Potential Midterm Exam Questions

1. What are the four ways in which AI is usually viewed? Which of the four is the preferred view of the authors of our textbook?
2. What does each of the lettered items in the following diagram represent?

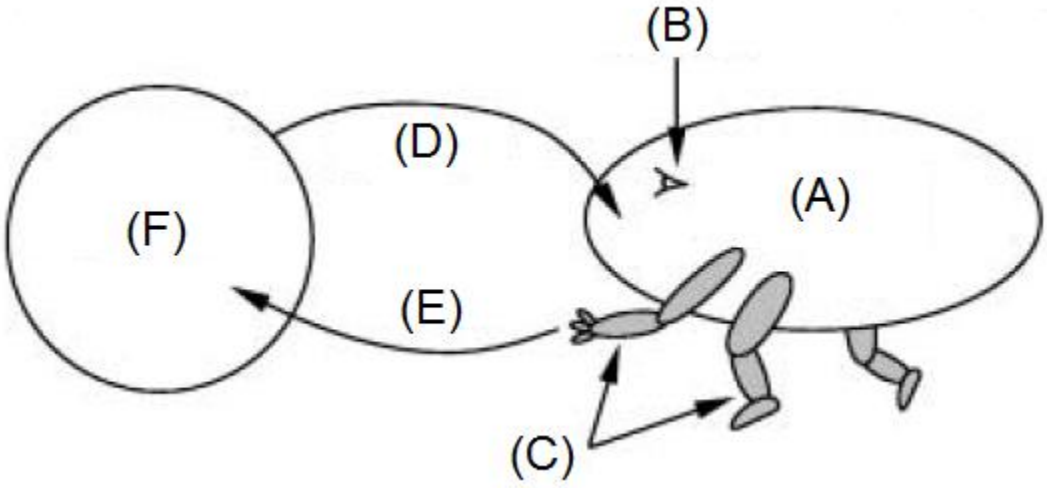

3. Rational agents are often specified by a PEAS description. What do each of the letters in the acronym "PEAS" stand for?
4. List, in order of increasing generality, the four basic types of agents we discussed in class.
5. You are in Arad and you want to find a path that will get you to Bucharest. Using the following map, construct the first two levels of a search tree diagram that would be used to solve the problem.

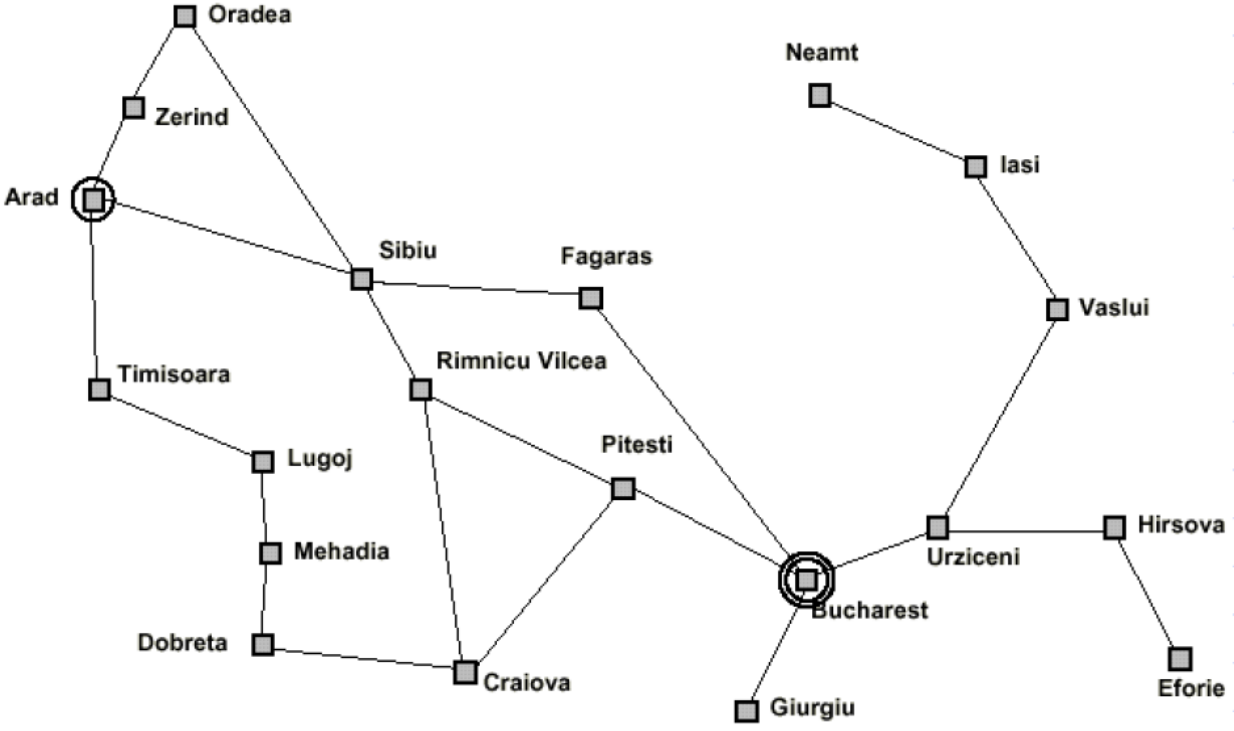

6. For the following trees, number the nodes in the order in which they are explored for breadth-first search and for depth-first search.

Breadth-first

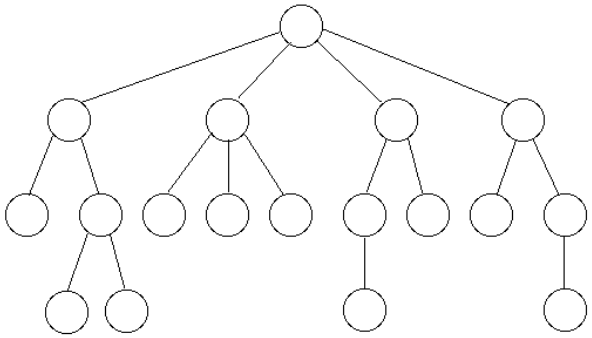

Depth-first

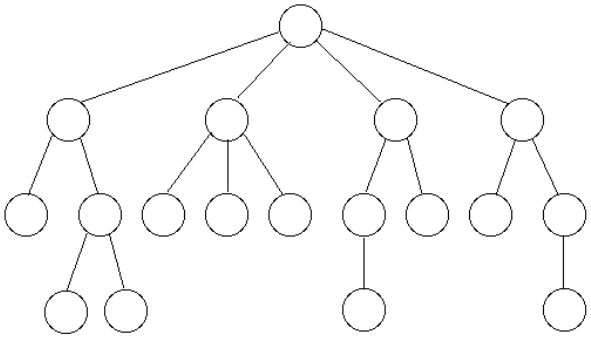

7. Fill in the following table:

Environment type \ Game	Solitaire	Chess	Soccer	Backgammon	Internet Shopping	Automated Uber
Observable (fully/partial)						
Deterministic/Stochastic						
Episodic/Sequential						
Static vs Dynamic						
Discrete/Continuous						

8. To evolve from uninformed search to intelligent search, one must add a(n)__ _____ function, which is simply an estimate of _____.
9. What must be true about a heuristic in order to call it admissible?
10. Regarding Minimax:
    a) Is it complete? (Is your answer conditional?)
    b) Is it optimal?
    c) What is the time complexity?
    d) What is the space complexity?
11. Give the following properties for Min-Max algorithm
    a) Complete?
    b) Optimal?
    c) Time complexity?
    d) Space Complexity?

12. Do minmax with alpha beta pruning for the following tree?

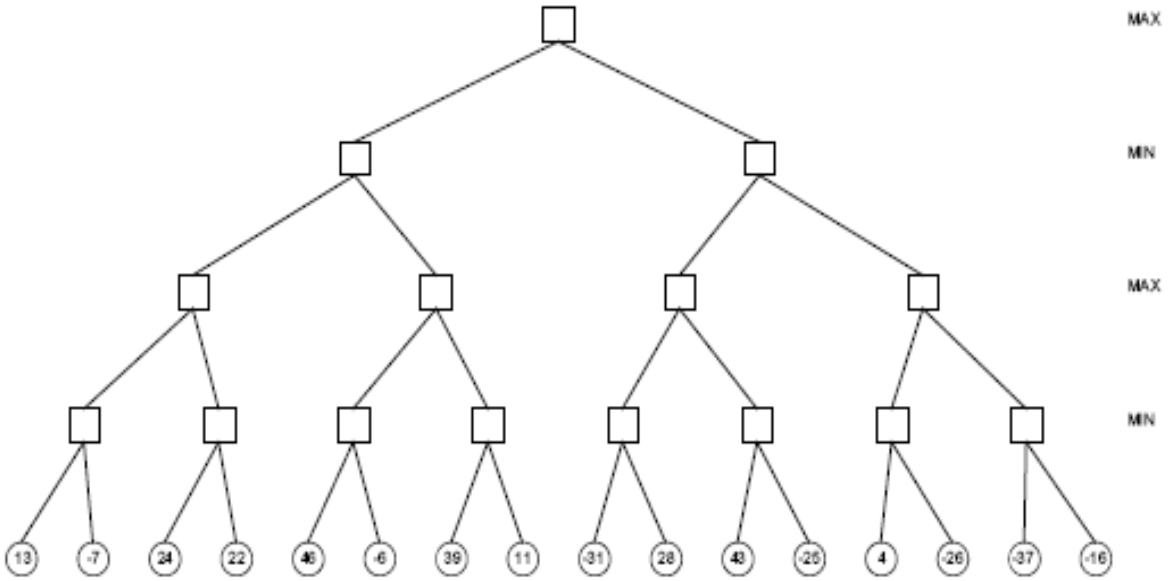

13. Find a Path from Arad to Bucharest using Greedy Search technique [for heuristic apply straight line distance]. Define the four characteristics Optimality, Space and Time Complexity and Completeness for Greedy Search?

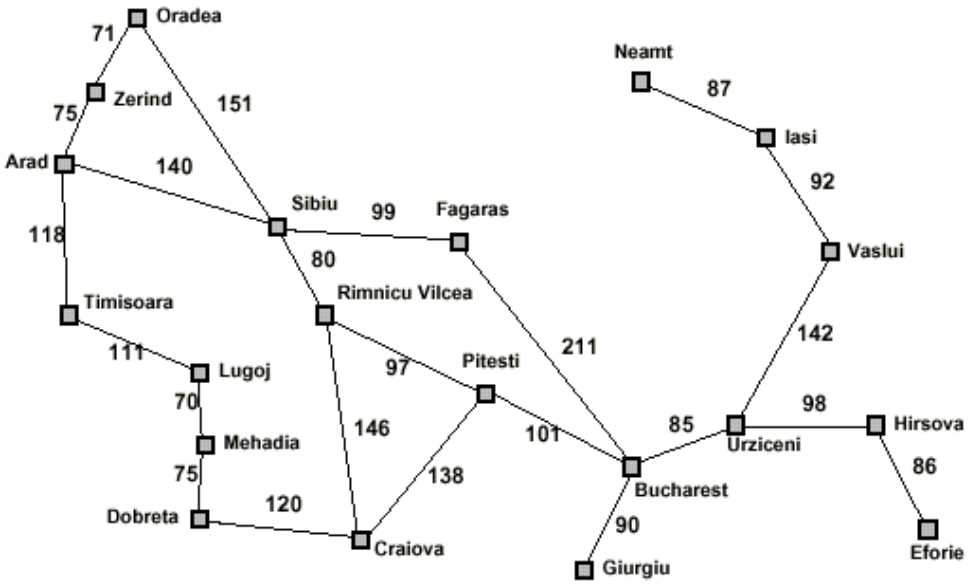

Straight−line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

14. Find a Path from Arad to Bucharest using A* search?

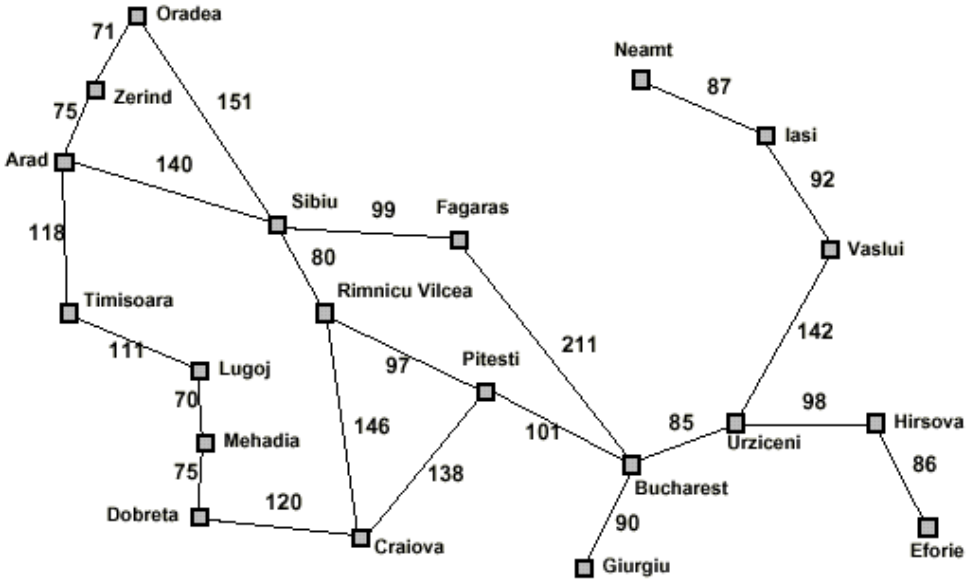

Straight-line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

15. Define the four characteristic Optimal, Size Complexity, Time Complexity and Completeness for A* search? Proof that A* is optimal? Explain what is admissible Heuristic?
16. How the breadth first searches work? Give example using the following graph? Give characteristic of this search?

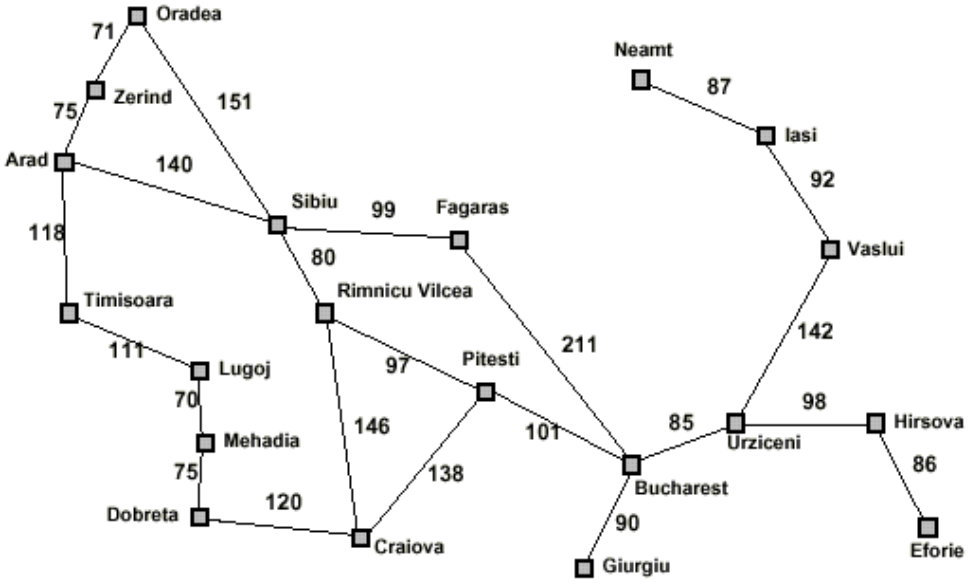

Straight-line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

17. What is depth first search? Give example with the following graph. Define the characteristic of this search?

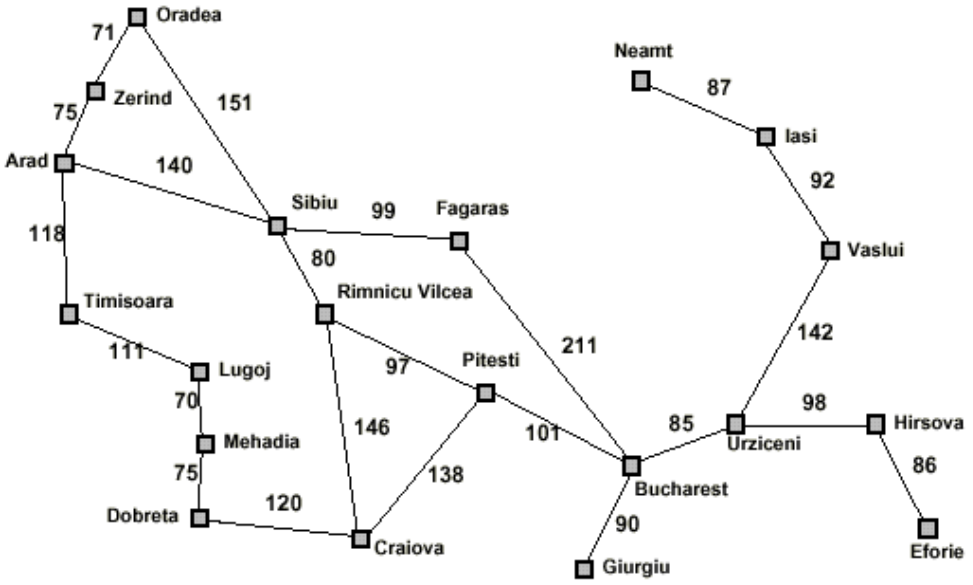

Straight-line distance to Bucharest	
Arad	366
Bucharest	0
Craiova	160
Dobreta	242
Eforie	161
Fagaras	178
Giurgiu	77
Hirsova	151
Iasi	226
Lugoj	244
Mehadia	241
Neamt	234
Oradea	380
Pitesti	98
Rimnicu Vilcea	193
Sibiu	253
Timisoara	329
Urziceni	80
Vaslui	199
Zerind	374

18. Give properties of the depth limited search and iterative deepening search? Give pseudo code for both of the searches.
19. What are the environment types we have seen in our textbook? Give one example of each environment type.
20. Describe the Turing Test. How does it compare with the current direction of artificial intelligence?
21. What four categories do the authors classify the views of artificial intelligence? Which view do the authors follow? Why?
22. What does PEAS stand for? Use it to design an automated vacuum cleaner.
23. Describe and compare the four main types of agents.
24. Compare and contrast breadth-1st search and depth-1st search in terms of completeness and optimality. When might one be used over the other? What are the problems with depth-1st search and how might they be compensated for?
25. Compare and contract greedy search and A* Search. What heuristics do they use? What advantages does A* search have over greedy search?
26. What makes an admissible heuristic? Why is this important for search? Explain why using the number of squares in the wrong position for the 8-puzzle is a bad idea even though it is admissible. What is a better heuristic?
27. Games can be classified into four categories using two characteristics. What are the characteristics and some examples of games that fall into the categories?
28. Using minimax, explain why chess is a more \interesting" game than tic-tac-toe. What method can be applied to minimax to improve the results? How does this affect the game of chess?
29. Enumerate the 4 components of the PEAS description of a rational agent and provide 2 examples for each component for the automated vacuum cleaner (i.e. Roomba) example.

30. What assumptions are made about the task environments for which the problem can be solved via search? Name at least two assumptions.
31. Given the following node of a search tree for the 8-puzzle problem, draw the output of the successor function on this node. Assume that repeated states are allowed.

7	2	4
5		6
8	3	1

32. Prove that uniform-cost search is a special case of A* search.
33. Assume there is an adversarial search problem where at every turn each player is given two choices. Next, consider the evaluation values at d=4 to be 12, 3, 13, 14, 10, 4, 1, 9 (left to right). Assuming that max's turn is next, and that both players are playing optimally, what is max's value?
34. Give two drawbacks to using a depth-first search algorithm.
35. What is the big-O time complexity for the greedy best-first search algorithm?
36. Why is the straight-line distance between two points always an admissible heuristic for the path cost between the two points? Assume that the path cost just takes into account the distance along the path.
37. Consider a tree search algorithm is to be solved via breadth first search. What is the worst-case memory footprint if at every stage of the search there are 3 choices and that the goal state is at a depth of 3? Assume that each node is 1000 bytes.
38. What are the 4 categories of AI?
39. Describe the Turing Test.
40. What is a rational agent?
41. What is the difference between a reflex-based agent and a goal-based agent?
42. Define the following types of environments, and give an example for each:
43. What is the basic difference between uninformed and informed search?
44. Describe A*. Is it optimal?
45. Describe Minimax. When is it ideal?
46. What is alpha-beta pruning?
47. What change do stochastic games require of the basic minimax? Describe how this change works.
48. Explain why problem formulation must follow goal formulation?
49. Consider a state space where the start state is number 1 and the successor function for state n returns two states, number 2n and 2n+1.
    a) Draw the portion of the state space for states 1 to 15.
    b) Suppose the goal state is 11. List the order in which nodes will be visited for breadth-first search, depth-limited search with limit 3, and iterative deepening search.
    c) Would bidirectional search be appropriate for this problem? If so describe in details how it works.
50. Describe a state space in which iterative deepening search performs much worse than depth-first search (for example, O (n^2) vs. O (n)?
51. Trace the operation of A* search applied to the problem of getting to Bucharest from Lugoj using the straight-line distance heuristic. That is, show the sequence of nodes that the algorithm will consider and the f and h score for each node?

52. The heuristic path algorithm is a best-first search in which the objective function is f (n) = (2-w) g (n) +wh (n). For what values of w is this algorithm guaranteed to be optimal? (You can assume that h is admissible). What kind of search does this perform when w=0? When w=1? When w=2?
53. Prove that Breadth search is a special case of uniform-cost search.
54. Define a game?
55. What is the main difference between minimax and Alpha-Beta algorithm?
56. Both the performance measure and the utility function measure how well an agent is doing. Explain the difference between the two?
57. Can there be more than one agent program that implements a given agent function?
58. Use breadth first search to create a tree from the following maze. Circle objects as they come off of the fringe so that fringe order can be checked.

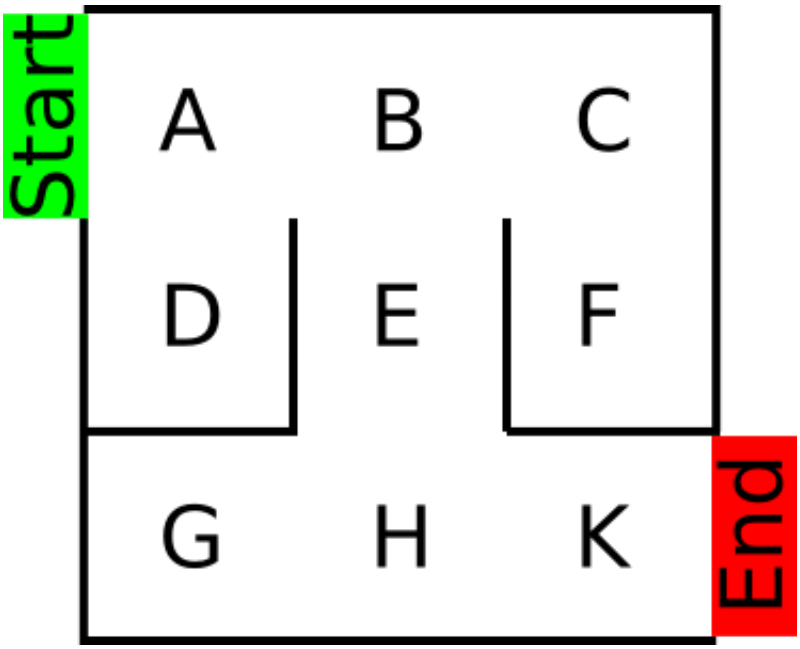

59. What is the advantage of counting the total cost to the current node when calculating a heuristic?
60. Under what conditions is breadth first search optimal? Why?
61. Is depth first complete? Why?
62. What is greedy search and its major disadvantage?
63. Explain minimax theory for game play.
64. Draw a Simple Reflex Agent.
65. What makes a heuristic admissible?
66. Explain depth limited search.
67. Explain completeness, time complexity, space complexity, and optimality with regard to search.
68. List the PEAS of the Wumpus World.
69. Fill in the chart characteristics

Characteristic	**Backgammon**
Observable (fully/partial)	
Deterministic/Stochastic	
Episodic/Sequential	
Static vs Dynamic	
Discrete/Continuous	

70. How would you logically determine the location of the Wumpus?

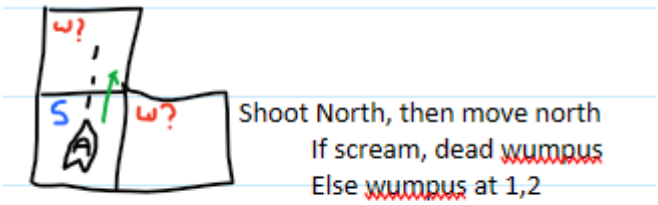

71. You are given an a x b Vacuum World. Formulate the problem.
72. PEAS for self-driving taxi.
73. What is the time and space complexity of breadth first search?
74. Explain PEAS for a self-driving car.
75. Describe the differences between:
    a. Deterministic vs Stochastic
    b. Episodic vs Sequential
    c. Static vs Dynamic
    d. Discrete vs Continuous
76. Name 4 Informed Search Strategies and describe them
77. Draw a simple reflex agent interacting with its environment. Explain its workings in terms of a reflex agent for the vacuum world problem.
78. For the Breadth-First, Depth-First, and A* algorithms, discuss if they are, in all situations, complete, optimal, and describe their time and space complexity
79. Describe alpha-beta pruning and how it works to improve the time complexity of minimax.
80. Describe how games of chance are evaluated by a minimax game player
81. Name the two methods a genetic algorithm uses to generate new states from selected ones and briefly describe how they work.
82. Name the five components used to define a search problem.
83. Describe the environment for a crossword puzzle using the seven properties of task environments. For each property, explain why this is the case.
84. What does it mean for a function h(n) to be an admissible heuristic?
85. What qualities separate a good heuristic from a bad heuristic? Give an example of a permissible heuristic for the Tiles game (Picture assembly game) and explain.
86. Explain and list the three major components that make up a genetic algorithm, how they work together to conduct a search and properties of this search.
87. What differentiates an uninformed search from an informed search? Give 3 examples of each.
88. Define state, state space, search tree, search node, goal, and action.
89. What are three types of logic?
90. Why do people have a hard time solving the missionaries and cannibals puzzle, given that the state space is so simple?
91. Consider a state space where start state is number 1 and each state k has two successors: numbers 2k & 2k+1. How well would bidirectional search work on this problem? What is the branching factor?
92. What does it mean for an A.I to act humanely or to act rationally? What does it mean for an A.I to think humanely or to think rationally?
93. What is the effect of heuristic accuracy on performance?
94. How are utility-based agents different from goal-based agents?
95. Can an agent that keeps no history of its precept sequence be rational? Explain.
96. What are the conditions for optimality in heuristic search strategies?

97. For each of the following activities, give a PEAS description of the task environment. (a) Playing soccer; (b) Walking to a shop; (c) Eating lunch; (d) Taking a quiz.
98. Give two examples each of single-agent and multi-agent systems.
99. How are problem solving agents solving contingency problems different from the ones solving exploratory problems?
100. What does the Acronym PEAS stand for? Explain each and give an example from the Wumpus World
101. Describe the Declarative Approach we can take when creating a logical agent. What are the steps?
102. What is the difference between regular A* search and SMA* search? Which algorithm has better performance?
103. What is the problem encountered when attempting to create a terminal test (i.e. a test for winning the game) for interesting games like chess? What is the solution to this problem?
104. What are some specific differences between search problems and game-playing problems? Do these differences increase or decrease the difficulty of solving such problems?
105. Define the PEAS for a tutor.
106. Random Search Game: Two players (A and B) each occupy a space on an n X n grid. A wins the game if they can find B in the grid within 2n turn cycles (otherwise, B wins). A starts in the bottom left corner, and B chooses their starting point. Both players can see each other's starting location, but once the game starts the B's location is hidden (but B can see A). A is allowed to move up/down/left/right up to n/2 times (round down) per turn, but must move at least one space per turn. B can move up to n/3 times per turn, but is not required to move at all. (Example: For 6X6 grid, A has 12 turns to find B. A can move up to 3 times per turn, and B can move up to 2 times per turn.) Create a formulation of a search problem for this game as player A.
107. Explain how competitive search is different from normal search.
108. Place each of the following games in the correct slot in the table below: Scrabble, Uno, Craps, Tic-Tac-Toe, Battleship, Risk

	Deterministic	**Chance**
Perfect Info		
Imperfect Info		

109. Wumpus Expansion: Recall/reformulate the characterization of the game of Wumpus (Fully/Partially Observable, Deterministic/Stochastic/etc). The grid is now 7X7, and there are no pits. Instead, there are 4 wumpuses (I guess? Not sure of the plural form) and you now have 3 arrows. You can now smell a wumpus two squares away as opposed to just one, and you can tell whether the wumpus is one or two squares away. At the end of your turn (after movement or shot), each wumpus moves in a random direction (up/down/left/right). How would the characterization of the game change?
110. Work through the action of the minimum-conflicts heuristic on a randomly placed board of the 5-queens puzzle
111. Formulate the 8-puzzle problem

112. What is the worst case time and space performance of iterative deepening search?
113. Describe a state space in which iterative deepening search performs worse than depth first search
114. Show that 8-puzzle states are divided into two disjoint sets, such that any state is reachable from another state in the same set, and no state is reachable from a state in the other set. Devise a procedure to decide what set a given state is in (3.4 in book).
115. Discuss how well the standard approach to game playing would apply to games such as tennis, pool, croquet, which take place in a continuous physical space.
116. Suppose you have an oracle OM(s) that correctly predicts the opponent's move in any state. Using this, formulate the definition of a game as a single-agent search problem. Describe an algorithm for finding the optimal move.
117. Fill in the blanks: _____ are always trying to maximize their _____ based on a _____ they receive from the environment.
118. Give an example of an environment that is dynamic but at the same time has discrete actions. Explain.
119. Differentiate between and explain search cost and solution cost.
120. Prove the scenario in which breadth-first search optimal.
121. Match the search algorithm with its corresponding data structure and explain how it uses them to carry out their searches: Breadth-first search, Depth-first search, Uniform-cost search : Priority queue, Queue, Stack
122. Assuming a thorough search, i.e., there is no goal node, determine the order the nodes will be expanded in above for depth-first search, greedy best-first search, and iterative-deepening search.

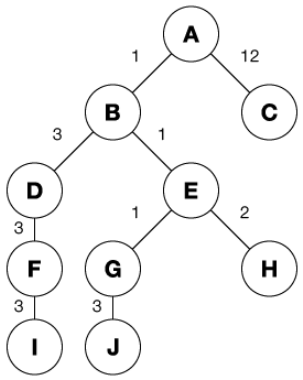

123. Determine, describe, and argue the admissibility of a possible heuristic function for the n-queens problem.