

Midterm 1 Study Guide

Helpful resource: <https://www.stepbystepml.com>. You can run many of the search algorithms we learned there! This is a good way to check your answers for the search questions below.

Agents, Environments, and Search


1. Define the term *agent* and give an example of an agent.
2. What properties make environments challenging to model in the kinds of search problems that we've been discussing?
3. What do the terms *optimal* and *rational* mean in relation to agents?
4. What is a *solution* in a search problem?

Solving Problems by Searching

5. Define in your own words the following terms: *state*, *state space*, *search tree*, *search node*, *goal state*, and *action*.
6. You have three pitchers, measuring 12 gallons, 8 gallons, and 3 gallons, and a water faucet. You can fill the pitchers completely, pour them completely from one to another, or empty them. You need to measure exactly 1 gallon into the 3-gallon pitcher.
 - a. List the possible actions for this problem.
 - b. Describe how you would represent states for this problem.
 - c. Draw a portion of the state space that shows what series of actions and states would lead to this result.

Search Algorithms and Properties

7. Using the basic state space search algorithm, we can do breadth-first search, depth-first search, uniform cost search, or A* search. What changes do you make to the algorithm in each case?
8. How do depth-limited search and iterative deepening search work?
9. We compare search methods based on completeness, optimality, time complexity, and space complexity.
 - a. Define what it means for a search algorithm to be *complete*.
 - b. Define what it means for a search algorithm to be *optimal*
 - c. What are the advantages and disadvantages of breadth-first search?
 - d. What are the advantages and disadvantages of depth-first search?
10. Consider the graph given below. List the nodes, in order, that will be visited by the search algorithms as they try to find the goal node, labeled **Z**, beginning from the root node, labeled **a**.
 - a. Depth-first search (assume children are added in right to left order).
 - b. Depth-bounded search, with bound 3
 - c. Breadth-first search (assume children are added in left to right order)



Informed Search

11. What is informed search? Define the term *heuristic* and describe how heuristics are used in informed search.
12. Explain the difference between uniform cost search, greedy best-first search, and A* search.
13. In A* search we calculate $f(n)$, which is the estimated cost of the cheapest solution through node n , and is given by the equation:

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

What are $g(n)$ and $h(n)$?

Adversarial Search

14. What is a *policy* in an adversarial search problem?
15. What does *zero-sum* mean in the context of games?
16. Consider the game tic-tac-toe. Assume that the current state of the game is as shown below, and it is X's turn.

```

-----
| X | 0 | 0 |
-----
|   | X |   |
-----
|   | 0 |   |
-----

```


From this node, draw the game tree showing all of X's possible immediate next moves.

17. Consider the following evaluation function for tic-tac-toe: $((3 * X2) + X1) - ((3 * O2) + O1)$

where "Xn" is the number of rows, columns, or diagonals with exactly n X's and no O's, and "On" is the number of rows, columns, or diagonals with exactly n O's and no X's.

Score of all of the leaves on the tree you created above.

18. Consider the following game tree, where P1 represents player1 and P2 represents player2. Using alpha-beta pruning, show the values of alpha and beta at each non-leaf node. Remember that MAX layers update alpha and MIN layers update beta.



Uncertainty in Search

19. Why is it important to discount future rewards in search problems that involve uncertainty?
20. Explain some sources of uncertainty in search problems.
21. In a Markov Decision Process, what does the transition function take as arguments, and what does it return?
22. Snake-eyes is a game that uses a deck with three cards: 1, 2, 3. Your goal is to get the highest score. You can draw as many cards as you like (cards are put back in the deck and reshuffled), but if you are dealt two 1s, the game ends and you lose all your points.
 - a) What is the transition function for this game?
 - b) What is the expected reward for the policy *[Draw, Draw, Pass]*?

Reinforcement Learning

23. Define *reward*, *policy*, *discounting*, *exploration*, and *exploitation* in the context of reinforcement learning.
24. Define the variables in the following equation: $\text{sample} = R(s, \pi(s), s') + \gamma V^\pi(s')$
25. Consider a pair of slot machines whose rewards are only partially known: after playing each machine 3 times, you have observed better rewards for Slot A. How would the relative weighting of exploration and exploitation in a reinforcement learning agent affect how much each machine is played?