CMPSCI 240: Reasoning Under Uncertainty
Final Exam

May 7, 2013.

Name: ______________________________ ID: ____________________

Instructions:

- Answer the questions directly on the exam pages.
- Show all your work for each question. Giving more detail including comments and explanations can help with assignment of partial credit.
- If the answer to a question is a number, you may give your answer using arithmetic operations, such as addition, multiplication, and factorial (e.g., "$9 \times 35! + 2$" or "$0.5 \times 0.3/(0.2 \times 0.5 + 0.9 \times 0.1)$" is fine) unless the problem says otherwise.
- If you need extra space, use the back of a page.
- No books, notes, calculators or other electronic devices are allowed. Any cheating will result in a grade of 0.
- If you have questions during the exam, raise your hand.
- The formulas for some standard random variables can be found on the last page.

Question	Value	Points Earned
1	10	
2	10	
3	10	
4	10	
5	10	
6	10	
7 (Extra Credit)	5	
Total	60+5	

Question 1. (*10 points*) Indicate whether each of the following five statements is TRUE or FALSE. No justification is required.

1.1 (*2 points*): *If A and B are disjoint events then $P(A \cap B) = 0$.*

Answer: TRUE.

1.2 (*2 points*): *If A and B are independent events then A and B are also independent conditioned on any other event C.*

Answer: FALSE.

1.3 (*2 points*): *Every two-player game where both players have a finite number of pure strategies has a Nash equilibrium.*

Answer: TRUE.

1.4 (*2 points*): *For any random variable X and $t > 0$, $P(X \geq E(X) + t) \leq var(X)/t^2$.*

Answer: TRUE.

1.5 (*2 points*): *If X is a geometric random variable with parameter p, then $P(X \geq t) = (1-p)^{t-1}$.*

Answer: TRUE.

Question 2. (*10 points*) Let X be a random variable where $P(X=1)=1/2, P(X=2)=1/4$, and $P(X=4)=1/4$. Let $Y=4X+1$ be another random variable that depends on X. To get full marks you need to simplify your answers to a single number.

2.1 (*3 points*): *Compute the following quantities:*

$$E(X) = 1/2 + 2/4 + 4/4 = 2 \quad E(X^2) = 1/2 + 2^2/4 + 4^2/4 = 5.5 \quad var(X) = 5.5 - 2^2 = 1.5$$

2.2 (*2 points*): *Compute the following quantities:*

$$E(Y) = 4E(X) + 1 = 9 = (5/2 + 9/4 + 17/4) \qquad var(Y) = 4^2 var(X) = 24 = (5^2/2 + 9^2/4 + 17^2/4 - 9^2)$$

2.3 (*3 points*): *Compute the following quantities:*

$$E(X+Y) = 2 + 9 = 11 \quad var(X+Y) = var(5X+1) = 25var(X) = 37.5 \quad \sigma(X+Y) = \sqrt{37.5}$$

2.4 (*2 points*): *Recall that the covariance of X and Y is defined as $cov(X,Y) = E((X-E(X))(Y-E(Y))) = E(XY) - E(X)E(Y)$. Compute the covariance of X and Y.*

$$\begin{aligned} cov(X,Y) &= E(XY) - E(X)E(Y) \\ &= \frac{1\times 5}{2} + \frac{2\times 9}{4} + \frac{4\times 17}{4} - 9\times 2 \\ &= \frac{10+18}{4} + 17 - 18 \\ &= 6 \end{aligned}$$

Question 3. (*10 points*) Consider the Markov chain with three states $\{s_1, s_2, s_3\}$ and transition matrix M:

$$M = \begin{pmatrix} 1/2 & 1/2 & 0 \\ 1/3 & 1/3 & 1/3 \\ 0 & 1 & 0 \end{pmatrix}$$

For example, if you are in state s_3 then the probability you would move to state s_2 is 1.

3.1 (*2 points*): *Draw the transition diagram of the Markov chain and label each edge with the appropriate transition probability.*

Answer:

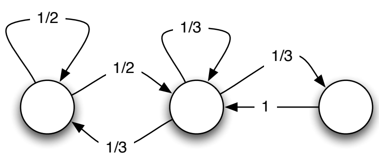

3.2 (*2 points*): *Circle the* two *following terms which apply to the Markov chain.*

irreducible reducible aperiodic periodic

3.3 (*2 points*): *If the chain is in state s_1, what's the probability it's in state s_1 after 2 steps?*

Answer: $1/2 \times 1/2 + 1/2 \times 1/3 = 1/4 + 1/6 = 5/12$

3.4 (*1 points*): *If the chain is in state s_1, what's the probability it's in state s_1 after 3 steps?*

Answer: $1/2\times1/2\times1/2+1/2\times1/2\times1/3+1/2\times1/3\times1/3+1/2\times1/3\times1/2 = 1/8+2/12+1/18 = 25/72$

3.5 (*3 points*): *What's the steady state distribution of the Markov chain?*

Answer: Let the steady state distribution be $(a, b, 1-a-b)$ for some $0 \leq a, b \leq 1$. Since $(a, b, 1-a-b)M = (a, b, 1-a-b)$ we get the following three equations.

$$\begin{aligned} a/2 + b/3 &= a & (1) \\ a/2 + b/3 + 1 - a - b &= b & (2) \\ b/3 &= 1 - a - b & (3) \end{aligned}$$

For the first equation we deduce $2b = 3a$. From the last equation we deduce $4b = 3 - 3a$ and hence $4b = 3 - 2b$. Therefore $b = 1/2$ and $a = 1/3$. Hence the stationary distribution is $(1/3, 1/2, 1/6)$.

Question 4. (*10 points*) Your magician friend gives you a coin. You have two hypotheses:

$$H_1 = \text{"it is a normal coin with equal chance of landing heads or tails."}$$

$$H_2 = \text{"it is a special coin where both sides are tails."}$$

Your prior probabilities for these hypotheses are $P(H_1) = 8/9$ and $P(H_2) = 1/9$. Let D be the event that when you toss the coin two times it lands tails twice.

4.1 (*2 points*): *What's the value of the likelihoods:*

$$P(D|H_1) = 1/4 \qquad P(D|H_2) = 1$$

4.2 (*2 points*): *Given that event D is observed, which is the maximum a posteri (MAP) hypothesis? Show your working.*

Answer: Since $P(D|H_1)P(H_1) = 1/4 \times 8/9 = 2/9$ is greater than $P(D|H_2)P(H_2) = 1/9$, the MAP hypothesis is H_1.

4.3 (*2 points*): *What is the value of $P(D)$ based on the priors above?*

Answer: By the law of total probability $P(D) = P(D|H_1)P(H_1) + P(D|H_2)P(H_2) = 2/9 + 1/9 = 1/3$.

4.4 (*2 points*): *Let F be the event that your friend tells you the coin is special. You don't fully trust your friend and assume the following likelihoods $P(F|H_1) = 1/3$ and $P(F|H_2) = 1$. If events D and F are both observed which hypothesis should you pick? You may assume that your friend talks to you before you start tossing the coin.*

Answer: It is natural to assume that F and D are independent given the hypotheses since conditioned on an hypothesis the coin tossing will be independent of what your friend tells you. Since $P(F|H_1)P(D|H_1)P(H_1) = 1/3 \times 2/9 = 2/27$ is less than $P(F|H_2)P(D|H_2)P(H_2) = 1 \times 1/9 = 1/9$, the MAP hypothesis is now H_2.

4.5 (*2 points*): *Let $X = i$ if the ith hypothesis is true. Let $Y = 1$ if event D occurs and $Y = 0$ otherwise. Let $Z = 1$ if event F occurs and $Z = 0$ otherwise. Draw the Bayesian network for X, Y, and Z based on the scenario above.*

Answer:

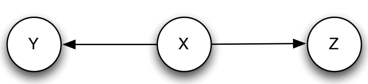

Question 5. (*10 points*) Suppose you store a 128 bit file, i.e., a binary string of length 128, on a really old hard drive that can support such small files. Because the hard drive is unreliable, the file might become corrupted. In particular, each bit is independently flipped (i.e., a 0 changes to a 1 or a 1 changes to a 0) with probability 1/128. Let X be the number of bits that are flipped.

5.1 (*3 points*): *What is the value of the following quantities:*

$P(X=0) = (127/128)^{128}$ $E(X) = (1/128) \times 128 = 1$ $var(X) = (1/128) \times (127/128) \times 128 = 127/128$

5.2 (*2 points*): *Let A be the event that the first bit was flipped. What are the following probabilities?*

$P(X=2) = \binom{128}{2}(1/128)^2(1-1/128)^{126}$ $P(A|X=2) = \frac{\binom{127}{1}(1/128)^2(1-1/128)^{126}}{\binom{128}{2}(1/128)^2(1-1/128)^{126}} = \frac{127}{128 \times 127/2} = \frac{2}{128}$

5.3 (*2 points*): *Suppose the file format for 128 bit files has some built-in error detection. In particular, a file is in the format if and only if the number of ones is even. Write an expression for the probability that a file with this property still satisfies the property after the bits are flipped.*

Answer:

$$\begin{aligned} P(\text{still satisfies the format}) &= P(\text{an even number of bits get flipped}) \\ &= \sum_{i \in \{0,2,4,\ldots n\}} \binom{128}{i}(1/128)^i(1-1/128)^{128-i} \end{aligned}$$

5.4 (*2 points*): *Because the hard drive is unreliable, you decide to store 5 separate copies of the file. What is the probability that the majority (i.e., 3 or more) of the copies do not become corrupted? You may assume that the bits in the copies are independently flipped with probability* 1/128.

Answer: Let $p = (127/128)^{128}$ be the probability a copy doesn't become corrupted. Then the probability that 3 or more copies don't become corrupted is $\binom{5}{3}p^3(1-p)^2 + \binom{5}{4}p^4(1-p)^1 + \binom{5}{5}p^5(1-p)^0$.

5.5 (*1 points*): *When you need to retrieve the file, what would you do if all 5 copies are no longer identical? In particular, how would you take into account the file format mentioned above?*

Answer: If x^1, x^2, x^3, x^4, x^5 are the copies, best solution is to return the valid file y that minimizes $d(x^1, y) + d(x^2, y) + d(x^3, y) + d(x^4, y) + d(x^5, y)$.

Question 6. (*10 points*) Suppose the towns of Amherst and Northampton go to war. They move their troops to either side of the Connecticut River and the situation is very tense.

- If both sides attack, both towns lose 10 points.
- If both sides withdraw, both towns win 0 points.
- If one side attacks and the other side withdraws, the attacking side wins 1 point and the withdrawing side loses 1 point.

Suppose Amherst attacks with probability p and withdraws with probability $1 - p$. Northampton attacks with probability q and withdraws with probability $1 - q$.

6.1 (*2 points*): *What is the pay-off matrix for this game?*

Answer:

	N attacks	N withdraws
A attacks	$(-10, -10)$	$(1, -1)$
A withdraws	$(-1, 1)$	$(0, 0)$

6.2 (*2 points*): *As a function of p and q, what is the expected payoff for Amherst?*

Answer: $-10 \times pq + 1 \times p(1 - q) - 1 \times q(1 - p) = -10pq + p - q$

6.3 (*1 points*): *If Amherst Intelligence finds out Northampton will definitely withdraw, i.e., $q = 0$, what strategy should Amherst adopt?*

Answer: They should attack.

6.4 (*2 points*): *If Amherst Intelligence finds out $q = 1/2$, what strategy should Amherst adopt and what would be the expected payoff for Amherst?*

Answer: When $q = 1/2$, Amherst's expected payoff is $-4p - 1/2$. This is maximized when $p = 0$. Hence, Amherst should withdraw and the expected payoff is $-1/2$.

6.5 (*3 points*): *Find* all *Nash equilibriums for the game.*

Answer: The two pure Nash equilibriums are (Amherst attacks, Northampton withdraws) and (Amherst withdraws, Northampton attacks). Since Amherst's expected payoff $-10pq + p - q = (1 - 10q)p - q$ doesn't depend on p when $q = 1/10$ and Northampton's expected payoff $-10pq + q - p = (1 - 10p)q - p$ doesn't depend on q when $p = 1/10$, each side attacking with probability $1/10$ is also a Nash equilibrium.

Question 7. (*5 points*) ***Extra Credit:*** Suppose you have a biased coin where the probability of heads is $p > 0$ and the probability of tails is $(1-p) > 0$. Unfortunately p is unknown.

7.1 (*3 points*): *Consider the following experiment for measuring the bias of the coin: toss the coin t times and let X be the number of times the coin lands heads. For your estimate of p, you return $\tilde{p} = X/t$. Use the Chebyshev bound to estimate $P(|p - \tilde{p}| \geq 1/100)$ and find a value of t such that $P(|p - \tilde{p}| \geq 1/100) \leq 1/100$.*

Answer:

$$\begin{aligned} P(|p-\tilde{p}| \geq 1/100) &= P(|tp - X| \geq t/100) \\ &\leq \frac{var(X)}{(t/100)^2} = \frac{p(1-p)t}{(t/100)^2} \leq \frac{100^2}{4t} \end{aligned}$$

by appealing to the Chebyshev bound and noting that $p(1-p)$ is always less than $1/4$. Hence if $t \geq 100^3/4$ then $P(|p - \tilde{p}| \geq 1/100) \leq 1/100$.

7.2 (*2 points*): *Design an experiment using your biased coin and define a random variable X based on the outcomes such that $P(X = 0) = 1/2$ and $P(X = 1) = 1/2$.*

Answer: Flip the coin twice. If the results are different and it was tails first, output 1. If the results are different and it was heads first, output 0. If the results were the same, repeat. In each round the probability of outputting 0 is $(1-p)p$ and this is equal to $p(1-p)$, i.e., the probability of outputting 1. Since you must output 0 or 1 eventually, the probability of outputting 0 is $1/2$ and the probability of outputting 1 is $1/2$.

Standard Random Variables

- **Bernoulli Random Variable with parameter** $p \in [0,1]$:

$$P(X=k) = \begin{cases} 1-p & \text{if } k=0 \\ p & \text{if } k=1 \end{cases} \quad , \quad E(X) = p \quad , \quad var(X) = p(1-p)$$

- **Binomial Random Variable with parameters** $p \in [0,1]$ **and** $N \in \{1,2,3,\ldots\}$:

For $k \in \{0,1,2,\ldots,N\}$: $P(X=k) = \binom{N}{k} p^k \,(1-p)^{N-k}$, $E(X) = Np$, $var(X) = Np(1-p)$

- **Geometric Random Variable with parameter** $p \in [0,1]$:

For $k \in \{1,2,3,\ldots\}$: $P(X=k) = (1-p)^{k-1} \cdot p$, $E(X) = \frac{1}{p}$, $var(X) = (1-p)/p^2$

- **Poisson Random Variable with parameter** $\lambda > 0$:

For $k \in \{0,1,2,\ldots\}$: $P(X=k) = e^{-\lambda}\frac{\lambda^k}{k!}$, $E(X) = \lambda$, $var(X) = \lambda$

- **Discrete Uniform Random Variable with parameters** $a, b \in \mathbb{Z}$ **and** $a < b$:

For $k \in \{a, a+1, \ldots b\}$: $P(X=k) = \frac{1}{b-a+1}$, $E(X) = \frac{a+b}{2}$, $var(X) = \frac{(b-a+1)^2 - 1}{12}$

Bayes Formula

- If $A_1, \ldots, A_n$ partition Ω then for any event B:

$$P(A_i|B) = \frac{P(A_i \cap B)}{P(B)} = \frac{P(B|A_i)P(A_i)}{\sum_{j=1}^{n} P(B|A_j)P(A_j)}$$