

Second Midterm

CSC 242

May 2008

Write your **NAME** legibly on the bluebook. Work all problems. You may use two double-sided pages of notes. Please hand your notes in with your bluebook. The best strategy is not to spend more than the indicated time on any question (minutes = points).

1. Bayes, Entropy, Mutual Information (15 + 10 Mins)

A) Propositions a and b are (absolutely) independent. Write down three mathematical statements of this independence and prove them all equivalent.

B) The Entropy of a random variable is defined

$$H(x) = - \sum_x P(x) \log P(x)$$

For binomial trials like coin tosses, there are two terms with probabilities q (heads, say) and $1 - q$ (tails). An example would be the entropy of a fair coin toss:

$$H(x) = -\left[\frac{1}{2} \log\left(\frac{1}{2}\right) + \frac{1}{2} \log\left(\frac{1}{2}\right)\right] = -\left[\frac{1}{2}(-1) + \frac{1}{2}(-1)\right] = 1.$$

Entropy is a measure of “surprise”, and if $q = 1/8$, for example, the outcome is mostly tails and so is more predictable; in fact in this case $H(x) = 0.5346$.

Show that if you represent q by a fraction $1/n$, the entropy of the binomial trial random variable is

$$-\left[\frac{n-1}{n} \log(n-1) - \log n\right].$$

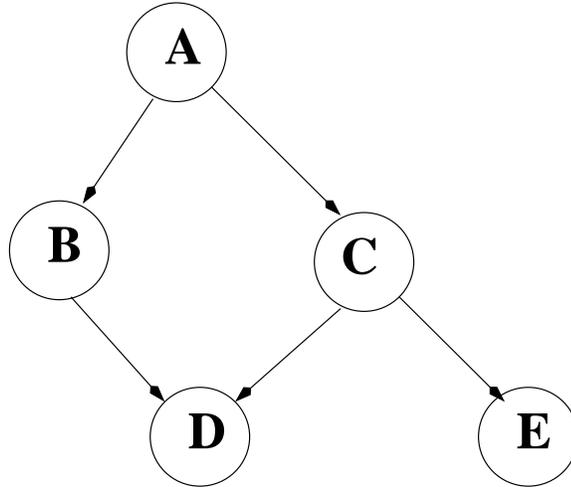
C) (extra credit 10 mins). The *mutual information* between two random variables is defined as

$$I(x, y) = \sum_{xy} p(x, y) \log \left[\frac{p(x, y)}{p(x)p(y)} \right].$$

Prove that $I(x, y) = H(x) + H(y) - H(x, y)$. (Hint: first notice that $I(x, y) = \sum_{xy} p(x, y)[\log(p(x, y)) - \log(p(x)) - \log(p(y))]$. Then use the definitions of $H(\cdot)$ and conditional probability, and a simple marginalization.)

2. Bayes Nets (15 Mins)

Given this Bayes net:



A) Use inference by enumeration to write down a formula for the query $P(C | b, e)$.

B) Using the properties of the Bayes Net, what is the formula for the probability of the particular atomic event $P(a, b, c, d, e)$?

C) A node n is conditionally independent of all other nodes in the net given its *Markov blanket*. What is the Markov blanket of node C ?

3. Decision Theory and Contingency Planning (25 Mins)

Should you buy a house? This one has a market value of \$100,000. You can get it for \$90,000, and you'll buy it if your expected total cost is less than \$100,000. However, if it has engineering problems (is not OK), you expect it would cost \$12,000 to fix. You estimate that the probability it is OK ($P(OK)$) is .75.

A) What is your expected net gain (or loss) in dollars if you buy?

B) Your brother-in-law offers to check over the house, giving you more information. Maybe you should make a contingency plan, with a decision that's conditioned on his results. You estimate these probabilities if the house passes his tests:

$$P(\text{pass} | OK) = .80$$

$$P(\text{pass} | \neg OK) = .40$$

What's the probability the house passes the tests? (Hint: marginalize).

C) What's the probability that the house is OK (or not) given each possible test outcome? That is, write out how to calculate the four terms $P(OK | \text{pass}), \dots, P(\neg OK | \neg \text{pass})$. (hint: Bayes law; don't evaluate fractions and note that two of the terms are 1 minus the other two).

D) Now you've got enough to calculate your expected value of buying given the house passes, and the expected value of buying given that it does not pass. Write out expressions for those values (no need to evaluate them).

E) Suppose that both your expected values from D) are positive (you'll buy) whether the house passes or not. What is the value of information for the test? And what is your optimal contingency plan?

Answers:

A) Expected gain is $P(OK)(Value - Price) + (1 - P(OK))(Value - Price - Fix)$

B) $P(pass)$ by marginalization is $P(pass | OK)P(OK) + P(pass | \neg OK)(1 - P(OK))$.

C)

$$P(OK | pass) = (P(pass | OK)(P(OK)))/P(pass)$$

$$P(\neg OK | pass) = 1 - P(OK | pass)$$

$$P(OK | \neg pass) = (P(\neg pass | OK)(P(OK)))/(1 - P(pass))$$

$$P(\neg OK | \neg pass) = 1 - P(OK | \neg pass)$$

D) Substitute $P(OK | pass)$ for $P(OK)$ in part A) for expected gain of buying given the houses passes test. Likewise substitute $P(OK | \neg pass)$ for $P(OK)$ in part A) for expected gain of buying given it fails.

E) Under the supposition, the test does not change your plan, so it tells you nothing, so its value is zero, and your contingency plan is not to test but to buy anyway.

4. Learning (20 Mins)

A) You're a temporal difference agent in a completely observable environment. Your initial estimate for the utility $U(s)$ of states you haven't visited (initially all) is 0. Your future reward discount $\gamma = 1.0$. Your learning rate $\alpha = 0.1$. Here is the history of one trial taking you from state 1 to an enjoyable state 2 to a mildly painful terminal state 3 (It reads down the columns, left to right).

STEP	0	1	2
state s	1	2	3(terminal)
reward R	-1	5	-4
action a	go2	go3	
new state s'	2	3	

What is your estimate of the utilities $U(1), U(2), U(3)$ at the end of step 1? At the end of step 2?

B) You're an Adaptive Dynamic Programming agent in the same environment. You make two trials, the first being the same as above in part A. Here is the second:

STEP	0	1	2	3	4
state s	1	2	2	2	3(terminal)
reward R	-1	5	5	5	-4
action a	go2	go2	go3	go3	
new state s'	2	2	2	3	

What is your estimate of the transition function (or table) $T(s, a, s')$ at the end of the first trial? At the end of the second trial? Technically, why do you care about $T(s, a, s')$ (*i.e.*, how will you use it?)

Answers: A) The T-D equation is

$$U(s) \leftarrow U(s) + \alpha(R(s) + \gamma U(s') - U(s)).$$

Thus after the first step we revise $U(1)$ to be $0 + .1(-1 + 0 - 0) = -0.1$ and similarly after second we get $U(2) = .5$. This update does not propagate back to $U(1)$ (this time). It also seems reasonable, though it's a case not covered by the update equation, to remember that $U(3) = -4$, since 3 is a terminal state.

B) This is just a matter of adding up what happens and computing averages; from the first trial we get the probabilities $T(1, go2, 2) = 1$, $T(2, go3, 3) = 1$. From the second we have no reason to change $T(1, go2, 2) = 1$, and we add $T(2, go2, 2) = 1$, then we have to revise $T(2, go3, 3)$ to $1/2$, and add $T(2, go3, 2) = 1/2$. Finally have to change $T(2, go3, 2)$ to $1/3$ and $T(2, go3, 3)$ to $2/3$.

As ADP agents we need $T(s, a, s')$ to estimate utilities, which are weighted (by $T(\cdot)$) sums of neighboring utilities.

5. Vision (15 Mins)

You want to classify images of textures like the attached into some smallish number of categories (say about ten). You are going to use some form of clustering, implemented it doesn't matter how (parametric, non-parametric, neural net, nearest-neighbor, whatever). Your job is to extract features from such images for use in such clustering. You've only got the following vision operators: a) sub-sample the image: make a new, "half-size" image (with $1/4$ as many pixels) by taking every other pixel from every other row. b) blur or average the image by convolving with a 2×2 mask of 1's. c) and d) differentiate the image vertically and horizontally by convolving with a 2×1 mask and a 1×2 of 1 and -1.

Describe how you will use these operations to construct features (numbers) from input textures and how these features will help in classifying the textures.