

3. Approximate path planning

Path Planning Approaches


- Combinatorial / geometric planers
 - **Exact: Compute C-free geometrically**
 - Comments
 - This gives **complete** planners
 - Works beautifully in 2D and for some simple cases in 3D
 - Worst-case bound for combinatorial complexity of C-objects in 3D is high
 - A complete planner in 3D runs in $O(2^{n^{\#dof}})$
 - **Impractical** for high #dof
- Approximate planners



Approximate path planning

Knowing C-obstacles is like having a map: You know the roads and you can make a plan on how to get to the goal.

Without knowing the C-obstacles you are in the dark.




Approximate path planning


- Idea: Approximate C-free
- Approaches
 - Space partitioning/grid-based planners with A*
 - and variants (weighted A*, D*, ARA*,...)
 - Sampling-based
 - Rapidly-Exploring Random Tree (RRT)
 - Probabilistic RoadMap (PRM)
 - Potential field planners
 - Hybrid methods combining ideas from all of the above

Grid-based planners with A*

Grid based planners “pixelize” the space



Let's say we have a robot moving in 2d without rotation and we want to implement a grid-based planner.




We decide what resolution we want on each axis, and from here we get the size of the grid in that dimension.

$$K_x = \frac{X_{max}}{\delta_x}$$

C-space=(x, y)

C-space sampled with a grid

$$K_y = \frac{Y_{max}}{\delta_y}$$



Grid-based planners with A*

- Sample C-space with a uniform grid/lattice
 - This “pixelizes” the C-space (pixels/voxels)
- Search for a path from start to end through “free” space
 - Dijkstra/A* and variants
 - Graph is implicit, given by lattice topology: move +/-1 in each direction, possibly diagonals

Dijkstra's algorithm

- It's basically a **best**-first search
- Initialize: $dist[v] = \infty$, $dist[s] = 0$,
- Repeat: select the **best** vertex (closest to start), and relax its edges

- Data structures

- PQ of $(u, dist[u])$
- $\text{priority}(v)$: $dist[v]$

- Keeps track of :

- $dist[v]$ = cost of getting from *start* to v
- $done[u]$: true if u has been explored
- $pred[v]$: predecessor of v on the (optimal) path from *start* to v

Dijkstra(vertex s)

- initialize
 - for all v : $dist[v] = \infty$, $done[v] = \text{false}$, $pred[v] = \text{null}$
 - $dist[s] = 0$, $\text{PQ.insert}(s, dist[s])$ ← insert the start
- while PQ not empty
 - $(u, dist[u]) = \text{PQ.deleteMin}()$
 - mark u as done ← //claim: $dist[u]$ is the shortest path from *s* to u
 - for each edge (u, v) , if v not done:
 - $alt = dist[u] + \text{edge}(u, v)$
 - if $alt < dist[v]$
 - $dist[v] = alt$, $\text{PQ.decreaseKey}(v, dist[v])$

↑ requires a structure that can search, or a PQueue with additional book-keeping

On a grid-graph

Dijkstra(vertex s)

- initialize
 - for all v : $dist[v] = \infty$, $done[v] = \text{false}$, $pred[v] = \text{null}$
 - $dist[s] = 0$, $\text{PQ.insert}(s, dist[s])$
- while PQ not empty
 - $(u, dist[u]) = \text{PQ.deleteMin}()$
 - if u is done, continue
 - mark u as done
 - for each neighbor v of u if v not done and $\text{isFree}(v)$:
 - $alt = dist[u] + \text{edge}(u, v)$
 - if $alt < dist[v]$
 - $dist[v] = alt$, $\text{PQ.insert}(v, dist[v])$


u is a placement,
and also a pixel
on the grid

$\text{isFree}(v)$: is v in C-free


Grid graphs in 2d

$u = (x, y)$ is a placement, and also a pixel on the grid

Neighbors of $u = (x, y)$



4-connected



8-connected

Grid graphs in 3d

neighbors of $v = (x, y, \theta)$

6-connected: $(x + \delta_x, y, \theta)$ $(x - \delta_x, y, \theta)$
 $(x, y + \delta_y, \theta)$ $(x, y - \delta_y, \theta)$
 $(x, y, \theta + \delta_\theta)$ $(x, y, \theta - \delta_\theta)$

or, connected diagonally as well

Motion planners assume the existence of a collision-detection routine that can check whether a given configuration, or path segment, is in free space.

Would my robot, if placed at this point p , intersect any obstacle?

```
//return true if placement p is in C-free
```

```
//put differently, return true if placing the robot R at p does not intersect any obstacle
```

```
bool isFree(placement p, the robot, the obstacles)
```

- translate and rotate the robot to p
- check whether any edge of the robot intersects any of the obstacles





p is a point in C-space: (x, y) or (x, y, θ) or $(x, y, z, \theta_x, \theta_y, \theta_z)$ or ...

EXAMPLE



We need to write: $\text{isFree}(p = (x, y, \theta), \text{Robot } R, \text{ Obstacles } S)$

2D: robot can translate and rotate

C-space: 3D configuration p: (x, y, θ)



EXAMPLE



Dijkstra => A*

<https://qiao.github.io/PathFinding.js/visual/>

Dijkstra

- Evaluates vertices based on their distance to the start
 - $\text{priority}(v) = \text{dist}(v)$
- Dijkstra will explore a large portion of the graph before reaching the target. Would be nice if we could cut down the number of nodes traversed before reaching the goal



A*

- Idea: Steer the search towards the goal (while keeping solution optimal)
- $\text{priority}(v) : f(v) = \text{dist}(v) + h(v)$
 - $\text{dist}(v)$: cost of getting from *start* to v
 - $h(v)$: estimate of the cost from v to *goal*
- Dijkstra is A* with $h(v) = 0$

Dijkstra => A*

Animations

- <https://qiao.github.io/PathFinding.js/visual/>
- https://en.wikipedia.org/wiki/Dijkstra%27s_algorithm#/media/File:Dijkstras_progress_animation.gif
- https://www.youtube.com/watch?v=DINCL5cd_w0
- https://www.google.com/search?client=firefox-b-1-d&q=dijkstra+vs+A%2a+algorithm&tbo=vid&sa=X&ved=2ahUKEwic2fXNvMr-AhV1FFkFHcSxB_kQ0pQJegQICxAB&biw=1313&bih=731&dpr=2#fpstate=ive&vld=cid:02bef27f,vid:g024lzsxnDo
-

A*

- The heuristic $h(v)$ is called **admissible** if $h(v)$ is smaller than the true cost of getting from v to the target.
- Theorem: If $h(v)$ is admissible then A* will return an optimal solution.
- Put differently if $h(v)$ is too high, the algorithm loses optimality
 - $h(v) = 0$ will always work
 - higher $h(v)$ will steer the search more => more efficient
 - The closer $h(v)$ is to the true cost of getting from v to goal, the more efficient
 - if $h(v)$ is too high => A* not optimal
- In many situations a safe admissible heuristic is

$$h(v) = \text{EuclidianDistance}(v, goal)$$

admissible because the cost of getting from a to b is \geq the Euclidian distance from a to b

(x, y) or (x, y, θ) or $(x, y, z, \theta_x, \theta_y, \theta_z)$ or ...



generic A* (placement start, goal)

- initialize
 - for all v : $dist[v] = \infty$, $f[v] = \infty$, $pred[v] = \text{null}$, $done[v] = \text{false}$
 - $dist[s] = 0$, $f[s] = h(s, \text{target})$, $\text{PQ.insert}(s, f[s])$
- while PQ not empty
 - $\langle u, f[u] \rangle = \text{PQ.deleteMin}()$
 - for each neighbor v , if not $done[v]$ and $\text{isFree}(v)$
 - $alt = dist[u] + \text{edge}(u, v)$
 - if $alt < dist[v]$
 - $dist[v] = alt$; $f[v] = dist[v] + h(v, \text{target})$;
 - $pred[v] = u$; $\text{PQ.decreaseKey}(v, f[v])$

isFree(v): is v in C-free



Grid-based planners with A*

- The paths may be longer than true shortest path in C-space
- Not complete, but **resolution complete**
 - probability of finding a solution, if one exists $\rightarrow 1$ as the resolution of the grid increases
- While searching, it finds what points are in C-free, so it constructs C-free. Can interleave the construction with the search (ie construct only what is necessary). Or can construct it all at once (occupancy grid).



- simple to understand/implement
- work in any dimension



- size and quality of path depends on the discretization of the problem
- not practical in high-d spaces
 - e.g. 6 dof: $1000 \times 1000 \times 1000 \times 360 \times 360 \times 360$

A* variants

- weighted A*
 - $c \cdot h() \implies$ solution is no worse than $(1 + c) \times$ optimal
- anytime A*
 - use weighted A* to find a first solution ; then use A* with first solution as upper bound to prune the search
- real-time replanning
 - if the underlying graph changes, it usually affects a small part of the graph \implies don't run search from scratch
 - D*: efficiently recompute SP every time the underlying graph changes
- Finer resolution \Rightarrow better paths but slower
- C-free can be pre-computed (occupancy grid) or computed incrementally
- One-time path planning vs many times; static vs dynamic environment
- fixed resolution vs. multi-resolution techniques

Sampling-based planning

- Geometric planners:
 - hard to construct C-obstacles except for simple cases (2d, no rotation)
- Grid-based planners:
 - grid has uniform resolution and uses too much large for high #dof
 - e.g. DOF= 6: $1000 \times 1000 \times 1000 \times 360 \times 360 \times 360$




Sampling-based planners

- Sample and generate a sparse representation of C-free
- Potential field planners

Sampling-based planners

We don't know the C-obstacles, but we'll assume that we have a function that can check whether a given configuration is free.






All planners need a collision detection function

Would my robot, if placed at this point p, intersect any obstacle?

```
//return true if placement p is in C-free
//put differently, return true if placing the robot R at p does not intersect any obstacle
bool isFree(placement p, the robot, the obstacles)
    • translate and rotate the robot to p
    • check whether any edge of the robot intersects any of the obstacles
```

p is a point in C-space: (x, y) or (x, y, θ) or $(x, y, z, \theta_x, \theta_y, \theta_z)$ or ...



Sampling-based planning

- Idea: Sample C-free and compute a roadmap that captures its connectivity
- Single-query, incremental search planners
 - Construct a graph/roadmap to connect *start* and *end*
 - Reconstruct for different (*start*, *end*) pairs
 - E.g. RRT (rapidly-exploring random tree) and variants
- Multiple-query planners
 - Construct a graph/roadmap and use it for **any** (*start*, *end*) pairs
 - E.g. PRM (probabilistic roadmap) and variants

History

- Dijkstra 1950s
- A* 1960s
- PRM 1996
- RRT 1998
- RRT* 2010

Probabilistic Roadmaps and RRTs

- Efficient, easy to implement, applicable to many types of scenes
- Well-suited for high #dof
- Shown to be **probabilistically complete**
 - Finds a solution, if one exists, with probability $\rightarrow 1$ as the nb. of samples increases
- Leading motion planning technique, embraced by many groups, many variants, used in many type of scenes/applications.
 - PRM*, FMT* (fast marching tree), ...
- No discretization (sample from a continuous space)
- But: Path not optimal, time may be unbounded

The RRT

(LaValle, 1998)

- Incrementally build a tree rooted at “start” outwards, while trying to determine if a path exists at each step
- Original paper:
 - https://www.cs.cmu.edu/~motionplanning/papers/sbp_papers/PRM/randtrees_02.pdf

NEW_CONFIG((q, q_{near}, q_{new}))

- if q is not free, return false
- if segment $q_{near}q_{new}$ is not in C-free, return false

```
BUILD_RRT( $q_{init}$ )
1  $T.init(q_{init});$ 
2 for  $k = 1$  to  $K$  do
3    $q_{rand} \leftarrow \text{RANDOM\_CONFIG}();$ 
4   EXTEND( $T, q_{rand}$ );
5 Return  $T$ 
```

EXTEND(T, q)

```
1  $q_{near} \leftarrow \text{NEAREST\_NEIGHBOR}(q, T);$ 
2 if NEW_CONFIG( $(q, q_{near}, q_{new})$ ) then
3    $T.add\_vertex(q_{new});$ 
4    $T.add\_edge(q_{near}, q_{new});$ 
5   if  $q_{new} = q$  then
6     Return Reached;
7   else
8     Return Advanced;
9 Return Trapped;
```

Figure 2: The basic RRT construction algorithm.






Figure 3: The EXTEND operation.

Random Trees, RRTs & RRT*

1001 nodes, goal not yet reached



1001 nodes, path length 34.94



Probabilistic Roadmaps and RRTs

PRM

- Roadmap adjusts to the density of free space and is more connected around the obstacles
- Size of roadmap can be adjusted as needed
- More time spent in the “learning” phase ==> better roadmap
- Built once, re-used many times, used in static environments



RRT

- Used in changing environments
- Faster to build than PRM

We need to write: $\text{isFree}(p = (x, y, \theta), \text{Robot } R, \text{ Obstacles } S)$

2D: robot can translate and rotate

C-space: 3D



configuration p: (x, y, θ)

$(8, 15, 45)$: not free

$(8, 5, 0)$: free

Also need a segment collision detection function

Is segment pq in C-free?

//p, q are points in C-space: (x, y) or (x, y, θ) or $(x, y, z, \theta_x, \theta_y, \theta_z)$ or ...

bool localPlanner(placement p, placement q, the robot, the obstacles)

Probabilistic roadmaps

Probabilistic Roadmaps (Kavraki, Svetska, Latombe, Overmars et al , 1996)

- **Roadmap construction phase**
 - Start with a sampling of points in C-free and try to connect them
 - Two points are connected by an edge if a simple quick planner can find a path between them
 - This will create a set of connected components
- **Roadmap query phase**
 - Use roadmap to find path between any two points

Probabilistic Roadmaps

- Generic-Sampling-Based-Roadmap:
 - for $i = 1$ to N :
 - generate a random point p_i in C
 - if $\text{isFree}(p)$, add p to R
 - add p_{start} to R
 - for each point p_i in R :
 - $N(p_i) = \{ \text{closest neighbors of } p_i \text{ in } R \}$
 - for each neighbor q in $N(p_i)$:
 - if there is a collision-free local path from p_i to q and there is not already an edge from p_i to q then add an edge from p_i to q in the roadmap R
- Variants
 - how they select the initial n samples from C
 - e.g. return a set of n points arranged on a regular grid in C , random points, etc
 - how they select the neighbors
 - return the k nearest neighbors of p in V
 - return the set of points lying in a ball centered at p of radius r
- Often used: samples arranged in a 2-dimensional grid, with nearest 4 neighbors (2^d)

Probabilistic Roadmaps (Kavraki, Svetska, Latombe, Overmars et al , 1996)

```
(1)    $N \leftarrow \emptyset$ 
(2)    $E \leftarrow \emptyset$ 
(3)   loop
(4)      $c \leftarrow$  a randomly chosen free
           configuration
(5)      $N_c \leftarrow$  a set of candidate neighbors
           of  $c$  chosen from  $N$ 
(6)      $N \leftarrow N \cup \{c\}$ 
(7)     for all  $n \in N_c$ , in order of
           increasing  $D(c,n)$  do
(8)       if  $\neg \text{same\_connected\_component}(c,n)$ 
            $\wedge \Delta(c,n)$  then
(9)          $E \leftarrow E \cup \{(c,n)\}$ 
(10)        update  $R$ 's connected
           components
```

- Start with a *random* sampling of points in C-free
- Roadmap is stored as set of *trees* for space efficiency
 - trees encode connectivity, cycles don't change it. Additional edges are useful for shorter paths, but not for completeness
- Augment roadmap by selecting additional sample points in areas that are estimated to be "difficult"

the local planner $\Delta(c,n)$: is segment cn in C-free?

Probabilistic Roadmaps

- Roadmap adjusts to the density of free space and is more connected around the obstacles
- Size of roadmap can be adjusted as needed
- More time spent in the “learning” phase ==> better roadmap

- Efficient, easy to implement, applicable to many types of scenes
- Well-suited for high #dof
- No discretization (sample from a continuous space)
- Shown to be probabilistically complete
 - finds a solution, if one exists, with probability $\rightarrow 1$ as the nb of samples increases
- Leading motion planning technique, Embraced by many groups, many variants of PRM's, used in many type of scenes/applications (PRM*, FMT* (fast marching tree), ...)

Sampling-based planning

- Geometric planners:
 - hard to construct C-obstacles except for simple cases (2d, no rotation)
- Grid-based planners:
 - grid has uniform resolution and uses too much large for high #dof
 - e.g. DOF= 6: $1000 \times 1000 \times 1000 \times 360 \times 360 \times 360$
- Sampling-based planners
 - Sample and generate a sparse representation of C-free
- Potential field planners

Potential field methods [Latombe et al, 1992]

- Define a potential field
 - Robot moves in the direction of steepest descent on potential function
- Ideally potential function has global minimum at the goal, has no local minima, and is very large around obstacles
- Algorithm outline:
 - place a regular grid over C-space
 - search over the grid with potential function as heuristic

<https://www.youtube.com/watch?v=r9FD7P76zJs>

Potential field methods

- Pro:
 - Framework can be adapted to any specific scene
- Con:
 - can get stuck in local minima
 - Potential functions that are minima-free are known, but expensive to compute
- Example: RPP (Randomized path planner) is based on potential functions
 - Escapes local minima by executing random walks
 - Successfully used to
 - perform riveting ops on plane fuselages
 - plan disassembly operations for maintenance of aircraft engines

Demos

DARPA challenges

- Fostered the development of self-driving vehicles
- 2004: noone finished the course
- 2005:
 - 132 mi course, in the desert in Nevada
 - 5 vehicles finished the race, with Stanford “Stanley” in first place, the first autonomous vehicle to ever finish a race (Stanley now at the Smithsonian Air & Space museum)
- 2007
 - Required teams to build an autonomous vehicle capable of driving in traffic and performing complex maneuevers such as merging, passing and parking
 - 5 vehicles finished the race, with CMU “Boss” in first place, and Stanford “Junior” in second.

DARPA challenges

- Planners: Both graph search and incremental tree-based
 - **CMU**: lattice graph in 4D (x,y, orientation, velocity), search with D*
 - **Stanford**: incremental sparse tree of possible maneuvers, hybrid A*
 - **Virginia Tech**: graph discretization of possible maneuvers, search with A*
 - **MIT**: variant of RRT with biased sampling
- talk by Sertac Karaman in Darpa 2007 MIT team:
- *A Survey of Motion Planning and Control Techniques for Self-driving Urban Vehicles, by Brian Paden, Michal Cáp, Sze Zheng Yong, Dmitry Yershov, and Emilio Frazzoli*
<https://arxiv.org/pdf/1604.07446.pdf>

DARPA 2007, Stanford team

- uses hybrid A*

<http://robots.stanford.edu/papers/junior08.pdf>

Junior: The Stanford Entry in the Urban Challenge

Michael Montemerlo¹, Jan Becker⁴, Suhrid Bhat², Hendrik Dahlkamp¹, Dmitri Dolgov¹,
Scott Ettinger³, Dirk Haehnel¹, Tim Hilden², Gabe Hoffmann¹, Burkhard Huhnke²,
Doug Johnston¹, Stefan Klumpp², Dirk Langer², Anthony Levandowski¹, Jesse Levinson¹,
Julien Marcil², David Orenstein¹, Johannes Paefgen¹, Isaac Penny¹, Anna Petrovskaya¹,
Mike Pflueger², Ganymed Stanek², David Stavens¹, Antone Vogt¹, and Sebastian Thrun¹

¹Stanford Artificial Intelligence Lab, Stanford University, Stanford CS 94305

²Electronics Research Lab, Volkswagen of America, 4009 Miranda Av., Palo Alto, CA 94304

³Intel Research, 2200 Mission College Blvd., Santa Clara, CA 95052

⁴Robert Bosch LLC, Research and Technology Center, 4009 Miranda Ave, Palo Alto, CA 94304

Abstract

This article presents the architecture of Junior, a robotic vehicle capable of navigating urban environments autonomously. In doing so, the vehicle is able to select its own routes, perceive and interact with other traffic, and execute various urban driving skills including lane changes, U-turns, parking, and merging into moving traffic. The vehicle successfully finished and won second place in the DARPA Urban Challenge, a robot competition organized by the U.S. Government.

1 Introduction

<https://www.youtube.com/watch?v=qXZt-B7iUyw>

DARPA

- Currently, RACER challenge for off-road autonomous vehicles



<https://www.darpa.mil/news-events/2023-04-11>




4

Understanding PATH PLANNING

A*, RRT, RRT*

MATLAB TECH TALKS


PRM: Probabilistic Roadmap Method for robotics



<https://www.youtube.com/watch?v=tIFVbHENPCl>

Comparison of RRT, PRM (MIT course project)

https://www.youtube.com/watch?v=gP6MRe_IHFo



Project 7

Heuristical motion planning

Implement (x, y, θ) -planning for a polygonal robot moving with translation and rotation in 2D using one of:

- A*
- A tree (RRT)
- A probabilistic roadmap (PRM)



