

Class 1 notes

January 13, 2011

Decision making problem:

Update state, formulate goal, formulate problem, then search for solution (action to take).

1 Russel and Norvig Chapters 3 & 4

1.1 Types of planning Problems:

- Deterministic, fully observable
 - Agent knows state it will be in
- Non-observable
 - Agent has no idea where (in state space) it is
 - Sometimes can still choose good actions, even without knowing anything about our state
- Partially observable
 - Observations provide info about state
 - Solution is a policy (what action to take, depending on what we see).
 - Often interleave search, execution.
- Unknown state space
 - Need to explore the state space, in addition to searching for a solution.

(Look at vacuum world example from Russel & Norvig section 3.2)

1.1.1 8-puzzle example problem

(from Russel & Norvig section 3.2)

- We ignore the intermediate problems of the tiles - this kind of simplifying of the state space is common when trying to solve complex problems.
- This problem is NP-hard. In simple terms, this is because we aren't able to solve it by solving several smaller instances of the problem.

1.2 Tree Search

- Recursively expand states to their possible next states, searching a tree of all possible states.
- Many algorithms are variants on tree search, but expand the states in a different order.
 - Breadth first search (BFS) (R&N pg. 73)
 - * Always finds lowest cost path first
 - * Requires more memory
 - Depth first search (DFS) (R&N pg. 75)
 - * Might find a non-optimal solution first
 - Iterative Deepening Depth First Search (R&N pg. 78)
- Evaluating search strategies
 - Criteria
 - * Completeness: does it always find a solution if there is one?
 - * Time complexity: how long it takes.
 - * Space complexity: how much memory it needs.
 - * Optimality: does it always find the least-cost solution?
 - Time and space complexity is defined in terms of some properties of the problem
 - * b: maximum branching factor of tree
 - * d: depth of least-cost solution
 - * m: maximum depth of state space
 - BFS and iterative deepening are complete
 - All strategies have exponential time complexity
 - For a more detailed comparison of the search strategies, see Fig. 3.17 on R&N pg. 81
- Avoiding repeated states: maintain a list of visited states so we don't repeat work

1.3 Best First Search

- Define some notion of “desirability” $h(n)$ for each node in the tree
- Expand the search tree in order of desirability
- Example: in path planning on a graph, the straight line distance can be used as “desirability”
- Greedy best first search is not complete; it can get stuck in loops

1.4 A*

- Variation on best first search.
- Uses cost to get to next node, plus $h(n)$ for that node, to determine which node to expand next.
- This heuristic is “admissible”, meaning it never overestimates cost
- Can prove optimality of A* using proof by contradiction (R&N pp. 97-99)
- Takes a lot of memory.
- Admissible heuristics for 8-puzzle
 - # of misplaced tiles
 - total Manhattan distance (always better than the # of misplaced tile)

2 LaValle Chapter 2

2.1 Discrete feasible planning

- Definition of a planning problem:
 - State space X .
 - For each state $x \in X$, there is an action space $U(x)$.
 - State transition function $f(x, u)$ that, given a current state and action, gives the next state.
 - An initial state $x_1 \in X$
 - A set of goal states $X_G \subset X$
- Forward search (LaValle pg. 33)
 - This is the idea of “Tree search” from R&N, but in a very nice and general notation.
 - Separate states into 3 types:

- * Unvisited: haven't expanded this state yet.
- * Dead: state and all of its children have been completely expanded.
- * Alive: haven't completely expanded all of this node's children.
- Expand the search tree in some order, keeping track of which states are unvisited, dead, and alive.
- LaValle uses the term “systematic” to mean the same thing as R&N's “complete”
- Dijkstra
 - Systematic, optimal planner.
 - Special form of dynamic programming.
 - A* can be viewed as an extension to Dijkstra
 - * Adds a heuristic using the an estimated cost-to-go, to guide the search towards paths likely to be fast
- IDA*
 - Like iterative deepening, but iteratively increases the value of the admissible heuristic, not the depth.
 - Runs DFS within each iteration.
 - Incrementing by too much between iterations can cause you to find a sub-optimal solution. Iterating by too little can cause running time to increase because you will have to run many iterations.
 - Wikipedia article has details of the algorithm: http://en.wikipedia.org/wiki/IDA*
- Value Iteration
 - Iterative way to compute optimal cost-to-go functions over X .
 - Essentially the same as the Bellman-Ford algorithm.
 - Works even with negative edge costs (unlike Dijkstra's algorithm).
 - Computes optimal paths of length k from all source states to all destination states.
 - Can add “termination states”, and then solve for paths of unbounded length.
 - * Keep running value iteration until the value of each node becomes stationary (doesn't change between iterations).
 - * Once the algorithm has terminated, just greedily follow the value function to get the optimal path.
 - * Doesn't work if the graph has negative cycles.