CS 343: Artificial Intelligence

Search



Profs. Peter Stone and Yuke Zhu

The University of Texas at Austin

[These slides are based on slides created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All CS188 materials are available at http://ai.berkeley.edu.]

Good morning colleagues!

- Past due:
 - Python tutorial
 - 3 reading responses: Al100 report; Chapters 1,2; Chapter 3
- HW1: Search
 - Due Monday 2/8 at 11:59 pm
- P1: Search
 - Due Wednesday 2/10 at 11:59 pm
 - Pair work allowed
- Readings: Constraint Satisfaction and Local Search
 - NOT just Chapter 4
 - Due Monday 2/1 at 9:30 am

Recap: Search

- Search problem:
 - States (configurations of the world)
 - Actions and costs
 - Successor function (world dynamics)
 - Start state and goal test
- Search tree:
 - Nodes: represent plans for reaching states
 - Plans have costs (sum of action costs)
- Search algorithm:
 - Systematically builds a search tree
 - Chooses an ordering of the fringe (unexplored nodes)
 - Optimal: finds least-cost plans



The Main Search Algorithms

- Uninformed Search:
 - Breadth First Search (BFS)
 - Depth First Search (DFS)
 - Uniform Cost Search (UCS) ~ [Dijkstra's Algorithm]
- Informed Search:
 - Greedy Search ~ [Best First Search]
 - A* Search
- Test your understanding!

Some of Your Questions

- Doesn't backward search require backward links? (Nayan)
- Any new search algorithms since the book was published? (Wentao)
- Doesn't path cost count as informed search? (Anish)
- Is space complexity still a concern? (Thomas)
- Can an agent have multiple goals (e.g. get to location with least gas)? What if they conflict?
- Could informed search be worse than uninformed? (Gautham)

The One Queue

- All these search algorithms are the same except for fringe strategies
 - Conceptually, all fringes are priority queues (i.e. collections of nodes with attached priorities)
 - Practically, for DFS and BFS, you can avoid the log(n) overhead from an actual priority queue, by using stacks and queues
 - Can even code one implementation that takes a variable queuing object



Search Algorithm Properties

- Complete: Guaranteed to find a solution if one exists?
- Optimal: Guaranteed to find the least cost path?
- Time complexity?
- Space complexity?
- Cartoon of search tree:
 - b is the branching factor
 - m is the maximum depth
 - solutions at various depths
- Number of nodes in entire tree?
 - $1 + b + b^2 + \dots b^m = O(b^m)$



Depth-First Search (DFS) Properties

- What nodes DFS expand?
 - Some left prefix of the tree.
 - Could process the whole tree!
 - If m is finite, takes time O(b^m)
- How much space does the fringe take?
 - Only has siblings on path to root, so O(bm)
- Is it complete?
 - m could be infinite, so only if we prevent cycles (more later)
- Is it optimal?
 - No, it finds the "leftmost" solution, regardless of depth or cost



Breadth-First Search (BFS) Properties

- What nodes does BFS expand?
 - Processes all nodes above shallowest solution
 - Let depth of shallowest solution be s
 - Search takes time O(b^s)
- How much space does the fringe take?
 - Has roughly the last tier, so O(b^s)
- Is it complete?
 - s must be finite if a solution exists, so yes!
- Is it optimal?
 - Only if costs are all 1 (more on costs later)



Iterative Deepening

- Idea: get DFS's space advantage with BFS's time / shallow-solution advantages
 - Run a DFS with depth limit 1. If no solution...
 - Run a DFS with depth limit 2. If no solution...
 - Run a DFS with depth limit 3.
- Isn't that wastefully redundant?
 - Generally most work happens in the lowest level searched, so not so bad!



Uniform Cost Search (UCS) Properties

- What nodes does UCS expand?
 - Processes all nodes with cost less than cheapest solution!
 - If that solution costs C^* and arcs cost at least ϵ , then the "effective depth" is roughly $C^*\!/\!\epsilon$
 - Takes time O(b^{C*/ε}) (exponential in effective depth)
- How much space does the fringe take?
 - Has roughly the last tier, so O(b^{C*/ε})
- Is it complete?
 - Assuming best solution has a finite cost and minimum arc cost is positive, yes!
- Is it optimal?
 - Yes! (Proof via A*)



Uniform Cost Issues

 Remember: UCS explores increasing cost contours

• The good: UCS is complete and optimal!

- The bad:
 - Explores options in every "direction"
 - No information about goal location
- We'll fix that soon!





Informed Search



Search Heuristics

- A heuristic is:
 - A function that *estimates* how close a state is to a goal
 - Designed for a particular search problem
 - Examples: Manhattan distance, Euclidean distance for pathing





Greedy Search

- Strategy: expand a node that you think is closest to a goal state
 - Heuristic: estimate of distance to nearest goal for each state
- Best case:
 - Best-first takes you straight to the nearest goal
- A common case:
 - Suboptimal route to goal due to imperfect heuristic
 - Does not lead to nearest goal
- Worst-case: like a badly-guided DFS





Combining UCS and Greedy

- Uniform-cost orders by path cost, or backward cost g(n)
- Greedy orders by goal proximity, or *forward cost* h(n)



h=0

A* Search orders by the sum: f(n) = g(n) + h(n)

When should A* terminate?

Should we stop when we enqueue a goal?



No: only stop when we expand a goal

Is A* Optimal?



- What will A* do here?
- What went wrong?
- Actual bad goal cost < estimated good goal cost
- We need estimates to be less than actual costs!

Idea: Admissibility





Inadmissible (pessimistic) heuristics break optimality by trapping good plans on the fringe Admissible (optimistic) heuristics can still help to delay the evaluation of bad plans, but never overestimate the true costs

Admissible Heuristics

• A heuristic *h* is *admissible* (optimistic) if:

 $0 \le h(n) \le h^*(n)$ where $h^*(n)$ the true cost to a nearest goal

+Example:



Coming up with admissible heuristics is most of what's involved in using A* in practice.

Test Your Understanding

- Practice problem in breakout rooms
- Work for a couple of minutes independently, but then quickly start comparing progress – even if you're not done yet.

Optimality of A* Tree Search



Optimality of A* Tree Search

Assume:

- + A is an optimal goal node
- ✤ B is a suboptimal goal node
- + h is admissible



Claim:

✦ A will exit the fringe before B

Optimality of A* Tree Search: Blocking

Proof:

- + Imagine B is on the fringe
- Some ancestor n of A is on the fringe, too (maybe A!)
- ✦ Claim: n will be expanded before B
 - 1. f(n) is less or equal to f(A)

nBf(n) = g(n) + h(n)Definition of f-cost $f(n) \le g(A)$ Admissibility of h g(A) = f(A)h = 0 at a goal

Optimality of A* Tree Search: Blocking

Proof:

- ✤ Imagine B is on the fringe
- Some ancestor n of A is on the fringe, too (maybe A!)
- ✦ Claim: n will be expanded before B
 - 1. f(n) is less or equal to f(A)
 - 2. f(A) is less than f(B) -

g(A) < g(B) B is suboptimal

h = 0 at a goal

f(A) < f(B)

Optimality of A* Tree Search: Blocking

Proof:

- + Imagine B is on the fringe
- Some ancestor n of A is on the fringe, too (maybe A!)
- ✦ Claim: n will be expanded before B
 - 1. f(n) is less or equal to f(A)
 - 2. f(A) is less than f(B)
 - 3. *n* expands before B
- + All ancestors of A expand before B
- ✦ A expands before B
- ✦ A* search is optimal



 $f(n) \le f(A) < f(B)$

UCS vs A* Contours

 Uniform-cost expands equally in all "directions"

 A* expands mainly toward the goal, but does hedge its bets to ensure optimality





A* Applications

- Video games
- Pathing / routing problems
- Resource planning problems
- Robot motion planning
- Language analysis
- Machine translation
- Speech recognition



Creating Admissible Heuristics

- Most of the work in solving hard search problems optimally is in coming up with admissible heuristics
- Often, admissible heuristics are solutions to *relaxed problems*, where new actions are available





Inadmissible heuristics are often useful too

Example: 8 Puzzle



Start State

Actions

2

Goal State



- How many states?
- What are the actions?
- How many successors from the start state?
- What should the costs be?

8 Puzzle I

- Heuristic: Number of tiles misplaced
- Why is it admissible?
- h(start) = 8
- This is a *relaxed-problem* heuristic







Start State

Goal State

| | Average nodes expanded when the optimal path has | | | |
|-------|---|---------|-----------------------|--|
| | 4 steps | 8 steps | 12 steps | |
| JCS | 112 | 6,300 | 3.6 x 10 ⁶ | |
| TILES | 13 | 39 | 227 | |

8 Puzzle II

- What if we had an easier 8-puzzle where any tile could slide any direction at any time, ignoring other tiles?
- Total Manhattan distance
- Why is it admissible?
- h(start) = 3 + 1 + 2 + ... = 18





Start State

Goal State

| | Average nodes expanded when the optimal path has | | | |
|-----------|--|---------|----------|--|
| | 4 steps | 8 steps | 12 steps | |
| TILES | 13 | 39 | 227 | |
| MANHATTAN | 12 | 25 | 73 | |

8 Puzzle III

- How about using the *actual cost* as a heuristic?
 - Would it be admissible?
 - Would we save on nodes expanded?
 - What's wrong with it?



- With A*: a trade-off between quality of estimate and work per node
 - As heuristics get closer to the true cost, you will expand fewer nodes but usually do more work per node to compute the heuristic itself

Trivial Heuristics, Dominance

• Dominance: $h_a \ge h_c$ if

 $\forall n : h_a(n) \geq h_c(n)$

- Heuristics form a semi-lattice:
 - Max of admissible heuristics is admissible

 $h(n) = max(h_a(n), h_b(n))$

- Trivial heuristics
 - Bottom of lattice is the zero heuristic (what does this give us?)
 - Top of lattice is the exact heuristic



A* Graph Search Gone Wrong?



Consistency of Heuristics



- Main idea: estimated heuristic costs ≤ actual costs
 - Admissibility: heuristic cost ≤ actual cost to goal

 $h(A) \leq actual cost from A to G$

✦ Consistency: heuristic "arc" cost ≤ actual cost for each arc

 $h(A) - h(C) \le cost(A to C)$

i.e. if the true cost of an edge from A to C is X, then the h-value should not decrease by more than X between A and C.

- Consequences of consistency:
 - The f value along a path never decreases

 $h(A) \le cost(A to C) + h(C)$

+ A* graph search is optimal

Optimality of A* Graph Search

- Sketch: consider what A* does with a consistent heuristic:
 - Fact 1: In tree search, A* expands nodes in increasing total f value (f-contours)
 - Fact 2: For every state s, nodes that reach s optimally are expanded before nodes that reach s suboptimally
 - Result: A* graph search is optimal



Optimality

- Tree search:
 - A* is optimal if heuristic is admissible
 - UCS is a special case (h = 0)
- Graph search:
 - A* optimal if heuristic is consistent
 - UCS optimal (h = 0 is consistent)
- Consistency implies admissibility
- In general, most natural admissible heuristics tend to be consistent, especially if from relaxed problems



A*: Summary



A*: Summary

- A* uses both backward costs and (estimates of) forward costs
- A* is optimal with admissible / consistent heuristics
- Heuristic design is key: often use relaxed problems

