

Search Methods

Tree Search	Expand nodes using gradients
Graph Search	Avoids revisiting nodes and ensure efficiency

Uninformed search

Uniform cost search	aka Cheapest-first Add visited node to Explored and add its neighbors to the frontier Visit cheapest node in the frontier Move to next cheapest if all neighbors are explored
Iterative deepening	Iteratively calls depth-limited search Initialize frontier with root If not goal, remove from frontier and expand If at the end of depth, terminate and start over Repeat until goal is found Guaranteed to find optimal path More efficient than DFS Time complexity: $O(b^d)$ Space complexity: $O(d)$
Bidirectional	Finds shortest path of a graph

Informed Search

Best-first search	Choose unvisited node with the best heuristic value to visit next Same as lowest cost BFS
Greedy best-first search	Uses heuristic to get the node closest to the goal Bad performance if heuristic is bad Does NOT consider cost of getting to the node

Informed Search (cont)

A* search	Always expand to node with minimum f Evaluate cost of getting to goal using heuristics $f(n) = g(n) + h(n)$ where g is cost to get to n Uses priority queue
Heuristics	Cost to get to the goal
Admissible heuristic	Optimistic model for estimating cost to reach the goal Never overestimates $h(n) \leq c(n)$ where c is actual cost
Consistent heuristic	$h(n) \leq c(n, a, n') + h(n')$ Immediate path costs less than longer path Consistent \implies Admissible

Consistent heuristic

Consistent heuristics

- A heuristic is **consistent** if for every node n , every successor n' of n generated by any action a ,
$$h(n) \leq c(n, a, n') + h(n')$$
- If h is consistent, we have
$$\begin{aligned} f(n) &= g(n) + h(n) \\ &= g(n) + c(n, a, n') + h(n') \\ &\geq g(n') + h(n') \\ &= f(n') \end{aligned}$$
- i.e., $f(n)$ is non-decreasing along any path.
- Theorem: If $h(n)$ is consistent, A* using GRAPH-SEARCH is optimal




Adversarial Search

Hill climbing	Method of local search Only move to neighbors to find the max Does NOT guarantee to find optimal
Simulated annealing	Method of local search Combine hill climbing and random walk Always find global max

Adversarial Search (cont)

Local beam	Generate k random states Generate successors of all k states If goal stop; else, pick k best successors and repeat Different from hill-climbing since information is shared between k points
Genetic algorithm	Cross-Over and mutation Decomposes strands of DNA and permute Produces children by: Selection, Cross-over, Mutation

Minimax Tree



⊙ Max node
□ Min node

α-β Pruning

```
function alphabeta(node, depth, α, β,
maximizingPlayer)
    if depth = 0 or node is a terminal node
        return the heuristic value of node
    if maximizingPlayer
        v := -∞
        for each child of node
            v := max(v, alphabeta (child,
depth - 1, α, β, FALSE))
            α := max(α, v)
            if β ≤ α
                β cutoff (prune)
        return v
    else
        v := ∞
        for each child of node
            v := min(v, alphabeta (child,
depth - 1, α, β, TRUE))
            β := min(β, v)
            if β ≤ α
                α cutoff (prune)
        return v
```

SAT

P	Q	¬P	P∨Q	P∧Q	P⇒Q	¬P∨Q
False	False	True	False	False	True	True
False	True	True	True	False	True	True
True	False	False	False	False	False	False
True	True	False	True	True	True	True

Propositional SAT: Graph coloring

At least 1 of K per i	$(C_i, 1 \vee C_i, 2 \vee \dots \vee C_i, k)$ $O(n)$ clauses
$1 \geq \text{color per } i$	$\forall k \neq k' (\neg C_i, k \vee \neg C_i, k')$ $O(n^2)$
If node i and j share an edge assign different colors	$\forall x \in k, (\neg C_i, x \vee \neg C_j, x)$