CS4480, Fall 2015 – Artificial Intelligence
Quiz 1
Due October 16, 2015

- This is a take-home exam. All answers must be your own work. You may:
  - Discuss this exam with me,
  - Use our textbook, class notes and slides.
  - **No other sources are allowed.**
- To receive full credit, **show your work** and **write legibly**.
- If you need clarification about any of the problems, please ask me.

Name: _____

1. Define artificial intelligence in your own words.

2. Our book gives four categories of AI definitions. Please state these four categories.

3. Explain the Turing test.

Is it a good test? (Justify your answer.)

3. What is a rational agent?

4. Specify the task environment, using PEAS, for a web-based shopping agent.

5. We discussed 6 properties of task environments for rational agents. Describe a task environment for a chess-playing agent, where the chess game is played with a clock. (Be sure to state assumptions and justify your choices.)

6. Explain the difference between:
   a. simple reflex agents and model-based agents:

   b. goal-based agents and utility-based agents:

7. Give the **initial state, goal test, successor function, and cost function** of the following two-player road trip game:

> You start with a set X of actors. The first player names an actor x1 who is an element of X, and the other player names an actor x2 who is also an element of X (but not x1) who has appeared in a movie with x1. The first player must then name an actor x3 who has appeared in a movie with x2, and so on. So the two players generate a sequence of actors such that each actor in the sequence has co-starred with the actor immediately preceding. A player loses when it's his or her turn to move, and s/he cannot name an actor of X who hasn't been named before. You are given a set of actors X, with complete information on who has appeared in a movie with whom, and you want to determine how to win the game.

8. Given the following tree, starting at the root, show the **frontier** at each node expansion as the tree is traversed using BFS and DFS, until node **f** is in the front of the queue.
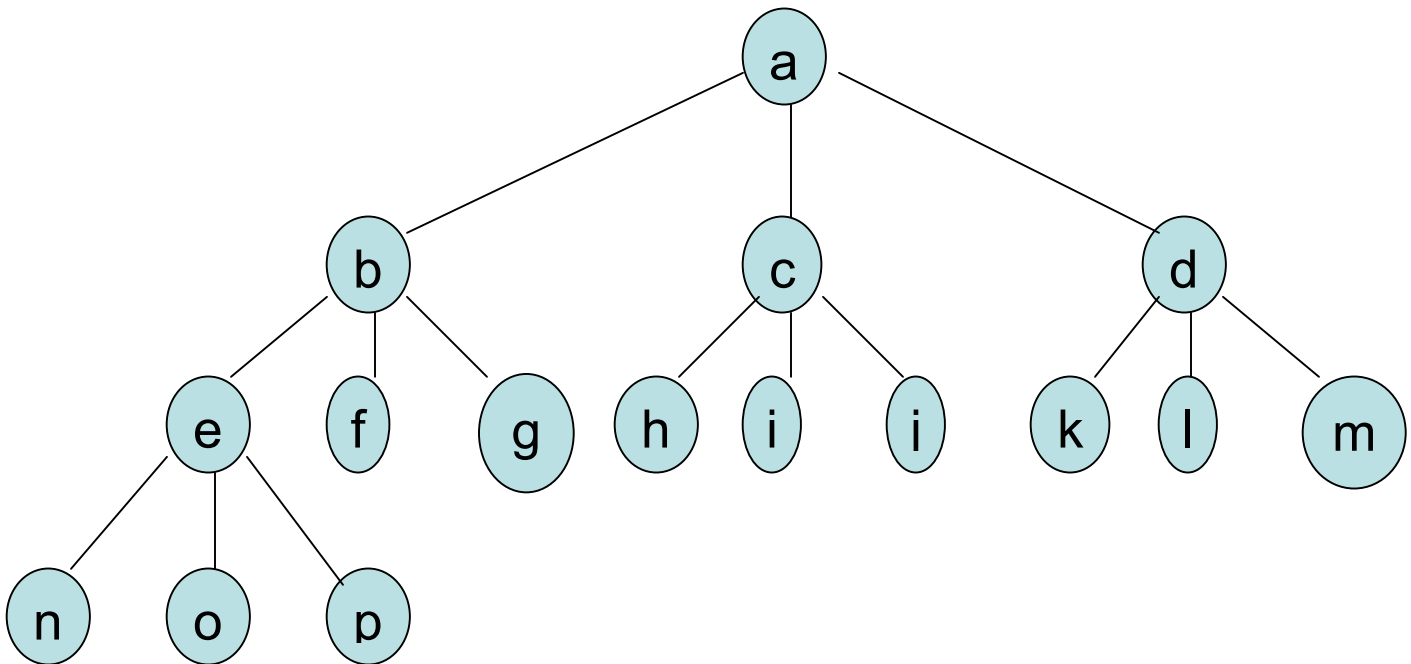
Breadth First Search Frontier at each node until node f is at the front of the queue:

\

Depth First Search Frontier at each node until node f is at the front of the queue:

9. How does Iterative Deepening improve over Depth-First Search?

10. In A* search:
    a. What are the functions g and h? (Please be precise.)

    b. By what criterion are the paths in the FRINGE ordered?

    c. What search algorithm is equivalent to A* when g = 0?

    d. What search algorithm is equivalent to A* when h = 0?

11. Suppose you are solving a problem with:
- most goal states lie very deep in the state space
- the state space is not finite
- there are many goals

Describe the behavior you would expect from the following algorithms. Please organize your answers carefully and show that you understand the search algorithms.
- Depth-first search
- A* search
- Iterative-deepening search
- Best-first search
- Hill climbing