University of Pittsburgh
CS 1571 Introduction to Artificial Intelligence — Final exam
Professor Milos Hauskrecht — December 13, 2001

Final exam

Instructions

- This quiz contains 7 problems. You have 100 minutes to earn 100 points.
- The quiz is a closed-book exam.
- Write your solutions in the space provided. If you need more space use additional sheet. Do not put part of the answer to one problem to the space for another problem. Remember to write your name and problem number on any additional sheet you use.
- Do not spend too much time on any one problem. Read them all through first and attack them in the order that allows you to make the most progress.
- Show your work as partial credit will be given. You will be graded not only on the correctness of your answer but also on the clarity with which you express it. Be neat.

Problem	Points	Grade
1	15	
2	20	
3	10	
4	10	
5	20	
6	15	
7	10	
Total	100	

Name:

Problem 1 (15 points). Consider the game search tree in the figure below

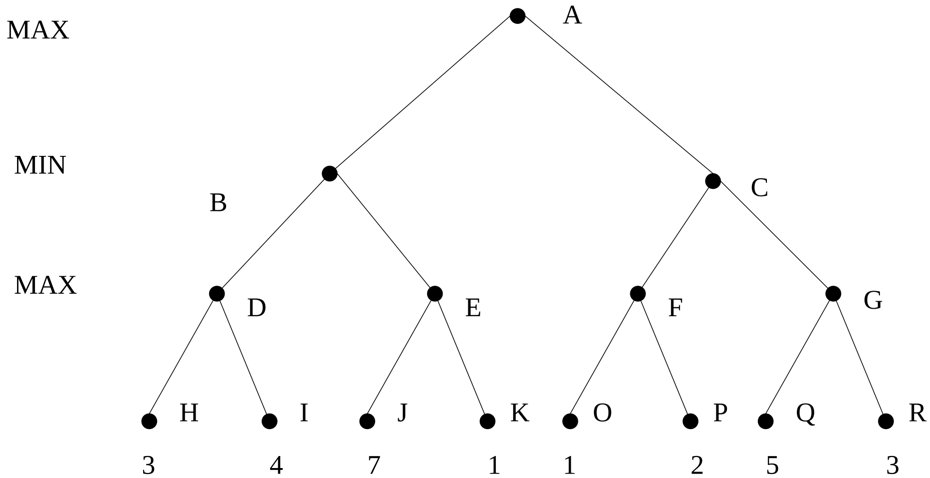

Assume the first player is the max player and the values at leaves of the tree reflect his/her utility. The opponent wants the same utility to be minimized.

Task a. (5 points). Compute the minimax values for each node in the tree? What move should the first player choose?

Task b. (10 points). Assume we use alpha-beta algorithm to explore the game tree and we do this in left-to-right order to determine the players strategies. Mark or list all nodes that are cut off from the tree and are never explored by the alpha beta procedure.

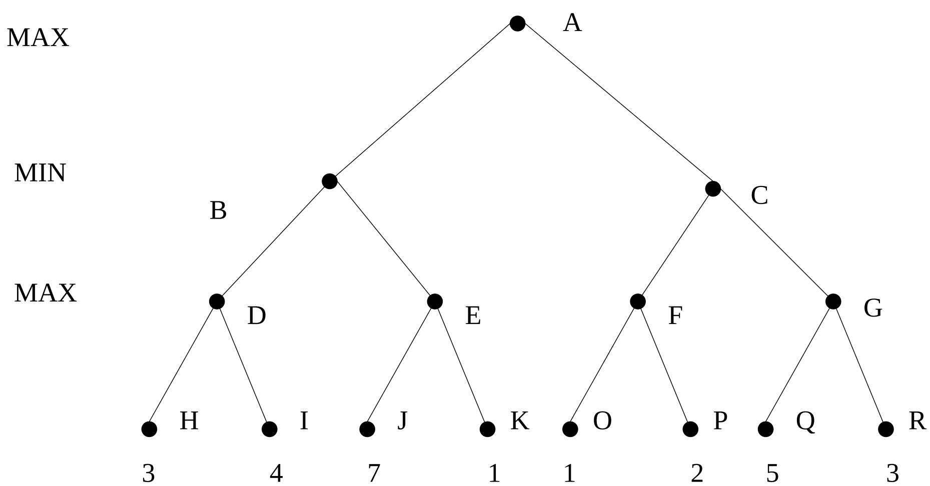

Problem 2 (10 points). Let the knowledge base (KB) consists of the following 2 sentences in the propositional logic:

$$P \Longrightarrow Q,$$

$$\neg P \Longrightarrow R.$$

Prove using resolution refutation that the KB entails: $(\neg Q \Longrightarrow R)$.

Problem 2 (20 points). Logical reasoning systems.

Part a. (6 points) Indicate whether the following statements are True or False:

1. Modus ponens is sound but not complete inference rule for the KB expressed in the first order logic. **T** **F**
2. Modus ponens is sound but not complete inference rule for the KB in the Horn normal form. **T** **F**
3. All KBs in the first order logic can be converted to the Horn normal form. **T** **F**

Part b. (7 points). Assume the KB that consists of the following rules:

- R1. $Soda(x) \wedge Chips(y) \Longrightarrow Cheaper(x, y)$
- R2. $Chips(x) \wedge Cereals(y) \Longrightarrow Cheaper(x, y)$
- R3. $Cheaper(x, y) \wedge Cheaper(y, z) \Longrightarrow Cheaper(x, z)$

and facts:

- F1. $Soda(Sprite)$
- F2. $Chips(CapeCod's)$
- F3. $Cereals(Cherios)$
- F4. $Cereals(MiniWheats)$

Consider that all facts F1-F4 are known at the beginning of the inference process. Illustrate the process of forward chaining by listing all newly inferred facts. Assume that both rules and facts are matched and tried in the order of their appearance.

Part c. (7 points) Show how to prove $Cheaper(Sprite, Cherios)$ using backward chaining and KB in part b. Draw the AND/OR search tree for the problem, assuming rules and facts are tried and matched in the order of their appearence.

Problem 4 (10 points). A lie detector test is known to be 60% reliable when the person is guilty of some crime (that is, $P(testguilty = T|guilty = T) = 0.6$) and 90% reliable when the person is innocent ($P(testguilty = F|guilty = F) = 0.9$). If a suspect is chosen from a group of suspects of whom only 10% has ever committed a crime, and the test indicates that he is guilty, what is the probability he is innocent?

Problem 5 (20 points). Assume the Bayesian belief network in the figure below. Assume that every variable in the network is binary and it can take two possible values T, F .

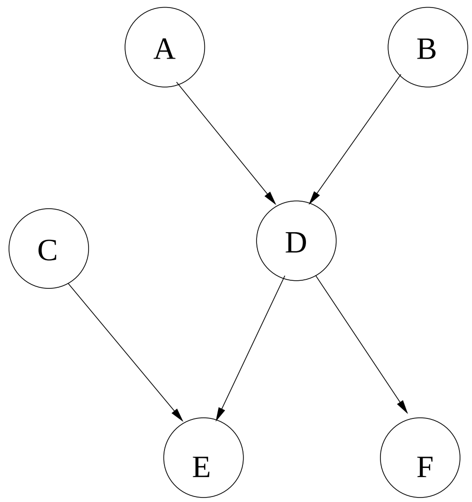

The belief network encodes the full joint distribution over random variables (represented by nodes) by exploiting conditional independences that hold among variables.

Part a. (2 points) What is the number of parameters needed to define the full joint distribution over variables in the problem domain without belief network representation? Remember all variables are binary.

Part b. (3 points) Show how to compute the joint probability P(A=T,B=T,C=F,D=F,E=T,F=F) using the belief network model.

Part c. (4 points) What is the number of parameters needed to define the belief network in the figure?

Part d. (5 points) Independences. Indicate whether the following statements are true or false.

1. B is independent of F given C. **T** **F**
2. A is independent of B given D. **T** **F**
3. F is independent of E given A. **T** **F**
4. B is independent of C given D. **T** **F**
5. A is independent of B given E. **T** **F**

Part e. (6 points) Outline briefly how would you compute conditional probability: $P(B = T|C = F, F = T)$ from the belief network and its parameters. You do not have to optimize the computation, just give the main idea.

Problem 6 (15 points). Assume you have to decide between 3 investment options A1, A2, A3. The monetary profits for different outcomes resulting from the three choices are summarized in the following decision tree.

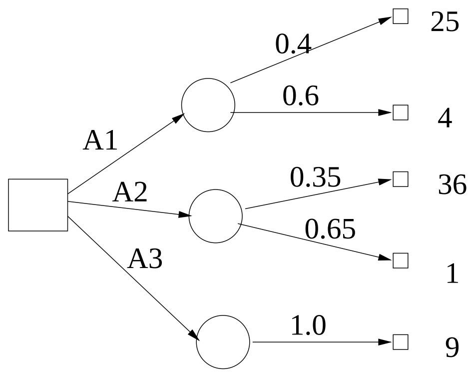

Part a. (7 points). Assume you want to maximize the expected monetary value of the investment. What action should you choose? Why?

Part b. (6 points). Assume your preferences towards different monetary outcomes are governed by the following utility function:

$$u(x) = \sqrt{x} + 2.$$

Compute the utilities of outcomes and add them to the tree below. What action should you choose under the expected utility criterion?

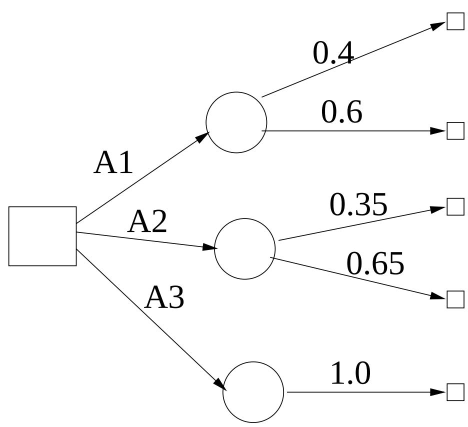

Problem 7 (10 points) Learning.

Part a. (3 points) What is overfitting?

Part b. (4 points) Describe how would you modify a linear regression model to find the degree m polynomial with the best (mean square error) fit to $< x, y >$ data pairs.

Part c. (3 points). Assume you have a coin. The probability of a head is represented by the parameter θ. Assume we observe the following sequence of coin tosses:
H, H, T, T, H, H, T, H, H, H.
Compute the maximum likelihood estimate of the parameter θ.