



Spring 26
ECE484
Lecture 14
Planning 1

Sayan Mitra

Search and planning problems appear in different levels and scales of autonomy stack

Global path planner --- invoked at each new *checkpoint*


- finds paths from every point in the map to next checkpoint
- dynamic programming

Road navigation

- For each path, the planner rolls out several discrete trajectories that are parallel to the smoothed center of the lane

Freeform navigation (parking lots)

- Generate arbitrary trajectories (irrespective of road structure) using modified A*



Planning as graph search

Search for collision free trajectories can be viewed abstractly as a graph search problem, where the nodes represent blocks of free space and the edges connect spatially adjacent nodes.

We can solve such problems using the graph search algorithms like (uninformed) Breadth-First Search and Depth-First Search

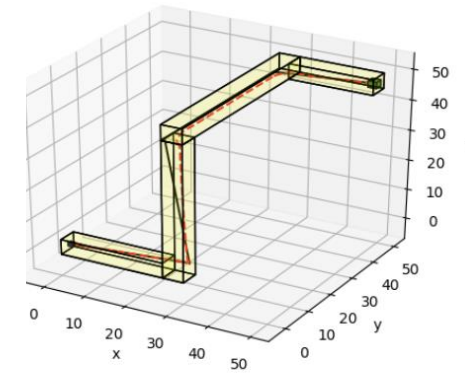

Roadmaps are not arbitrary graphs: Some paths are more preferable than others (e.g., shorter, faster, less costly in terms of fuel/tolls/fees, more stealthy, etc.)

Good guesses for distances can be made, even without knowing the full graph or the optimal paths

Can we *utilize this information* to find efficient paths, efficiently?



Making a Drone Smarter With Motion Planning [Nicholas Rehm](#)



<https://kmmille.github.io/FACTEST/>

Problem statement: find shortest path

Input: $\langle V, E, w, start, goal \rangle$


- V : (finite) set of vertices
- $E \subseteq V \times V$: (finite) set of edges
- $w : E \rightarrow \mathbb{R}_{>0}$: associates to each edge e to a positive weight $w(e)$
- $start, goal \in V$: start and end vertices.

A path is a sequence of vertices $p = v_0 \dots v_k$ such that $(v_i, v_{i+1}) \in E$

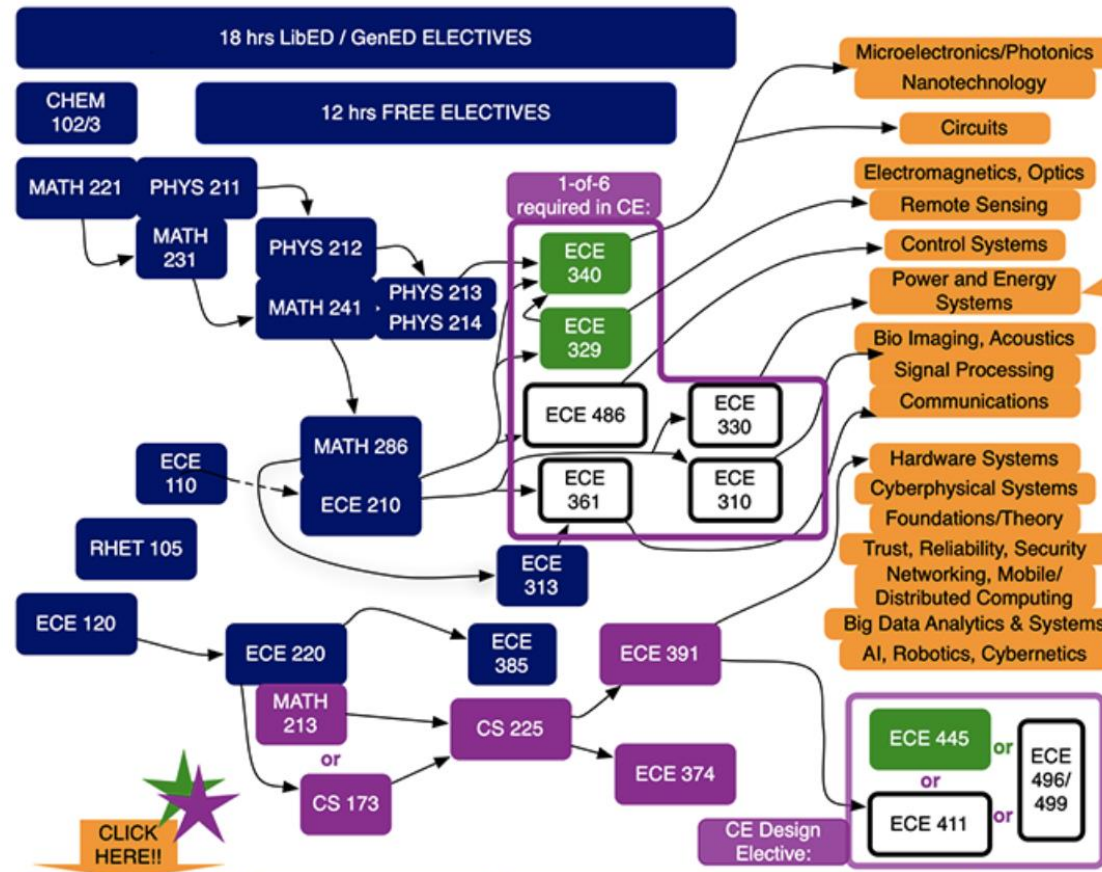
$$w(p) = \sum_{i=0}^{k-1} w(v_i, v_{i+1}) \quad \text{head}(p) = v_0 \quad \text{end}(p) = v_k$$

Output: a path p with $head(p) = start$ and $end(p) = goal$, such that its $w(p)$ is minimal among all such paths

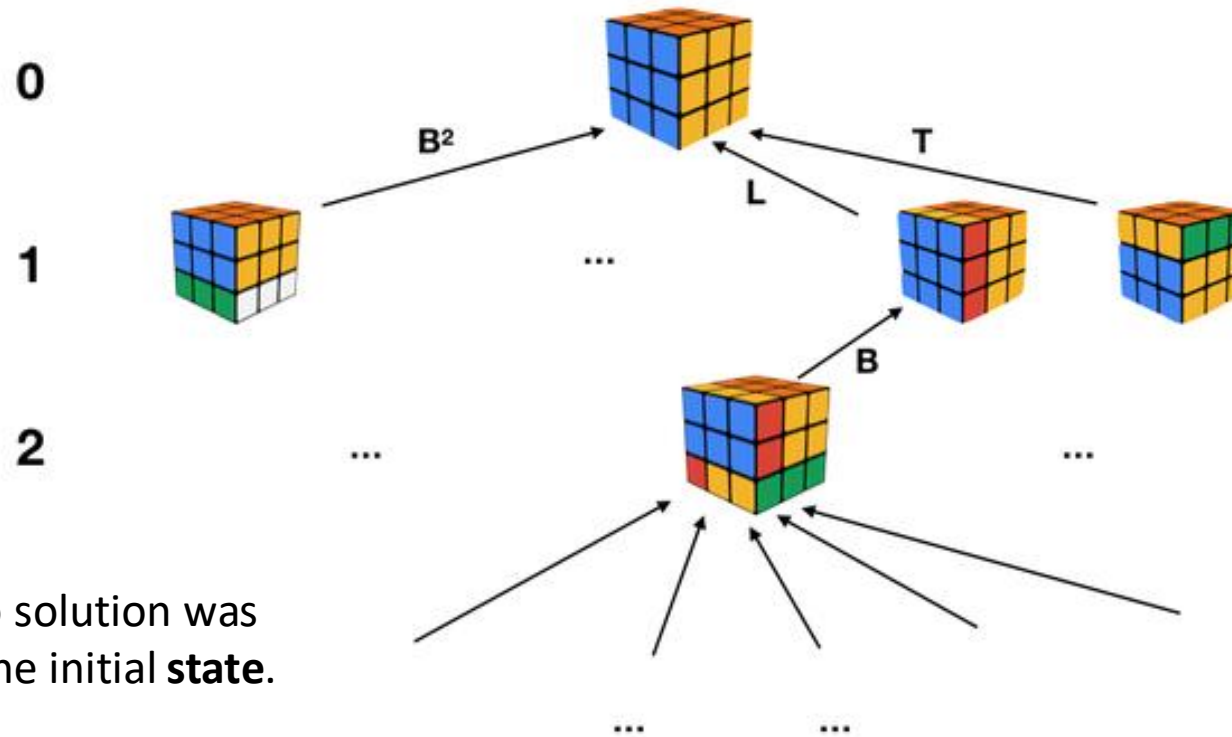
Algorithms: Dijkstra $O((|V| + |E|)\log|V|)$, Bellman-Ford $O(|V| \times |E|)$,...



Example graph with imperfectly known information



The Graph Can be Large



Number of states or vertices
43,252,003,274,489,856,000

Yet, maximum length of path to solution was shown to be 20, regardless of the initial **state**.

T. Rokicki, working with Google, proved "[God's number](#)" to be 20, in 2010.



Outline

- Deterministic search
 - Uninformed search
 - Informed search
 - Optimal search: A, A*

Search Algorithm Performance Metrics

Soundness: when a solution is returned, is it guaranteed to be correct path?

Completeness: is the algorithm guaranteed to find a solution when one exists?

Optimality: How close is the found solution to the best solution?

Space complexity: how much memory is needed?

Time complexity: what is the running time? Can it be used for online planning?

Uniform cost search (Uninformed search)

```
Q ← ⟨start⟩ // maintains paths sorted by cost
// initialize queue with start


while Q ≠ ∅:
    pick (and remove) from Q the path P with lowest cost w(P)
    if head(P) = goal then return P ; // Reached the goal
    foreach vertex v such that (head(P), v) ∈ E, do // for all neighbors
        add ⟨v, P⟩ to Q ; // Add expanded paths
return FAILURE ; // nothing left to consider
```

Note no ***visited*** list; Use no information obtained from the environment

Example of Uniform-Cost Search

Q:


| Path | Cost |
|---------------------|------|
| $\langle s \rangle$ | 0 |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |



Example of Uniform-Cost Search

Q:


| Path | Cost |
|------------------------|------|
| $\langle a, s \rangle$ | 2 |
| $\langle b, s \rangle$ | 5 |
| | |
| | |
| | |
| | |
| | |
| | |
| | |



Example of Uniform-Cost Search

Q:


| Path | Cost |
|---------------------------|------|
| $\langle c, a, s \rangle$ | 4 |
| $\langle b, s \rangle$ | 5 |
| $\langle d, a, s \rangle$ | 6 |
| | |
| | |
| | |
| | |
| | |
| | |
| | |



Example of Uniform-Cost Search

Q:


| Path | Cost |
|------------------------------|------|
| $\langle b, s \rangle$ | 5 |
| $\langle d, a, s \rangle$ | 6 |
| $\langle d, c, a, s \rangle$ | 7 |
| | |
| | |
| | |
| | |
| | |
| | |



Example of Uniform-Cost Search

Q:

| Path | Cost |
|------------------------------|------|
| $\langle d, a, s \rangle$ | 6 |
| $\langle d, c, a, s \rangle$ | 7 |
| $\langle g, b, s \rangle$ | 10 |
| | |
| | |
| | |
| | |
| | |




Notice a path to goal has been found but the algorithm does not terminate

Example of Uniform-Cost Search

Q:


| Path | Cost |
|------------------------------|------|
| $\langle d, c, a, s \rangle$ | 7 |
| $\langle g, d, a, s \rangle$ | 8 |
| $\langle g, b, s \rangle$ | 10 |
| | |
| | |
| | |
| | |



Example of Uniform-Cost Search

Q:

| Path | Cost |
|---------------------------------|------|
| $\langle g, d, a, s \rangle$ | 8 |
| $\langle g, d, c, a, s \rangle$ | 9 |
| $\langle g, b, s \rangle$ | 10 |
| | |
| | |
| | |
| | |



Algorithm stops when smallest cost path in Q has head = g

Properties of Uniform Cost Search

UCS is an extension of BFS to the weighted-graph case (UCS = BFS if all edges have the same cost)

UCS is *sound, complete* and *optimal* (assuming costs bounded away from zero)

- Exercise: Prove the above claims
- Soundness: For any entry p in Q , is a path and $\text{head}(p) = \text{start}$

UCS is *guided by path cost rather than path depth, so it may get in trouble if some edge costs are very small*

Worst-case time and space complexity $O(b^{W^/\epsilon})$, where W^* is the optimal cost, and ϵ is such that all edge weights are no smaller than*

Greedy or Best-First Search

UCS explores paths in all directions, with no bias towards the goal state

What if we try to get “closer” to the goal?

We need a **measure of distance to the goal**

It would be ideal to use the length of the shortest path...

but this is exactly what we are trying to compute!

We can **estimate** the distance to the goal through a **heuristic function**

$h : V \rightarrow \mathbb{R}_{\geq 0}$. E.g., the Euclidean distance to the goal (as the crow flies)

$h(v)$ is the estimate of the distance from v to goal

A reasonable strategy is to always try to move in such a way to minimize the estimated distance to the goal: this is the basic idea of the **greedy (best-first) search**


Greedy/Best-first search

```
Q ← ⟨start⟩ // initialize queue with start
while Q ≠ ∅:
    pick (and remove) the path P with lowest heuristic cost h(head(P)) from Q
    if head(P) = goal then return P // Reached the goal
    foreach vertex v such that (head(P), v) ∈ E, do // for all neighbors
        add ⟨v, P⟩ to Q ; // Add expanded paths
return FAILURE ; // nothing left to consider
```

Example of Greedy search

Q:


| Path | Cost | h |
|---------------------|------|----|
| $\langle s \rangle$ | 0 | 10 |



Example of Greedy search

Q:


| Path | Cost | h |
|------------------------|------|---|
| $\langle a, s \rangle$ | 2 | 2 |
| $\langle b, s \rangle$ | 5 | 3 |



Example of Greedy search

Q:


| Path | Cost | h |
|---------------------------|------|---|
| $\langle c, a, s \rangle$ | 4 | 1 |
| $\langle d, a, s \rangle$ | 6 | 4 |
| $\langle b, s \rangle$ | 5 | 3 |



Example of Greedy search

Q:


| Path | Cost | h |
|------------------------------|------|---|
| $\langle d, c, a, s \rangle$ | 7 | 4 |
| $\langle d, a, s \rangle$ | 6 | 4 |
| $\langle b, s \rangle$ | 5 | 3 |



Example of Greedy search

Q:

| Path | Cost | h |
|------------------------------|------|---|
| $\langle d, c, a, s \rangle$ | 7 | 4 |
| $\langle d, a, s \rangle$ | 6 | 4 |
| $\langle g, b, s \rangle$ | 10 | 0 |



Terminates with suboptimal path

Remarks on greedy/best-first search

Greedy (Best-First) search is similar to Depth-First Search

keeps exploring until it has to back up due to a dead end

Not complete (why?) and not optimal, but is often fast and efficient, depending on the heuristic function h

Exercise: Find a counter-example where path exists but bad heuristic function makes the algorithm loop forever

Worst-case time and space complexity?

A search: informed search

The problems

UCS is optimal, but may wander around a lot before finding the goal

Greedy is not optimal, but can be efficient, as it is heavily biased towards moving towards the goal. *The non-optimality comes from neglecting “the past.”*

The idea

Keep track *both of the cost of the partial path to get to a vertex, say $g(v)$, and of the heuristic function estimating the cost to reach the goal from a vertex, $h(v)$*

In other words, choose as a “ranking” function the sum of the two costs:

$$f(v) = g(v) + h(v)$$

$g(v)$ cost-to-come (from the start to v)

$h(v)$: cost-to-go estimate (from v to the goal)

$f(v)$: estimated cost of the path (from the start to v and then to the goal)

A search


open set and closed set

```
Q ← ⟨start⟩ // initialize queue with start
while Q ≠ ∅:
    pick (and remove) path P with lowest estimated cost  $f(P) = g(P) + h(\text{head}(P))$  from Q
    if head(P) = goal then return P // Reached the goal
    foreach vertex v such that (head(P), v) ∈ E, do // for all neighbors
        add ⟨v, P⟩ to Q ; // Add expanded paths
return FAILURE ; // nothing left to consider
```

Example of A search

Q:


| | | | |
|---------------------|---|----|----|
| Path | g | h | f |
| $\langle s \rangle$ | 0 | 10 | 10 |



Example of A search

Q:

| Path | g | h | f |
|------------------------|---|---|---|
| $\langle a, s \rangle$ | 2 | 2 | 4 |
| $\langle b, s \rangle$ | 5 | 3 | 8 |




Example of A search

Q:

| Path | g | h | f |
|--|--------------|--------------|--------------|
| $\langle a, s \rangle$ | 2 | 2 | 4 |
| $\rightarrow \langle b, s \rangle$ | 5 | 3 | 8 |

$\rightarrow \langle c, a, s \rangle$ 4 1 5
 $\langle d, a, s \rangle$ 6 5 11
 $\langle d, c, a, s \rangle$ 7 5 12
 $\langle g, b, s \rangle$ 10 0 10

not optimal



Remarks on A search

A search is similar to UCS, with a bias induced by the heuristic h
If $h = 0$ then A search = UCS.

The A search is complete, but is *not optimal*

What is wrong? (Recall that if $h = 0$ then A = UCS, and hence optimal...)

A* Search

Choose an *admissible heuristic*, i.e., such *that* $h(v) \leq h^*(v)$

$h^*(v)$ is the “optimal” heuristic---perfect cost to go

To be admissible $h(v)$ should be at most $h^*(v)$

A search with an admissible heuristic is called A* --- guaranteed to find optimal path


Example of A* search

Q:


| | | | |
|---------------------|---|---|---|
| Path | g | h | f |
| $\langle s \rangle$ | 0 | 6 | 6 |

changed h

*finds $\langle g, d, a, s \rangle$
optimal path*



Proof of optimality of A*



Let w^* be the cost of the optimal path

Suppose for the sake of contradiction, that A* returns P with $w(P) > w^*$

Find the first unexpanded node on the optimal path P^* ; call it n

$f(n) > w(P)$, otherwise n would have been expanded

$$f(n) = g(n) + h(n)$$

$$= g^*(n) + h(n) \quad [\text{since } n \text{ is on the optimal path}]$$

$$\leq g^*(n) + h^*(n) \quad [\text{since } h \text{ is admissible}]$$

$$= f^*(n) = w^* \quad [\text{by def. of } f, \text{ and since } w^* \text{ is the cost of the optimal path}]$$

Hence $w^* \geq f(n) = w(P)$, which is a contradiction

Admissible heuristics

- How to find an admissible heuristic? i.e., a heuristic that never overestimates the cost-to-go.
- Examples of admissible heuristics
 - $h(v) = 0$: this always works! However, it is not very useful, $A^* = \text{UCS}$
 - $h(v) = \text{distance}(v, g)$ when the vertices of the graphs are physical locations
 - $h(v) = \|v - g\|_p$, when the vertices of the graph are points in a normed vector space
- A general method
 - Choose h as the optimal cost-to-go function for a relaxed problem, that is easy to compute
 - Relaxed problem: ignore some of the constraints in the original problem

Admissible heuristics for the 8-puzzle

Initial state:

| | | |
|---|---|---|
| 1 | | 5 |
| 2 | 6 | 3 |
| 7 | 4 | 8 |

Goal state:

| | | |
|---|---|---|
| 1 | 2 | 3 |
| 4 | 5 | 6 |
| 7 | 8 | |

Which of the following are admissible heuristics?

- $h = 0$
- $h = 1$
- $h =$ number of tiles in the wrong position
- $h =$ sum of (Manhattan) distance between tiles and their goal position

YES, always good

not valid in goal state

YES, “teleport” each tile to the goal in one move

YES, move each tile to the goal ignoring other tiles.

A partial order of heuristic functions

- Some heuristics are better than others
 - $h = 0$ is an admissible heuristic, but is not very useful
 - $h = h^*$ is also an admissible heuristic, and it the “best” possible one (it give us the optimal path directly, no searches/backtracking)
- Partial order
 - We say that h_1 **dominates** h_2 if $h_1(v) \geq h_2(v)$ for all vertices v
 - h^* dominates all admissible heuristics, and 0 is dominated by all admissible heuristics
- **Choosing the right heuristic**
 - In general, we want a heuristic that is as close to h^* as possible
 - However, such a heuristic may be too complicated to compute
 - There is a tradeoff between complexity of computing h and the complexity of the search

Consistent heuristics

- An additional useful property for A* heuristics is called **consistency**
 - A heuristic $h : X \rightarrow \mathbb{R}_{\geq 0}$ is said **consistent** if $h(u) \leq w(e = (u, v)) + h(v), \forall (u, v) \in E$
 - In other words, a consistent heuristics satisfies a triangle inequality
- If h is a consistent heuristics, then $f = g + h$ is non-decreasing along paths: $f(v) = g(v) + h(v) = g(u) + w(u, v) + h(v) \geq f(u)$
- Hence, the values of f on the sequence of nodes expanded by A* is non-decreasing: the first path found to a node is also the optimal path \Rightarrow no need to compare costs!

Summary

- A* algorithm combines cost-to-come $g(v)$ and a heuristic function $h(v)$ for cost-to-go to find shortest path
 - informed search
- heuristic function must be *admissible* $h(v) \leq h^*(v)$
 - Never over-estimate the actual cost to go
 - Are all $h(v)$ values needed ?
 - What if h is not admissible
 - How to find heuristics



From Graph Search to Optimal Control

Dynamic Programming and LQR

ECE 484 · Planning 2

Spring 2026