

EECS 492, Winter 2010
Introduction to Artificial Intelligence
Final Exam (105 points)

Name: _____

Uniqname: _____

This exam is closed book: you may not use the textbook, lecture slides, notes, or any type of computing device. Write your name at the top of every page.

“I have neither given nor received aid on this examination,
nor have I concealed any violations of the honor code.”

signed _____

1	/ 20
2	/ 15
3	/ 15
4	/ 15
5	/ 10
6	/ 20
7	/ 10
total	/ 105

Problem 1. Quickies

1.1. (4 points) Uninformed search methods (such as DFS and BFS) *can* be used to solve which of the following problems? (Circle all that apply.)

- A) Map coloring
- B) Path planning
- C) 7 Queens
- D) 3SAT
- E) Linear regression

1.2. (4 points) Which are advantages of depth first search over breadth first search? (Circle all that apply.)

- A) Explores more promising subtrees before less promising ones
- B) Exploits problem structure
- C) Lower computational complexity
- D) Lower memory complexity

1.3. (4 points) Which of the statements below are true of Iterative Deepening Search (IDS)? (Circle all that apply.)

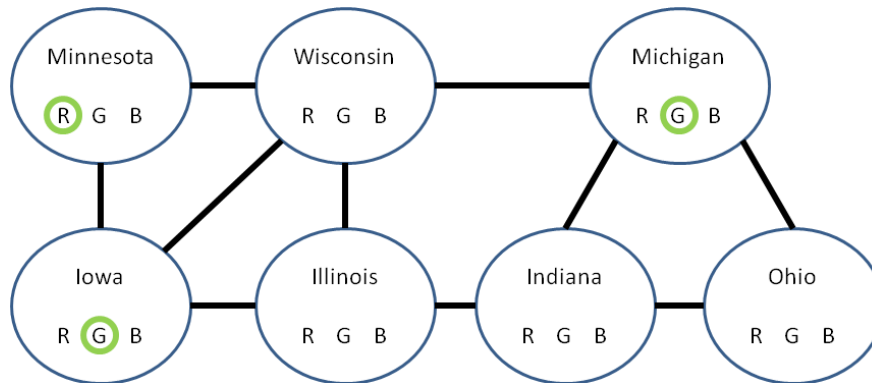
- A) Is an informed search algorithm.
- B) Has the same memory complexity as BFS
- C) Has an exponential computational complexity
- D) Is *both* complete and optimal.

1.4. (4 points) A secretive researcher invents a box that somehow manages to invert matrices in $O(1)$ time. Which of the following problems/algorithms would benefit from this? (Circle all that apply).

- A) kNN Regression
- B) kNN Classification
- C) Value Iteration
- D) Policy Iteration
- E) Extended Kalman Filter: Propagation step
- F) Extended Kalman Filter: Observation step

1.5. (4 points) The DPLL algorithm, which you implemented for the Wumpus programming challenge, is (circle all that apply):

- A) Breadth-first search
- B) Satisfiability checker
- C) Resolution-based inference algorithm
- D) Capable of solving 3-SAT, and therefore an NP-complete solver.

Problem 2. Constraint satisfaction

2.1. (3 points) We wish to find a 3-coloring of the graph above. Some colors have already been assigned. Apply the *Arc Consistency* algorithm to the graph above, crossing out any eliminated values.

(show your work in the graph above)

2.2. (6 points) We now wish to assign a value to another node.

A) Which node would we expand next if we used the minimum remaining values (MRV) heuristic? (Use the the Arc Consistency results from above.)

B) Which node would we expand next if we used the degree heuristic?

2.3. (2 points) Let us assume, as shown in the figure above, that values have been assigned for Minnesota, Iowa, and Michigan. Which of the following is a tight asymptotic bound on the computational complexity of searching the remaining $N=4$ values using forward checking?

- A) $O(N)$
- B) $O(N^3)$
- C) $O(3^N)$
- D) $O(N!)$

2.4. (4 points). Suppose we wish to solve this constraint satisfaction problem from scratch (i.e., without the three values in the figure above). We want to use a hill-climbing local search method. Which of the below would be an appropriate objective function? (Assume that we will either minimize or maximize as appropriate).

- A) The number of color assignments that are different from the solution.
- B) The number of nodes colored "red".
- C) Number of constraints that are satisfied.
- D) The number of disjoint sub-graphs induced in the original graph.

Problem 3. Logic

Consider the KB below:

$\text{Wolf}(x) \Rightarrow \text{Canine}(x)$
 $\text{Wolf}(x) \Leftrightarrow \text{HasSharpTeeth}(x)$
 $\neg \text{Wolf}(\text{GrandMother}) \wedge \neg \text{Wolf}(\text{LittleRedRidingHood})$

3.1. (2 points) The KB above is written in (check one):

Propositional Logic

First Order Logic

3.2. (4 point) "GrandMother" is what: (Circle all that apply)

- A) A function of zero arguments
- B) A proposition
- C) A predicate of zero arguments
- D) An object
- E) An existentially quantified object.
- F) A horn clause

3.3. (6 points) Which of the following statements are entailed by the KB above?

- A) $\text{Wolf}(\text{RogerTheWolf})$
- B) $\neg \text{Wolf}(\text{LittleRedRidingHood})$
- C) $\text{Wolf}(x) \vee \neg \text{HasSharpTeeth}(x)$
- D) $\neg \text{HasSharpTeeth}(\text{GrandMother})$
- E) $\neg \text{Canine}(x) \Rightarrow \neg \text{Wolf}(x)$
- F) $\neg \text{HasSharpTeeth}(\text{LittleRedHood}) \vee \text{Wolf}(\text{LittleRedRidingHood})$

3.4. (3 points) What substitutions, if any, unify the pairs of sentences below? You may standardize-apart as useful.

- a) Sentence 1: $A(x) \wedge B(\text{Bob})$
Sentence 2: $A(y) \wedge B(\text{Roger})$
- b) Sentence 1: $A(x) \wedge B(x,y) \wedge C(\text{Bob}, y)$
Sentence 2: $A(\text{Bob}) \wedge B(x, \text{Roger}) \wedge C(z, \text{Roger})$
- c) Sentence 1: $A(\text{Bob}) \wedge \neg B(x)$
Sentence 2: $A(x) \wedge \neg B(\text{Charlie})$

Problem 4. Probability

4.1. (2 points) In terms of $P(A|B)$, and the marginal distributions for A and B, what is $P(B|A)$?

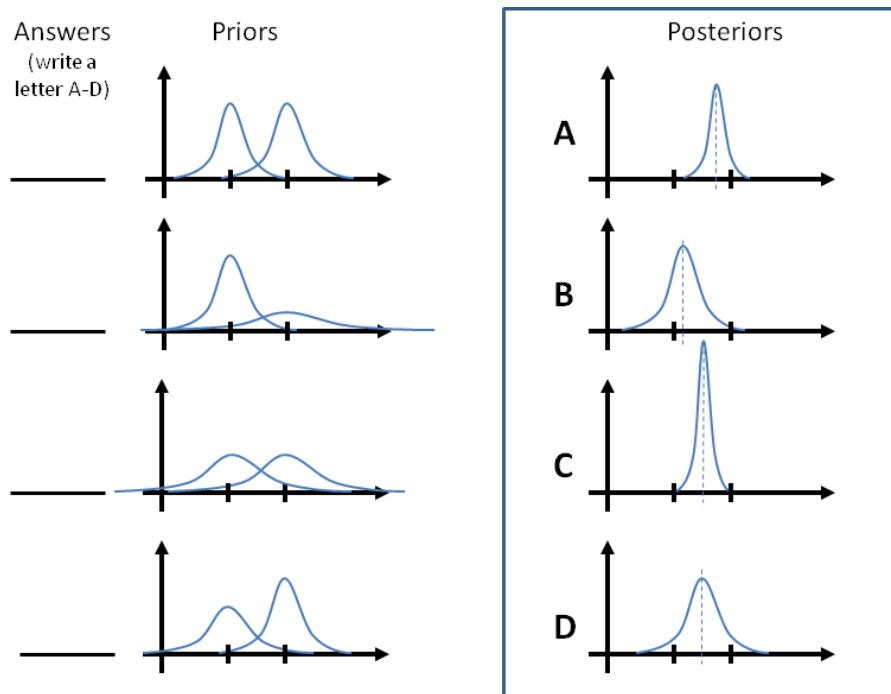
4.2. (6 points) The random variable x is normally distributed with mean 2 and covariance 3. The random variable w is normally distributed with mean 0 and covariance 1. What is the distribution for random variable y , which is given as $y = 5x - 2w$?

Mean of y :

Covariance of y :


4.3. (3 points) A fair coin is flipped four times. What is the probability that at least one "heads" occurs? (Write your answer as a simple expression.)

4.4 (4 points) Suppose that a robot obtains two observations of the position of a one-dimensional landmark. The distributions of these observations are given on the left. Select the appropriate posterior distribution on the right. Note that each answer on the right is used exactly once.



Problem 5. Bayes Net

Consider the Bayes net below.



5.1. (2 points) What is $P(A, B, -C, -D)$? (You can write your answer as a product.)

The 12 samples below were obtained by random sampling from the distribution above.

- | | | | | | | | | | | | |
|---|----|----|----|----|----|----|----|----|----|----|----|
| A | B | -C | -D | A | -B | C | D | -A | -B | C | D |
| A | -B | -C | D | -A | -B | C | -D | A | -B | -C | D |
| A | -B | -C | -D | A | -B | C | -D | -A | B | C | D |
| A | B | -C | D | A | -B | -C | -D | A | B | -C | -D |

5.2. (2 points) Using these samples, estimate the distribution $P(A, B, -C, -D)$. Write your answer as a non-reduced fraction (e.g., 46/112).

5.3. (2 points) Using rejection sampling, estimate the distribution $P(-A | D)$. Write your answer as a non-reduced fraction.

5.4. (4 points) Suppose we wish to compute $P(A | D)$ using likelihood sampling. Compute the weights associated with the samples below. Write your answer as a simplified decimal (e.g., 0.58).

A -B -C D weight: _____

-A B C D weight: _____

Problem 6. Decision Making, Learning, and Evaluation

6.1 (5 points) When performing value iteration, the utilities of every state will (non-strictly) monotonically increase when the immediate rewards of each state satisfy what property?

6.2 (5 points) What factors motivate the discount factor, gamma, being less than 1?

- A) A dollar today is worth more than a dollar tomorrow.
- B) Long plans are more likely to go wrong than short plans.
- C) Allowing for distant future rewards requires more memory.
- D) Algorithms for computing policies converge more quickly.
- E) The value of rewards is stationary: it does not matter when they are received.

6.3. (6 points) Which of the statements below are true? (Circle all that apply).

- A) For any combinational circuit, there is a network of perceptrons that computes the same function.
- B) A kNN query can be accelerated by using a linked list data structure.
- C) Adaboost is an example of a genetic algorithm
- D) Perceptrons and SVMs are similar in that they both compute hyperplane separators.
- E) The primary advantage of kNN classifiers versus neural networks is their ability to automatically determine the correct scaling of the input features.
- F) In a decision tree, it is sometimes advantageous to *not* split on an attribute, even when that attribute perfectly classifies the training data.

6.4 (4 points) Rank the challenge problems from your favorite (#1) through your least favorite (#6).

_____ PS1:Borealis

_____ PS4:SimElevator

_____ PS2:Decryption

_____ PS5:Image Denoising

_____ PS3:Wumpus

_____ PS6:Netflix

Optional Comments (How could we make them better? What was good about the ones you liked? Did they help you gain a deeper understanding of the material?)

Problem 7. Matching

(10 points, -1 for each wrong.) Match the terms (letters on left) with the most appropriate phrase on the right:

- | | |
|--------------------------------|--|
| _____ A) Cross-Validation | 1. When multiplied together, the joint of an MRF. |
| _____ B) Moralizing | 2. Principle component of EKF Propagation step |
| _____ C) Dominant Eigenvector | 3. A desirable property of a leaf node in adversarial search |
| _____ D) Crossover | 4. All true statements can be proved in a finite number of steps. |
| _____ E) Valid | 5. Estimate of "Cost to go" as used by A* |
| _____ F) Potential Function | 6. With mutation, mechanism for genetic reproduction. |
| _____ G) Covariance Projection | 7. Probably Approximately Correct |
| _____ H) Quiescence | 8. Convert a Bayes net into a MRF |
| _____ I) Herbrand's Theorem | 9. True in every model |
| _____ J) Independence | 10. Joint distribution can be written as product of marginal distributions |
| _____ K) Modus Ponens | 11. If A is true, and A implies B, then B. |
| _____ L) Heuristic Function | 12. Direction of major axis of covariance ellipse |
| | 13. Dividing data into multiple overlapping training/testing datasets |

Thanks everyone for a great term. Don't forget to submit your online course evaluation.