

COS 402: Artificial Intelligence

Sample Final Exam

Fall 2008

Print your name _____

General directions: This exam is closed book. However, you may use a one-page “cheat sheet” as explained in the instructions posted prior to the exam. You also may use a calculator (but only the standard math functions, not any programming or text capabilities). You may not use the text book, your notes, a computer, or any other materials during the exam. Cell phones and all other devices that can be used for any form of communication must be turned off throughout the entire exam.

No credit will be given for questions left unanswered, so you should be sure to answer all questions, even if you are only taking your best guess.

Write your answer to each question or problem in the space provided. If necessary, you may also attach extra sheets. Make sure your name is printed on all of these.

Please write neatly and answer all questions unambiguously since we cannot give credit to answers that we can't read or decipher.

This exam has a total of _____ points on _____ pages (including this cover page). Before beginning, please double-check that you have received all of the pages. You have 180 minutes. Be sure to pace yourself carefully, and leave yourself ample time for the longer problems at the end of the exam.

Good luck!

Honor pledge: “I pledge my honor that I have not violated the Honor Code during this examination.”

Please write the pledge in full and sign:

Short answer A

Directions: This section contains modified true-false questions. Each of these is a statement, part of which has been underlined. Your response to the statement should be either “*Correct*”, meaning that the statement is correct as is, or you should cross out part or all of what is underlined and write in the correct word or phrase. For instance, the statement might be:

A0. The capital of New Jersey is Albany.

Here, the correct response is to cross out the word “Albany” and to write the word “Trenton” in its place.

Note that if the statement is already correct, you should not modify it, even if there is a “better” answer. Also, any question left entirely blank will receive no credit, so be sure to write “*Correct*” if no changes are necessary.


A1. The iterative deepening search algorithm calls depth-limited search repeatedly with the depth limit ℓ set to 1, 2, 4, 8, 16, . . .

A2. The graph-search version of A^* will be optimal if an admissible heuristic function is used.

A3. By definition, α equals β , written $\alpha \equiv \beta$, means that for every model m , if α is true in m , then β is also true in m .

A4. An inference algorithm that never concludes $KB \models \alpha$, unless it is so, is said to be complete.

A5. In a Bayesian model like this one



all variables are assumed to be conditionally independent of one another given a “hidden cause” variable. Such a model is called a Markov chain model.

- A6. Let a and b be events. Then $P(a|b)$ can be expressed in terms of $P(b|a)$ using Markov's rule.
- A7. Human speech is composed of 40-50 sounds called signs.
- A8. Given evidence up to time t , the forward-backward algorithm runs in time quadratic in t .
- A9. The utility U^* of the optimal policy π^* must satisfy a set of equations called the Markov conditions.
- A10. In an MDP, the larger the discount factor, the more strongly favored are short-term rewards over long-term rewards.
- A11. Typical decision tree algorithms grow a tree by greedily minimizing an impurity function.
- A12. SVM's sidestep the statistical difficulty of working in very high dimensional spaces using the kernel trick.

Short answer B

Directions: This section contains modified multiple choice questions. Unless otherwise indicated, these questions may have *several* correct answers. You should *circle all that apply*. For instance:

B0. The colors of the American flag include which of the following?

- (a) red
- (b) pink
- (c) blue
- (d) purple

Obviously, the right answer is to circle both (a) and (c). If none of the answers are correct, write the word “None” next to the choices. Questions left entirely blank will receive no credit.

B1. Which of the following are used in typical chess programs such as Deep Blue?


- (a) alpha-beta pruning
- (b) evaluation functions
- (c) forward chaining
- (d) genetic algorithms
- (e) MCMC
- (f) transposition tables

B2. Let P and Q be proposition symbols.

Which of the following are models of $\neg P \vee Q \Rightarrow \neg P \wedge Q$?

- (a) $P = \text{false}, Q = \text{false}$
- (b) $P = \text{false}, Q = \text{true}$
- (c) $P = \text{true}, Q = \text{false}$
- (d) $P = \text{true}, Q = \text{true}$


B3. Consider the following game tree:



In the figure, \triangle and ∇ indicate that it is the MAX and MIN player's turn, respectively, and the numbers at the bottom show the payoff for the given terminal positions. Which of the following nodes would *not* be generated or explored if alpha-beta pruning were applied? Assume nodes are generated in the *optimal order* for alpha-beta pruning.

- (a) A
- (b) B
- (c) C
- (d) D
- (e) E
- (f) F

B4. Consider the following Bayes net:



Which of the following *must* be true? [Note: A complete answer to this question requires more information than was covered in class, so don't worry if you cannot fully answer it. For the actual final, you only need to know the independencies discussed in class and on page 499 of the R&N textbook.]

- (a) B and F are conditionally independent given A, C, D and E .
 - (b) B and F are conditionally independent given A, C and E .
 - (c) B and F are conditionally independent given A and E .
 - (d) B and F are conditionally independent given C and D .
 - (e) B and F are conditionally independent given C and E .
- B5. Let a and b be any two events. Which of the following *must* be true?
- (a) $P(a) \geq 0$
 - (b) $P(a \wedge b) = P(a)P(a|b)$
 - (c) $P(a \vee b) = P(a) + P(b)$
 - (d) $P(\neg a) + P(a) = 1$
- B6. Which of the following are classification tasks appropriate for classification learning algorithms?
- (a) predicting if a credit card transaction is fraudulent or legitimate
 - (b) predicting how much it will rain tomorrow
 - (c) predicting the letter of the alphabet represented by an image of a handwritten character
 - (d) breaking a database of customers into clusters based on their buying patterns (where the nature of the clusters is determined automatically by the computer, not in any way provided by a human)

Short answer C

Directions: In each of the following questions, fill in the blank next to each item using the code provided. For instance:

C0. Which of the following colors appear on the American flag and which on the Canadian flag?

- A = American but not Canadian
- C = Canadian but not American
- B = both
- N = neither

- (a) _____ red
(b) _____ purple
(c) _____ blue

Here, the right answer is (obviously) to mark (a) with the letter B , (b) with the letter N and (c) with the letter A .

C1. Which of the following search algorithms are complete and which are optimal?

- C = complete but not optimal
- O = optimal but not complete
- B = both complete and optimal
- N = neither complete nor optimal

- (a) _____ breadth-first search
(b) _____ depth-first search
(c) _____ depth-limited search
(d) _____ iterative deepening search

C2. Let P and Q be proposition symbols. Which of the following are literals and which are sentences?

- L = is a literal but not a sentence
- S = is a sentence but not a literal
- B = is both a literal and a sentence
- N = is neither a literal nor a sentence

- (a) _____ P
(b) _____ $\neg Q$
(c) _____ $P \wedge Q$
(d) _____ $P \Rightarrow \neg Q$
(e) _____ $P \models Q$

C3. Which of the following are true of MDP's and which are true of HMM's?

- H = true of HMM's only
- M = true of MDP's only
- B = true of both
- N = true of neither

- (a) _____ at the foundation of modern speech recognition systems
- (b) _____ at the foundation of the world's best backgammon program
- (c) _____ are useful for a robot navigating through its environment
- (d) _____ the agent has control of its actions
- (e) _____ the underlying state is unseen

C4. Mark each of the following applications with the *single* family of tools that would be *most* appropriate. Use the following code:

- B = Bayesian networks
- D = MDP's (including reinforcement learning)
- K = logical reasoning and knowledge representation
- L = machine learning (not including reinforcement learning)
- S = search algorithms
- T = temporal models (such as HMM's)

- (a) _____ classifying news articles by topic
- (b) _____ controlling a bank of elevators
- (c) _____ logistics planning for airlines and airports
- (d) _____ medical diagnosis using substantial expert knowledge
- (e) _____ optical character recognition
- (f) _____ robot localization
- (g) _____ software verification
- (h) _____ theorem proving
- (i) _____ tracking enemy targets

Problems

Directions: To obtain full credit, be sure to *show all your work, and justify your answers*. If necessary, you may attach extra sheets. Make sure your name is printed on all of these.

P1. [Propositional logic]

Consider three blocks A , B and C which can be stacked on top of each other, or on a table, for instance, in configurations like the following:



In this problem, we do not move the blocks around, but rather consider a *fixed* configuration of the blocks. In such a configuration, let $S_{A,B}$ be a proposition symbol representing the fact that A is stacked above B , and similarly define $S_{B,A}$, $S_{A,C}$, $S_{C,A}$, $S_{B,C}$ and $S_{C,B}$. For instance, in the figure on the left, $S_{C,A}$ would be the only one of these symbols that is true. In the figure on the right, the only true symbols are $S_{A,B}$, $S_{B,C}$ and $S_{A,C}$.

For parts (a-d) below, be sure to *explicitly* list *all* of the required clauses or sentences. Do not just give examples of a general pattern.

- Show how to represent using propositional logic the fact that if one block is stacked above another, then the second cannot be above the first (for instance, if A is above B , then B cannot be above A).
- Show how to represent the fact that if one block is stacked above a second block, which is in turn stacked above a third block, then the first block must be stacked above the third block (for instance, if A is above B , and B is above C , then A must be above C).
- Show how to represent the proposition that at least one block must have no blocks stacked above it.

(d) Let KB denote the conjunction of sentences in parts (a) and (b), and let α denote the sentence (or conjunction of sentences) in part (c). Convert $KB \wedge \neg\alpha$ to CNF and list all of its clauses.

(e) Choose three pairs of clauses from part (d) that are eligible for resolution and show the result of resolving them. (Be sure to indicate which pairs you are resolving.) Do not pick pairs that yield tautologies.

P2. [Markov decision processes]

Consider an MDP defined on a two-bit register with two actions: f , which flips the leftmost bit 90% of the time, but has no effect 10% of the time; and r , which swaps the two bits succeeding 100% of the time. (For instance, f applied to register contents 01 will, 90% of the time, change the register contents to 11; r applied to 01 will change the register contents to 10.) Denote the states by the register contents: 00, 01, 10 and 11. Let the discount factor γ be equal to 0.8, and let the reward at every state be equal to the number of ones in the register.

- (a) Write down a set of equations that the utility U^* of the optimal policy must satisfy. These equations should be in terms of $U^*(00)$, $U^*(01)$, $U^*(10)$ and $U^*(11)$ only.

- (b) Starting with $U_0(s) = 0$ for all states s , show how value iteration would compute U_1 and U_2 (the next two approximations of U^*) in this case.

P3. [Machine learning]

We are given the following labeled dataset:

x_1	+
x_2	-
x_3	+
x_4	+
x_5	-

On the first round of AdaBoost, the following weak hypothesis h_1 is found:

$h_1(x_1)$	=	+
$h_1(x_2)$	=	+
$h_1(x_3)$	=	-
$h_1(x_4)$	=	+
$h_1(x_5)$	=	-

(a) What is the training error of h_1 ?

(b) Compute the next distribution D_2 on the examples for round 2. (This distribution, or set of weights on the training examples, was denoted \mathbf{w} in the book.)

(c) Suppose now that the following weak hypothesis h_2 is found:

$$\begin{aligned}h_2(x_1) &= - \\h_2(x_2) &= - \\h_2(x_3) &= + \\h_2(x_4) &= + \\h_2(x_5) &= +\end{aligned}$$

What is the weighted training error of h_2 on D_2 ?

(d) Suppose we stop boosting after just these two rounds. Show how the final combined classifier produced by AdaBoost will predict on each of the five examples.