Aalto University, Department of Computer Science
ICS-E4000 Artificial Intelligence

Jussi Rintanen (TB329)
September 6, 2016

Problem 1 Explain briefly (1-4 sentences) (10 pts)

(a) Optimality of a search algorithm

(b) Explain what are *mixed equilibria* and in games in normal form (strategic form)? What kind of mixed equilibria can be obtained by iterated strict dominance?

(c) Discount factor γ in Markov decision processes

(d) Incoherent preferences in utility theory

(e) Reinforcement learning

Problem 2 Formulas that represent sets (10 pts)

As shown in the lecture, propositional formulas can be understood as representations of *sets* of bit-vectors. Let bit-vectors of length 4 be represented with atomic propositions a, b, c and d.

(a) Give a formula that represents the set $\{0000, 1111\}$.

(b) How many elements are there in the set represented by $a \wedge b \wedge c$? Explain.

(c) Consider two sets represented respectively by $(a \rightarrow b) \vee c$ and $a \vee (\neg c \rightarrow b)$. Give a formula that tests whether the sets are the same.

(d) Same, but give a formula that tests whether the first set is a subset of the second set.

Problem 3 Dynamic programming (10 pts)

(a) What is the Dynamic programming principle? (2 pts)

(b) Edit distance is a way of quantifying how dissimilar two strings (e.g., words) are to one another by counting the minimum number of operations required to transform one string into the other. Compute the edit distance between two strings "intention" and "execution". Each character insertion and deletion costs 1. Use dynamic programming and show all intermediate results. (8 pts)

Problem 4 Markov decision process (10 pts)

Consider an undiscounted ($\gamma = 1$) MDP having three states $s \in \{1,2,3\}$, with rewards $R(s,a,s') = R(s')$ = -1, -3, 0, respectively. State 3 is a terminal state. In states 1 and 2 there are two possible actions: a and b. The transition model $P(s,a,s')$ is as follows:

- In state $s = 1$, action a moves the agent to state $s' = 2$ with probability 0.7 and makes the agent stay put with probability 0.3.
- In state $s = 2$, action a moves the agent to state $s' = 1$ with probability 0.7 and makes the agent stay put with probability 0.3.
- Action b moves the agent to state $s' = 3$ with probability 0.1 and makes the agent stay put with probability 0.9.

Answer/do the following:

(a) Write the Bellman equation by expanding the following: $U(s) = \max_a EU(a \mid s) = \ldots$

(b) Apply 2 steps of either *value iteration* or *policy iteration* to determine a policy for states 1 and 2. State clearly whether you use value or policy iteration. If you use value iteration, assume that the initial utility function is all zero. If you use policy iteration, assume that the initial policy has action b in both states.

(c) If the policy does not change between two steps, can it be concluded that the algorithm has converged? (Please justify your answer with about one sentence.)

Problem 5 Decision theory (10 pts)

You are selling an apartment and are considering whether to use a real estate agent who takes 4 percent fee of the sales price. When you sell your apartment, you will either get a good price (100 k€) or a bad price (80 k€). The real estate agent increases the chance of getting a good price from 40% to 60%. Assume utility equals money.

(a) Draw the expectimax search tree that represents the problem.

(b) Calculate the expected utility (money) for each node in the tree.

(c) Should you hire the real estate agent?

The name of the course, the course code, the date, your name, your student number, and your signature must appear on every sheet of your answers.

Please note the following: your exam will be graded only if you have completed at least three of the obligatory home assignments before the exam.