


Recap: Search

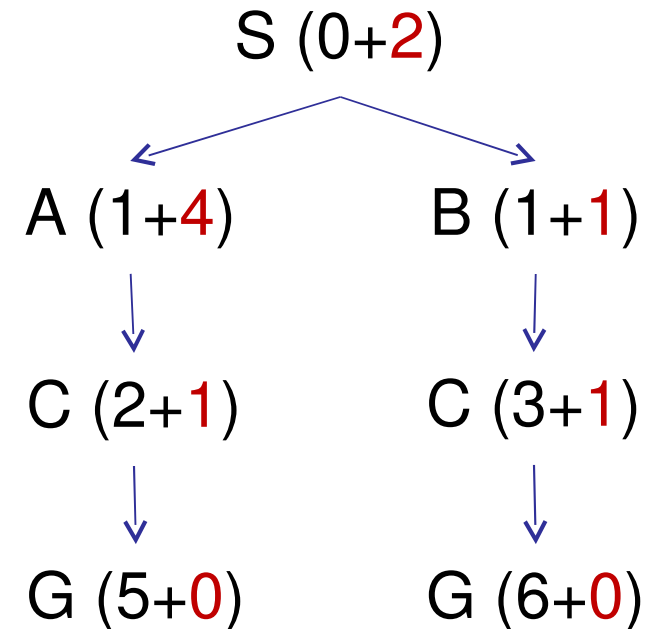
- **Search problem:**
 - States (configurations of the world)
 - Transition function: a function from states and actions to lists of (state, cost) pairs; drawn as a graph
 - Start state and goal test
- **Search tree:**
 - Nodes: represent plans for reaching states
 - Plans have costs (sum of action costs)
- **Search Algorithm:**
 - Systematically builds a search tree
 - Chooses an ordering of the fringe (unexplored nodes)

A* Graph Search Gone Wrong?


State space graph



Search tree



Consistency of Heuristics




- Stronger than admissibility
- Definition:
 - $\text{cost}(\text{A to C}) + h(\text{C}) \geq h(\text{A})$
 - $\text{cost}(\text{A to C}) \geq h(\text{A}) - h(\text{C})$
 - real cost \geq cost implied by heuristic
- Consequences:
 - The f value along a path never decreases
 - A* graph search is optimal

Optimality of A* Graph Search

Proof:

- New possible problem: some n on path to G^* isn't in queue when we need it, because some worse n' for the same state dequeued and expanded first (disaster!)
- Take the highest such n in tree
- Let p be the ancestor of n that was on the queue when n' was popped
- $f(p) < f(n)$ because of **consistency**
- $f(n) < f(n')$ because n' is suboptimal
- p would have been expanded before n'
- Contradiction!



Optimality

- Tree search:
 - A* is optimal if heuristic is admissible (and non-negative)
 - 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

Summary: A^*

- 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

Local Search Methods

- Tree search keeps unexplored alternatives on the fringe (ensures completeness)
- Local search: improve what you have until you can't make it better
- Generally much faster and more memory efficient (but incomplete)

Types of Search Problems


- Planning problems:

- We want a path to a solution (examples?)
- Usually want an optimal path
- *Incremental formulations*



- Identification problems:


- We actually just want to know what the goal is (examples?)
- Usually want an optimal goal
- *Complete-state formulations*
- Iterative improvement algorithms



Hill Climbing

- Simple, general idea:
 - Start wherever
 - Always choose the best neighbor
 - If no neighbors have better scores than current, quit
- Why can this be a terrible idea?
 - Complete?
 - Optimal?
- What's good about it?

Hill Climbing Diagram



- Random restarts?
- Random sideways steps?

Simulated Annealing

- Idea: Escape local maxima by allowing downhill moves
 - But make them rarer as time goes on

function **SIMULATED-ANNEALING**(*problem*, *schedule*) **returns** a solution state

inputs: *problem*, a problem

schedule, a mapping from time to “temperature”

local variables: *current*, a node

next, a node

T, a “temperature” controlling prob. of downward steps

current \leftarrow MAKE-NODE(INITIAL-STATE[*problem*])

for *t* \leftarrow 1 **to** ∞ **do**

T \leftarrow *schedule*[*t*]

if *T* = 0 **then return** *current*

next \leftarrow a randomly selected successor of *current*

$\Delta E \leftarrow$ VALUE[*next*] – VALUE[*current*]

if $\Delta E > 0$ **then** *current* \leftarrow *next*


else *current* \leftarrow *next* only with probability $e^{\Delta E/T}$

Simulated Annealing


- Theoretical guarantee:
 - Stationary distribution: $p(x) \propto e^{-\frac{E(x)}{kT}}$
 - If T decreased slowly enough, will converge to optimal state!
- Is this an interesting guarantee?
- Sounds like magic, but reality is reality:
 - The more downhill steps you need to escape, the less likely you are to ever make them all in a row
 - People think hard about *ridge operators* which let you jump around the space in better ways

Beam Search

- Like greedy hillclimbing search, but keep K states at all times:




Greedy Search



Beam Search


- Variables: beam size, encourage diversity?
- The best choice in MANY practical settings
- Complete? Optimal?
- Why do we still need optimal methods?

Genetic Algorithms



- Genetic algorithms use a natural selection metaphor
- Like beam search (selection), but also have pairwise crossover operators, with optional mutation
- Probably the most misunderstood, misapplied (and even maligned) technique around!

Example: N-Queens



- Why does crossover make sense here?
- When wouldn't it make sense?
- What would mutation be?
- What would a good fitness function be?