


Motion Planning I

D.A. Forsyth

(with a lot of H. Choset, and some J. Li)

What is motion planning?


- The automatic generation of motion
 - Path + velocity and acceleration along the path



Li slides

Basic Problem Statement

- Motion planning in robotics
 - Automatically compute a path for an object/robot that does not collide with obstacles.



Why is this not just optimization?

- Find minimum cost set of controls that
 - take me from A to B
 - do not involve
 - collision
 - unnecessary extreme control inputs
 - unnecessary extreme behaviors

$$\text{minimize } f(\mathbf{x}) \quad (1a)$$



$$\text{subject to} \quad (1b)$$

These will have to deal with collisions, etc. \longrightarrow $g_i(\mathbf{x}) \leq 0, \quad i = 1, 2, \dots, n_{ineq} \quad (1c)$


$$h_i(\mathbf{x}) = 0, \quad i = 1, 2, \dots, n_{eq} \quad (1d)$$

Is motion planning hard?

Basic Motion
Planning Problems



Degrees of Freedom




- The geometric configuration of a robot is defined by p degrees of freedom (DOF)
- Assuming p DOFs, the geometric configuration A of a robot is defined by p variables:

$$A(\mathbf{q}) \text{ with } \mathbf{q} = (q_1, \dots, q_p)$$


- Examples:
 - Prismatic (translational) DOF: q_i is the amount of translation in some direction
 - Rotational DOF: q_i is the amount of rotation about some axis

Our car has 3

Examples




Allowed to move only
in x and y : 2DOF




Allowed to move in x
and y and to rotate:
3DOF (x, y, θ)

Configuration Space (C-Space)



$$\mathbf{q} = (x, y, \theta)$$

$\mathcal{C} = \mathcal{R}^2 \times \text{set of 2-D rotations}$




$$\mathbf{q} = (q_1, q_2)$$

$\mathcal{C} = \text{2-D rotations} \times \text{2-D rotations}$


- Configuration space $\mathcal{C} = \text{set of values of } \mathbf{q}$ corresponding to **legal configurations** of the robot
- Defines the set of possible parameters (the search space) and the set of allowed paths

Free Space: Point Robot




- $\mathcal{C}_{\text{free}} = \{\text{Set of parameters } \mathbf{q} \text{ for which } A(\mathbf{q}) \text{ does not intersect obstacles}\}$
- For a point robot in the 2-D plane: \mathbb{R}^2 minus the obstacle regions

Free Space: Symmetric Robot



- We still have $\mathcal{C} = \mathbb{R}^2$ because orientation does not matter
- Reduce the problem to a point robot by expanding the obstacles by the radius of the robot

Free Space: Non-Symmetric Robot



- The configuration space is now three-dimensional (x, y, θ)
- We need to apply a different obstacle expansion for each value of θ
- We still reduce the problem to a point robot by expanding the obstacles

Any Formal Guarantees? Generic Piano Movers Problem




- Formal Result (but not terribly useful for practical algorithms):
 - p : Dimension of \mathcal{C}
 - m : Number of polynomials describing $\mathcal{C}_{\text{free}}$
 - d : Max degree of the polynomials
- A path (if it exists) can be found in time *exponential in p* and *polynomial in m and d*

[From J. Canny. "The Complexity of Robot Motion Planning Plans". MIT Ph.D. Dissertation. 1987]

Observation

- Generally, searching a graph is pretty straightforward
 - Dijkstra, A*, etc - know how to do this
- Strategy
 - get a graph we can search

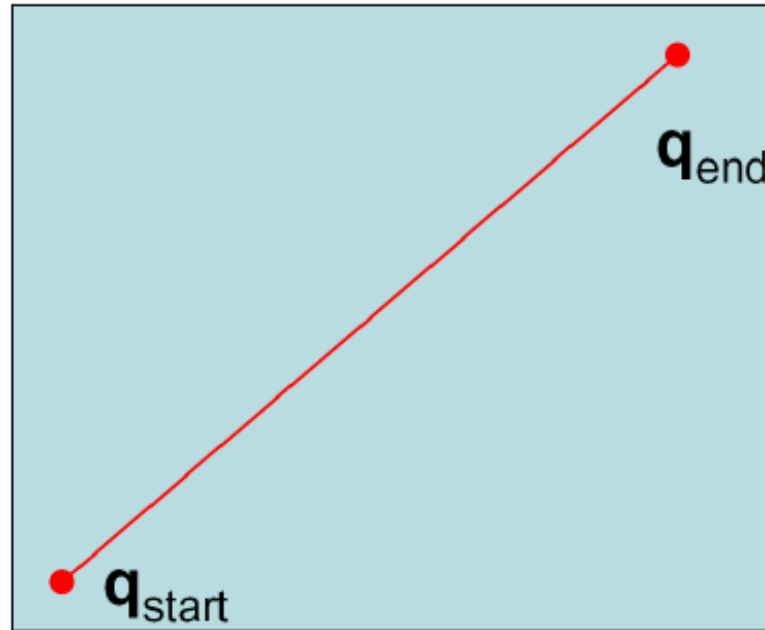
Roadmaps



General idea:


- Avoid searching the entire space
- Pre-compute a (hopefully small) graph (the roadmap) such that staying on the “roads” is guaranteed to avoid the obstacles
- Find a path between q_{start} and q_{goal} by using the roadmap

Visibility Graphs



In the absence of obstacles, the best path is the straight line between q_{start} and q_{goal}


Visibility Graphs



- Visibility graph G = set of unblocked lines between vertices of the obstacles + $\mathbf{q}_{\text{start}}$ and \mathbf{q}_{goal}
- A node P is linked to a node P' if P' is visible from P
- Solution = Shortest path in the visibility graph

Issues


- Constructing
 - Relatively straightforward with a sweep algorithm
 - Variant (visibility complex) root cause of early computer games
 - Wolfenstein 3D, Doom II, etc
- What if configuration space is not 2D
 - You can still construct, MUCH harder
- MANY locally optimal paths
 - topology of free space clearly involved



Visibility Graphs: Weaknesses

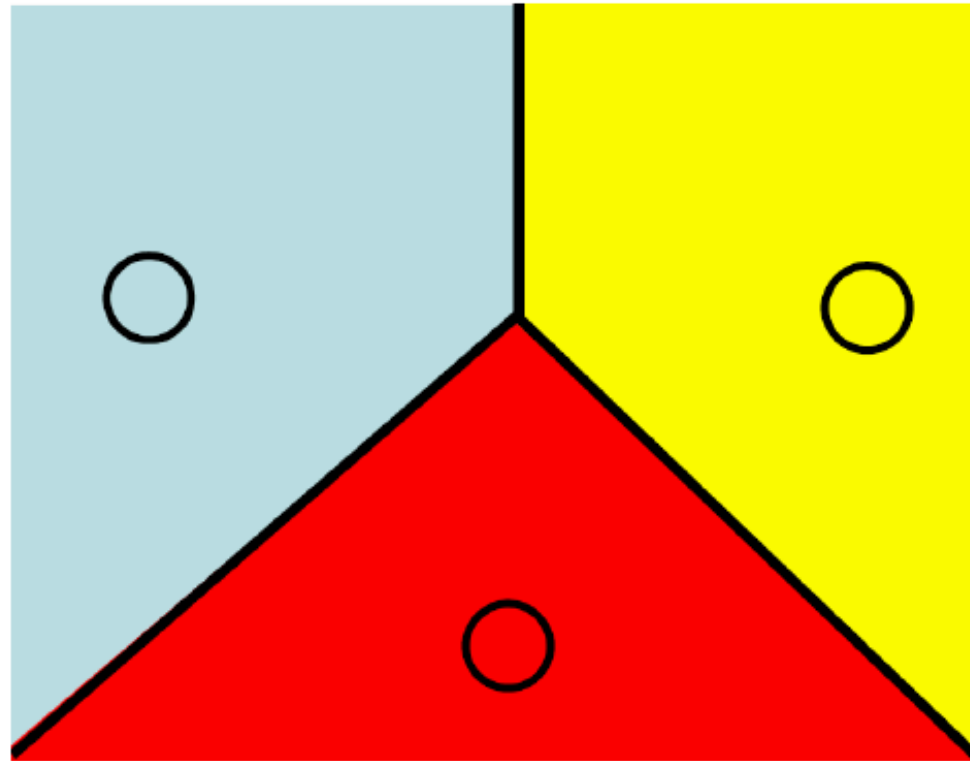
- Shortest path but:
 - Tries to stay as close as possible to obstacles
 - Any execution error will lead to a collision
 - Complicated in $\gg 2$ dimensions
- We may not care about strict optimality so long as we find a safe path. Staying away from obstacles is more important than finding the shortest path
- Need to define other types of “roadmaps”

Voronoi Diagrams




- Given a set of data points in the plane:
 - Color the entire plane such that the color of any point in the plane is the same as the color of its nearest neighbor

Voronoi Diagrams



- Voronoi diagram = The set of line segments separating the regions corresponding to different colors
 - Line segment = points equidistant from 2 data points
 - Vertices = points equidistant from > 2 data points


Voronoi Diagrams



- Complexity (in the plane):
- $O(N \log N)$ time
- $O(N)$ space


(See for example <http://www.cs.cornell.edu/Info/People/chew/Delaunay.html> for an interactive demo)

Voronoi Diagrams (Polygons)



- Key property: The points on the edges of the Voronoi diagram are the *furthest* from the obstacles
- Idea: Construct a path between $\mathbf{q}_{\text{start}}$ and \mathbf{q}_{goal} by following edges on the Voronoi diagram
- (Use the Voronoi diagram as a roadmap graph instead of the visibility graph)

Voronoi Diagrams: Planning



- Find the point $\mathbf{q}^*_{\text{start}}$ of the Voronoi diagram closest to $\mathbf{q}_{\text{start}}$
- Find the point $\mathbf{q}^*_{\text{goal}}$ of the Voronoi diagram closest to \mathbf{q}_{goal}
- Compute shortest path from $\mathbf{q}^*_{\text{start}}$ to $\mathbf{q}^*_{\text{goal}}$ on the Voronoi diagram


Voronoi: Weaknesses

- **Difficult to compute in higher dimensions** or nonpolygonal worlds
- Approximate algorithms exist
- Use of Voronoi is not necessarily the best heuristic (“stay away from obstacles”) Can lead to paths that are much too conservative
- Can be unstable → Small changes in obstacle configuration can lead to large changes in the diagram

Approximate Cell Decomposition: Limitations

- Good:
 - Limited assumptions on obstacle configuration
 - Approach used in practice
 - Find obvious solutions quickly
- Bad:
 - No clear notion of optimality (“best” path)
 - Trade-off completeness/computation
 - Still difficult to use in high dimensions


Exact Cell Decomposition



Any path within one cell is guaranteed to not intersect any obstacle

Choset slides

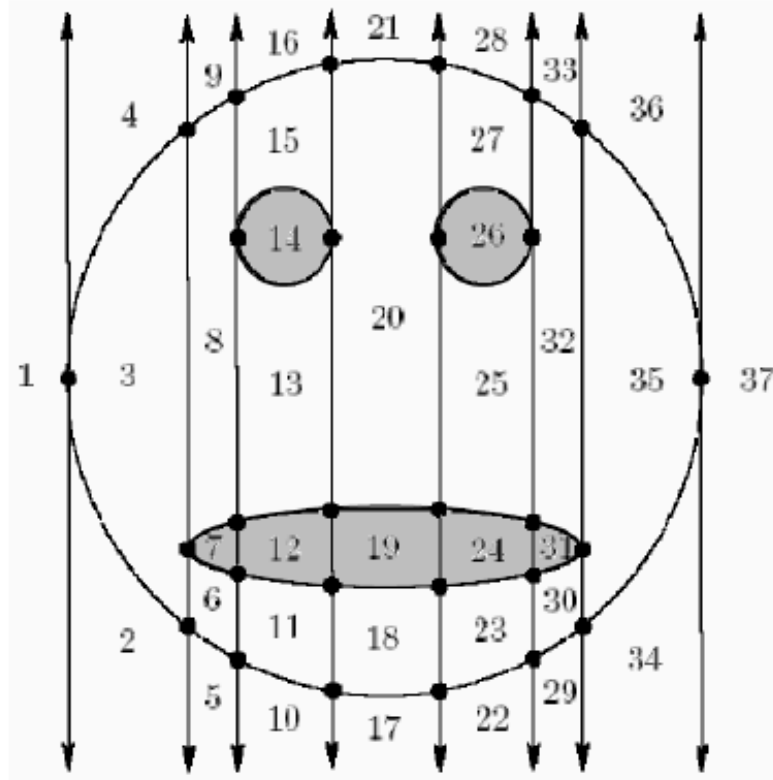
Exact Cell Decomposition



- The graph of cells defines a roadmap


Choset slides

Exact Cell Decomposition




- A version of exact cell decomposition can be extended to higher dimensions and non-polygonal boundaries (“cylindrical cell decomposition”)
- Provides exact solution → completeness
- Expensive and difficult to implement in higher dimensions

Potential Fields



- Stay away from obstacles: Imagine that the obstacles are made of a material that generate a *repulsive* field
- Move closer to the goal: Imagine that the goal location is a particle that generates an *attractive* field



Choset slides

$$U_g(\mathbf{q}) = d^2(\mathbf{q}, \mathbf{q}_{goal})$$

Distance to goal state


$$U_o(\mathbf{q}) = \frac{1}{d^2(\mathbf{q}, Obstacles)}$$

Distance to nearest obstacle point.
Note: Can be computed efficiently by
using the *distance transform*

$$U(\mathbf{q}) = U_g(\mathbf{q}) + \lambda U_o(\mathbf{q})$$

λ controls how far we
stay from the obstacles

Potential Fields: Limitations




Potential field

Zoomed in view

- Completeness?
- Problems in higher dimensions

Local Minimum Problem



- Potential fields in general exhibit local minima
- Special case: Navigation function
 - $U(\mathbf{q}_{\text{goal}}) = 0$
 - For any \mathbf{q} different from \mathbf{q}_{goal} , there exists a neighbor \mathbf{q}' such that $U(\mathbf{q}') < U(\mathbf{q})$

Getting out of Local Minima I

- Repeat
 - If $U(\mathbf{q}) = 0$ return Success
 - If too many iterations return Failure
 - Else:
 - Find neighbor \mathbf{q}_n of \mathbf{q} with smallest $U(\mathbf{q}_n)$
 - If $U(\mathbf{q}_n) < U(\mathbf{q})$ OR \mathbf{q}_n has not yet been visited
 - Move to \mathbf{q}_n ($\mathbf{q} \leftarrow \mathbf{q}_n$)
 - Remember \mathbf{q}_n

May take a long time to explore region “around” local minima

Getting out of Local Minima I

- Repeat
 - If $U(\mathbf{q}) = 0$ return Success
 - If too many iterations return Failure
 - Else:
 - Find neighbor \mathbf{q}_n of \mathbf{q} with smallest $U(\mathbf{q}_n)$
 - If $U(\mathbf{q}_n) < U(\mathbf{q})$ OR \mathbf{q}_n has not yet been visited
 - Move to \mathbf{q}_n ($\mathbf{q} \leftarrow \mathbf{q}_n$)
 - Remember \mathbf{q}_n

May take a long time to explore region “around” local minima

- Think of this the following way:
 - impose a grid
 - do depth first search on the potential
- Idea:
 - other kinds of search
 - randomization should help a lot
- Concern:
 - what if \mathbf{q} has lots of neighbors?

Getting out of Local Minima II

- Repeat
 - If $U(\mathbf{q}) = 0$ return Success
 - If too many iterations return Failure
 - Else:
 - Find neighbor \mathbf{q}_n of \mathbf{q} with smallest $U(\mathbf{q}_n)$
 - If $U(\mathbf{q}_n) < U(\mathbf{q})$
 - Move to \mathbf{q}_n ($\mathbf{q} \leftarrow \mathbf{q}_n$)
 - Else
 - Take a random walk for T steps starting at \mathbf{q}_n
 - Set \mathbf{q} to the configuration reached at the end of the random walk

Similar to stochastic search and simulated annealing:
We escape local minima faster

Getting out of Local Minima II

- Repeat
 - If $U(\mathbf{q}) = 0$ return Success
 - If too many iterations return Failure
 - Else:
 - Find neighbor \mathbf{q}_n of \mathbf{q} with smallest $U(\mathbf{q}_n)$
 - If $U(\mathbf{q}_n) < U(\mathbf{q})$
 - Move to \mathbf{q}_n ($\mathbf{q} \leftarrow \mathbf{q}_n$)
 - Else
 - Take a random walk for T steps starting at \mathbf{q}_n
 - Set \mathbf{q} to the configuration reached at the end of the random walk

Similar to stochastic search and simulated annealing:
We escape local minima faster

- Intuition:
 - random walk should get you out of local minima
 - then slide down the potential function
- Concern:
 - what if dimension is high?
 - random walk may not get out of local minima efficiently