{free}}$ with a small volume $vol(\mathcal{R}) < vol(\mathcal{C})$
- Probability that a random sample falls to \mathcal{R} is $\sim vol(\mathcal{R})/vol(\mathcal{C})$
- NP are problematic if their removal changes connectivity of $\mathcal{C}_{\text{free}}$
- NP are regions in $\mathcal{C} \rightarrow$ they are given implicitly
- Location/size/volume/shape of NPs is not known!

Consequences of having NP


- PRM builds unconnected roadmaps \rightarrow no solution
- RRT/EST cannot enter NP \rightarrow no solution
- Number of samples must be significantly increased
- Runtime is increased




narrow passage (NP)




PRM & NP



RRT/EST & NP








- Narrow passages are in \mathcal{C}
- Sometimes, we cannot (easily) see/estimate them from workspace!
- What makes the narrow passage in the Alpha-puzzle benchmark?

How does \mathcal{C}_{obs} appears?

- Can we guess shape of \mathcal{C}_{obs} based on workspace?
- $\text{vol}(\mathcal{A}) \ll \text{vol}(\mathcal{O})$




Workspace



Configuration space

- $\text{vol}(\mathcal{A}) < \text{vol}(\mathcal{O})$



Workspace



Configuration space


- When obstacles \mathcal{O} dominate, they mostly influence the shape of \mathcal{C}_{obs}

- Let $X, Y \subset \mathbb{R}^n$, X and Y are nonempty
- Brunn-Minkowski theorem:


$$\text{vol}(X \oplus Y) \geq (\text{vol}(X)^{\frac{1}{n}} + \text{vol}(Y)^{\frac{1}{n}})^n$$

- $\text{vol}(\mathcal{C}_{\text{obs}})$ is larger than $\min(\text{vol}(\mathcal{A}), \text{vol}(\mathcal{O}))$
- $\text{vol}(\mathcal{C}_{\text{obs}})$ can be much larger!

Example: $\text{vol}(\mathcal{A}) = \text{vol}(\mathcal{O})$



Workspace



Configuration space

- To cope with the narrow passage problem, improve path quality, speed-up planning, to enable planning in specific cases

Main tricks

- Change distribution of random samples
- Dedicated metrics
- Improved nearest-neighbor search
- Use suitable local planners
- Improve collision-detection

```
1 initialize tree  $\mathcal{T}$  with  $q_{init}$ 
2 for  $i = 1, \dots, l_{max}$  do
3    $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
4    $q_{near} = \text{find nearest node in } \mathcal{T} \text{ towards}$ 
    $q_{rand}$ 
5    $q_{new} = \text{localPlanner from } q_{near} \text{ towards}$ 
    $q_{rand}$ 
6   if  $\text{canConnect}(q_{near}, q_{new})$  then
7      $\mathcal{T}.\text{addNode}(q_{new})$ 
8      $\mathcal{T}.\text{addEdge}(q_{near}, q_{new})$ 
9     if  $\rho(q_{new}, q_{goal}) < d_{goal}$  then
10      return path from  $q_{init}$  to  $q_{new}$ 
```

- Many existing modifications of sampling-based planners¹²³⁴⁵

¹L. Zhang et al. “Motion Planning for Robotics: A Review for Sampling-Based Planners”. In: *Biomimetic Intelligence and Robotics* 5.1 (Mar. 2025), p. 100207. DOI: [10.1016/j.birob.2024.100207](https://doi.org/10.1016/j.birob.2024.100207).

²T. McMahan et al. “A survey on the integration of machine learning with sampling-based motion planning”. In: *Foundations and Trends® in Robotics* 9.4 (2022), pp. 266–327.

³X. Xiao et al. “Motion planning and control for mobile robot navigation using machine learning: a survey”. In: *Autonomous Robots* 46.5 (2022), pp. 569–597. DOI: <https://doi.org/10.1007/s10514-022-10039-8>.

⁴J. Wang et al. “A survey of learning-based robot motion planning”. In: *IET Cyber-Systems and Robotics* 3.4 (2021), pp. 302–314. DOI: <https://doi.org/10.1049/csy2.12020>.


⁵L. G. D. O. Várás et al. “Systematic Literature Review of Sampling Process in Rapidly-Exploring Random Trees”. In: *IEEE Access* 7 (2019), pp. 50933–50953. DOI: [10.1109/ACCESS.2019.2908100](https://doi.org/10.1109/ACCESS.2019.2908100).

Observation

- RRT tree grows towards random samples
- If we sample some region more dense, the tree is “attracted” to grow there

Goal-bias

- Random sample q_{rand} is generated in \mathcal{C} with probability $(1 - p_{goal})$, otherwise it is set to $q_{rand} = q_{goal}$
- The rest of RRT algorithm is the same
- Improves the performance if the tree can directly reach the goal
- Decreases the performance if the tree is hindered by obstacles




Observation





- RRT tree grows towards random samples
- If we sample some region more dense, the tree is “attracted” to grow there

Goal-bias

- Random sample q_{rand} is generated in \mathcal{C} with probability $(1 - p_{\text{goal}})$, otherwise it is set to $q_{\text{rand}} = q_{\text{goal}}$
- The rest of RRT algorithm is the same
- Improves the performance if the tree can directly reach the goal
- Decreases the performance if the tree is hindered by obstacles




- Goal-bias may improve or even worsen the performance!



Observation


- Goal-bias attracts the tree towards q_{goal} , but the tree may be blocked by obstacles
- Generalization: we can attract the tree toward any region $\mathcal{R} \subseteq \mathcal{C}$ if we sample \mathcal{R} densely




RRT + goal-bias

Guided-based sampling


- Estimate a path that can “guide” the tree in the \mathcal{C} -space
- Generate q_{rand} around the path-waypoints (starting from first waypoint) until the tree reaches the waypoint
- Then generate q_{rand} around the next waypoint




Guiding path




Sampling at 1





Sampling at 2



Sampling at 3



Sampling at 4



How to compute the guiding path?



- Generally, the guiding path has to be located in \mathcal{C} !!
- Finding a good guiding path has the same complexity as the original planning problem!
- (i.e., guiding sampling is 'planning solved by planning')
- Practically, we have two options

Guiding path in \mathcal{W}

- Path is computed in workspace — geometric planning (Voronoi diagram, Visibility graph, etc.)
- Suitable for low-dimensional problems
- The remaining dimensions are sampled uniformly

Guiding path in \mathcal{C}

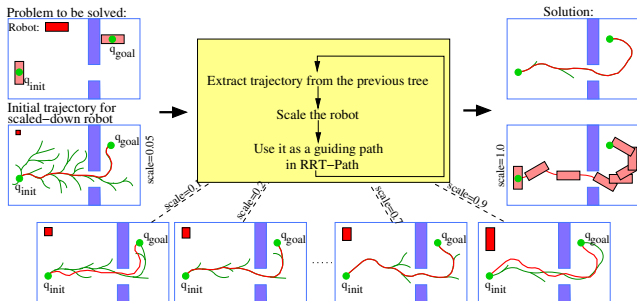
- Path is computed in \mathcal{C} by a simplified search




Guiding path in \mathcal{W}


$q = (x, y, \varphi)$
 (x, y) from the path
 φ randomly


- Problem is simplified — relaxation of constraints
- For example, robot is scaled-down
- Solve simplified planning problem
- Use the solution to generate random samples along it
- The process can be iterative





The final solution for the original robot











RRT modification: bidirectional search

- Use two trees: \mathcal{T}_i rooted at q_{init} , \mathcal{T}_g rooted at q_{goal}
- One tree expands towards q_{rand} , second tree expands towards q_{new} of the first tree

```
1  $\mathcal{T}_i.addNode(q_{init})$ 
2  $\mathcal{T}_g.addNode(q_{goal})$ 
3 for  $i = 1, \dots, I_{max}$  do
4      $q_{rand} = \text{generate randomly in } \mathcal{C}$ 
5      $q_{near} = \text{find nearest node in } \mathcal{T}_i \text{ towards } q_{rand}$ 
6      $q_{new} = \text{localPlanner from } q_{near} \text{ towards } q_{rand}$ 
7     if  $canConnect(q_{near}, q_{new})$  then
8          $\mathcal{T}_i.addNode(q_{new})$ 
9          $\mathcal{T}_i.addEdge(q_{near}, q_{new})$ 
10         $q'_{near} = \text{find nearest node in } \mathcal{T}_g \text{ towards } q_{new}$ 
11         $q'_{new} = \text{localPlanner from } q_{near} \text{ towards } q_{rand}$ 
12        if  $canConnect(q'_{near}, q'_{new})$  then
13             $\mathcal{T}_g.addNode(q_{new})$ 
14             $\mathcal{T}_g.addEdge(q_{near}, q_{new})$ 
15            if  $canConnect(q'_{new}, q_{new})$  then
16                joint trees
17                return path from  $q_{init}$  to  $q_{goal}$ 
18     $\mathcal{T}_i, \mathcal{T}_g = \mathcal{T}_g, \mathcal{T}_i$  // swap trees
```



- Use two trees: \mathcal{T}_i rooted at q_{init} , \mathcal{T}_g rooted q_{goal}
- One tree expands towards q_{rand} , second tree expands towards q_{new} of the first tree
- Helps to enter narrow passages (sometimes)
- Connection of two trees
 - Computationally intensive
 - To speed up, performs only if $\varrho(q_{new}, q'_{new})$ is small enough
 - Difficult if motion model/constraints have to be considered
- Balanced trees: swap trees if $|\mathcal{T}_i| > |\mathcal{T}_g|$



Original PRM/sPRM

- Uniform sampling $q \sim U(\mathcal{C})$

Gaussian sampling: two-samples

- Uniform sample $q_1 \sim U(\mathcal{C})$, then another sample $q_2 \sim N(q, \Sigma)$ (around q_1 from Gaussian distribution)
- Ignore if $q_1, q_2 \in \mathcal{C}_{\text{free}}$ or $q_1, q_2 \in \mathcal{C}_{\text{obs}}$, otherwise
- add the collision-free one to the roadmap
- Generates the random samples near \mathcal{C}_{obs} only!

Gaussian + uniform


- Combination of two previous methods
- More dense sampling around \mathcal{C}_{obs} than basic PRM

Bridge test


- Generate q_1 and q_2 using the Gaussian method
- Determine the midpoint q' on the line segment $|q_1, q_2|$
- Use q' if $q' \in \mathcal{C}_{\text{free}}$ and $q_1, q_2 \in \mathcal{C}_{\text{obs}}$




Uniform



Gaussian



Gaussian + Uniform



Bridge-Test

Main idea: postpone time-consuming routines until they are needed

Lazy sampling-based planning⁶⁷⁸⁹

- Postpone collision detection on edges/configurations
- Approximate collision detection on edges
- Do exact collision detection only when the edge/configuration is part of a solution

⁶A. Mandalika et al. “Generalized Lazy Search for Robot Motion Planning: Interleaving Search and Edge Evaluation via Event-based Toggles”. In: *CoRR* abs/1904.02795 (2019). arXiv: 1904.02795.


⁷J. D. Hernández et al. “Lazy evaluation of goal specifications guided by motion planning”. In: *2019 International Conference on Robotics and Automation*. 2019, pp. 944–950.

⁸K. Hauser. “Lazy collision checking in asymptotically-optimal motion planning”. In: *IEEE International Conference on Robotics and Automation*. 2015, pp. 2951–2957. DOI: 10.1109/ICRA.2015.7139603.

⁹J. Denny et al. “Lazy toggle prm: A single-query approach to motion planning”. In: *2013 IEEE International Conference on Robotics and Automation*. IEEE. 2013, pp. 2407–2414.


Lazy PRM¹⁰

- Build PRM roadmap without collision detection of edges
- Find path in it, check collisions only for path edges
- Recalculate path if some edge is not collision free
- If no path is found, extend the roadmap by new samples/edges; repeat
- Otherwise, the path is collision-free




- Faster planning in certain scenarios, but not always!

¹⁰R. Bohlin and L. E. Kavraki. "Path planning using lazy PRM". In: *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*. Vol. 1. IEEE, 2000, pp. 521–528.



Batch informed Tree (BIT*)¹¹ — most advanced planner


- Samples are added to the tree in batches
- Batch of samples is processed using priority queue (using the estimated cost of solution)
- Batches contain multiple samples drawn using Informed-RRT* principle
- Edges are evaluated in lazy manner: and pruned if they do not improve solution
- Only promising edges are checked for collisions



¹¹J. D. Gammell et al. "Batch Informed Trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 2015, pp. 3067–3074. DOI: [10.1109/ICRA.2015.7139620](https://doi.org/10.1109/ICRA.2015.7139620).

Batch informed Tree (BIT*)¹¹ — most advanced planner

- Samples are added to the tree in batches
- Batch of samples is processed using priority queue (using the estimated cost of solution)
- Batches contain multiple samples drawn using Informed-RRT* principle
- Edges are evaluated in lazy manner: and pruned if they do not improve solution
- Only promising edges are checked for collisions




¹¹J. D. Gammell et al. "Batch Informed Trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 2015, pp. 3067–3074. DOI:


[10.1109/ICRA.2015.7139620](https://doi.org/10.1109/ICRA.2015.7139620).

Batch informed Tree (BIT*)¹¹ — most advanced planner


- Samples are added to the tree in batches
- Batch of samples is processed using priority queue (using the estimated cost of solution)
- Batches contain multiple samples drawn using Informed-RRT* principle
- Edges are evaluated in lazy manner: and pruned if they do not improve solution
- Only promising edges are checked for collisions




(a) \mathbb{R}^2 : Runs solved vs. time



(b) \mathbb{R}^8 : Runs solved vs. time



(c) \mathbb{R}^2 : Median solution cost vs. time



(d) \mathbb{R}^8 : Median solution cost vs. time



