

CS 430: Introduction to Artificial Intelligence
Fall 2006
Final Exam

Name: _____

This final exam consists of 16 pages and a total of 110 marks. The exam is open book and open notes. You have 110 minutes to complete this exam. If you are taking too much time on one question, please move on to other questions.

Section	Marks
Pre-midterm questions	/ 20
Propositional Logic	/ 20
Bayesian Networks (Parameter Learning)	/ 10
Bayesian Networks (Structure Creation)	/ 10
Bayesian Networks (Inference)	/ 20
MDPs and Reinforcement Learning	/ 30
Total	/ 110

Section I: Pre-midterm questions [20 points]

1. Answer the following questions using true or false. [1 point each]

a) Recall the different characteristics of environments (ie. fully vs partially observable, deterministic vs stochastic, episodic vs sequential, static vs dynamic, discrete vs continuous, and single vs multi-agent). The most difficult type of environment for an agent is a fully observable, deterministic, episodic, static, discrete and single-agent environment.

b) In a deterministic environment, if you are in state S and you perform action A , the probability of reaching the next state S' is 1.

c) Rationality depends on the following 4 things: a performance measure, the agent's prior knowledge of the environment, the actions the agent can perform, and the agent's current percept.

d) A simple-reflex agent will work well in a partially observable environment.

e) Local search techniques such as hillclimbing, simulated annealing, and genetic algorithms not only find the optimal state but also return the path to that state.

f) A program implementing breadth-first search can be easily converted to depth-first search just by changing the fringe implementation from a FIFO queue to a LIFO queue.

g) Suppose you have a heuristic function $h^*(n)$ which tells you the actual cost of the cheapest path from the node n to the goal. Then $h^*(n)$ dominates all other admissible heuristics for the search problem.

- h) A-star search has a space complexity that is linear in the number of states.

- i) Depth first search is optimal.

- j) Both simulated annealing and genetic algorithms are a pain to implement because they require a lot of parameter-tweaking in order to get them to work on a particular problem.

- k) Evaluation functions are only useful for minimax or alpha-beta if we are unable to keep the entire game tree in memory.

- l) The entire game tree for chess can be stored in memory for a standard desktop machine with 4 GB of RAM.

- m) It is perfectly rational to play a dominant strategy.

- n) If Player 1's payoff in a zero-sum, two-player game is +100, then the payoff for Player 2 is -100.

- o) If a normal-form game has a single Nash equilibrium and you tell your opponent that you are playing your part of the Nash Equilibrium strategy, then your opponent (assuming he/she is rational) has no choice but to play his/her part of the Nash Equilibrium strategy.

p) In the Tragedy of the Commons, the farmers would make more money if they acted in their own self-interest and did not cooperate with each other when deciding on the number of goats to graze.

2. Circle the pure strategy Nash equilibrium in the normal-form game below or write down “None” if one doesn’t exist. [2 points]

	Player B: S1	Player B: S2
Player A: S1	A = 1, B = 3	A = -2, B = 2
Player A: S2	A = -2, B = 2	A = 3, B = -1

3. Write out the gradient descent rule for the function $f(x) = 4x^5 + 2x^3 + 2$, assuming the learning rate $\alpha = 0.1$. [2 points]

Section II: Propositional Logic [20 points]

1. You would like to order a pizza for an extremely picky friend. Fortunately, you've stored all of her preferences in a knowledge base and you can use that knowledge base to infer if she will be happy with the pizza you just ordered for her.

If your pizza has pepperoni, your friend will be happy. If the pizza has anchovies and feta cheese, your friend will get sick. If your friend gets sick, she will not be happy. If you order olives, the pizzeria will automatically put olives and feta cheese on the pizza. You order a pizza with anchovies and olives.

a) The paragraph above has been converted into propositional logic symbols below. However, the knowledge base below is not in CNF. Convert it to CNF. [10 points]

$$P \Rightarrow H$$

$$(A \wedge FC) \Rightarrow S$$

$$S \Rightarrow \neg H$$

$$O \Rightarrow FC \wedge O$$

$$A \wedge O$$

b) Based on the knowledge base above, does it entail that your friend will **NOT** be happy with the pizza you ordered? Show all your resolution steps or you won't get full credit. [5 points]

3. If $\alpha \models \beta$ and $\beta \models \gamma$, does $\alpha \models \gamma$? Explain your answer. [5 points]

Section III: Bayesian Networks (Naïve Bayes) [10 points]

In this question, you will be answering questions regarding implementing a naïve Bayes classifier that predicts a person's favorite team based on their favorite color and home town. The training data set is shown below. The FavoriteColor feature only takes two possible values of Green and Orange. The HomeTown feature only takes two possible values of Corvallis and Eugene. Finally, the FavoriteTeam class label only takes two possible values of Ducks and Beavers.

FavoriteColor	HomeTown	FavoriteTeam
Green	Eugene	Ducks
Orange	Corvallis	Beavers
Orange	Corvallis	Beavers
Green	Corvallis	Ducks
Orange	Eugene	Beavers
Green	Eugene	Beavers
Green	Corvallis	Ducks
Green	Eugene	Ducks
Orange	Corvallis	Beavers
Orange	Eugene	Beavers

1. Write down all the probabilities you need to implement a naïve Bayes classifier along with their values based on this training data set. Use the Dirichlet priors we mentioned in class. [8 points]

2. Some of the probabilities above can be derived from the other probabilities. What is the minimum number of the probabilities above that you must store in order to implement a naïve Bayes classifier? [2 points]

Section IV: Bayesian Networks (Structure creation) [10 points]


1. We would like to model the event that Bob is late for class. If we know that Bob is late because he slept in, then this knowledge would “explain away” the event that Bob was abducted by aliens. Draw the Bayesian network that represents this information using the three nodes of Late, AlienAbduction, and SleptIn. [3 points]

2. You have been asked to build a Bayesian network for a widget manufacturing process. The quality of the widget depends on the skill of the worker and the quality of the machine. The worker's skill is in turn influenced by the training the worker has received and the aptitude of the worker. Knowing the worker's skill makes the worker's training and aptitude independent of the widget quality. Two factors affect the quality of the machine: its age and the frequency of servicing. Knowing the quality of the machine makes the age and the frequency of servicing independent of the quality of the widget.

Draw a Bayesian network which captures this manufacturing process. You will use the nodes WidgetQuality, WorkerSkill, WorkerTraining, WorkerAptitude, MachineQuality, MachineAge, MachineServicing. [7 points]

Section V: Bayesian Networks (Inference) [20 points]


1. Calculate $P(A = \text{True} \mid B = \text{True})$ using the Bayesian network structure and conditional probability tables below. [8 points]



A	P(A)
False	0.2
True	0.8

A	B	P(B A)
False	False	0.5
False	True	0.5
True	False	0.7
True	True	0.3

2. Use the Bayesian network structure and conditional probability tables below to calculate $P(C = \text{false})$. [8 points]




A	P(A)
False	0.6
True	0.4

B	P(B)
False	0.2
True	0.8

A	B	C	P(C A,B)
False	False	False	0.75
False	False	True	0.25
False	True	False	0.75
False	True	True	0.25
True	False	False	0.75
True	False	True	0.25
True	True	False	0.75
True	True	True	0.25

C	D	P(D C)
False	False	0.75
False	True	0.25
True	False	0.9
True	True	0.1

3. Answer true or false to the conditional independence statements using the graph below. Write out the blocked or non-blocked paths for partial credit.

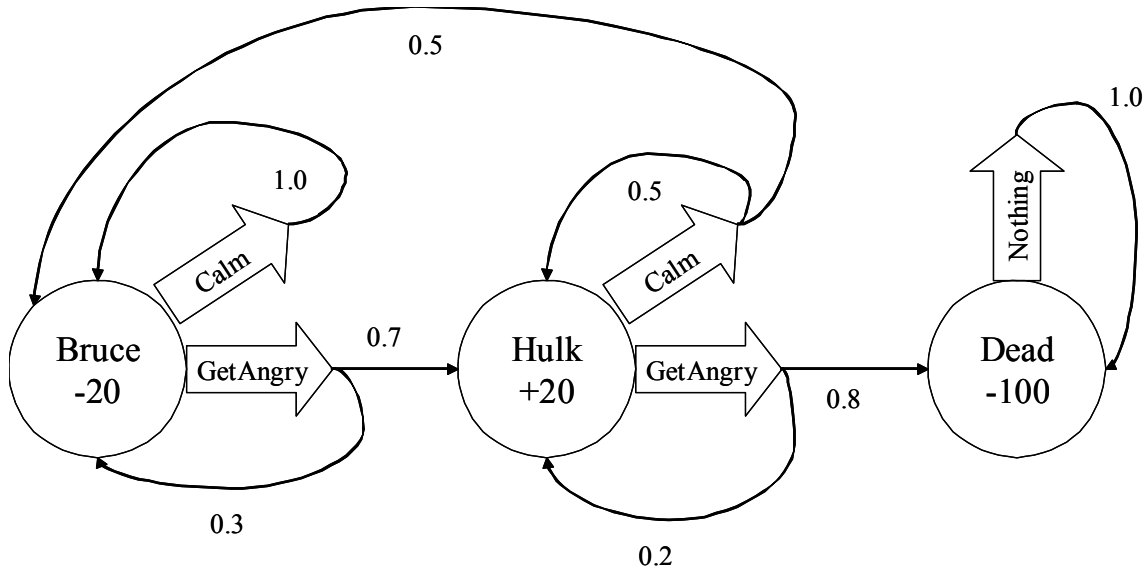


a) $I(B, D | A)$ [2 points]

b) $I(H, I | E)$ [2 points]

Section VI: MDPs and Reinforcement Learning [30 points]

1. The Markov Decision Process below represents a model of how Bruce Banner becomes the Incredible Hulk.



a) If we perform value iteration on this MDP, what are the initial values of the utilities of each state ie. what are $U_1(\text{Bruce})$, $U_1(\text{Hulk})$, and $U_1(\text{Dead})$? [2 points]

b) If we perform one round of value iteration on this MDP, what are the values for the utilities of each state ie. what are $U_2(\text{Bruce})$, $U_2(\text{Hulk})$, and $U_2(\text{Dead})$? Use $\gamma = 0.9$. [6 points]

c) Now suppose we do policy iteration using an initial policy of $\pi_1(\text{Bruce}) = \text{Calm}$, $\pi_1(\text{Hulk}) = \text{Calm}$ and $\pi_1(\text{Dead}) = \text{Nothing}$. Do one round of the Policy Evaluation phase and calculate the utilities of each state. Use $\gamma = 0.9$. [8 points]

2. Explain the difference between the reward for a state and the utility of a state. [2 points]

3. In a reinforcement learning algorithm with an ϵ -greedy (ie. epsilon greedy) exploration scheme, what happens when we make $\epsilon = 0$? What happens when we make $\epsilon = 1$? [2 points]

4. The grid world below is the setting for a predator-prey game. At the beginning, the predator is placed at the location marked “Pred” while the prey is placed at the location marked “Prey”. The goal of the predator is to catch the prey by moving into the same square as the prey. The predator can move but the prey remains stationary. The black squares correspond to pits. If the predator falls into a pit, then it will suffer a large negative reward.

We will make our agent the predator and we will learn the optimal policy for the predator using reinforcement learning. In order to formulate the predator-prey game as a reinforcement learning problem, we will first describe the action space. The agent can perform actions of N, E, S, and W in this world. However, these actions are stochastic, meaning with some probability, the agent will succeed and with some probability it will move in a random unintended direction.

								Pred
	Prey							

a) Define the state space in this problem. How many possible states are there? (To help you, there are 12 black squares). [2 points]

b) Define a reward function for this game, keeping in mind that the predator should capture the prey as quickly as possible. Remember that the pits are undesirable for the predator, so it should avoid them. [2 points]

c) The state with the highest q-value will be the state where the predator catches the prey. What state-action pairs should have the next highest q-values? You can describe this with sentences. [2 points]

d) Without changing the transition model or the reward function, how else can we “encourage” the predator to take the risky path between the pits to catch the prey? [2 points]

e) Suppose the prey could move and the game ends when the predator occupies the same square as the prey. How would you formulate the state space now and how many possible states are there? [2 points]

The End