# Q520: Take-Home Final Exam

This take-home exam is completely open book/open notes. But please do not discuss the exam with anyone: I'm interested to see how everyone does on their own on it.

My intention with this exam is to show some of the connections between different topics in the course. I know that it might have seemed that we wandered from topic to topic. So this exam and my last two lectures are an attempt to show some of the many connections.

All of the individual parts of this exam will count the same. I might change this if it looks like one or another parts turn out too hard.

This exam is due on Wednesday, April 30 at 9:00 PM. The best way to get it to me is to put it under the door of my office, Rawles 323. (Rawles Hall is near the corner of Third St. and Woodlawn.) Alternatively, you could put it in the mailroom on the first floor, near the elevators. But if you do that, please be sure you put it in the right box.

(Needless to say, you are allowed to be early!)

## 1 A Shorter Problem from our work on Linear Algebra

**a** What does it mean to say that vectors $v_1, \ldots, v_n$ are *linearly independent*?

**b** Define the notion of a *rank* of a matrix.

**c** Give an example of a $4 \times 4$ matrix of rank 1, another of rank 2, another of rank 3, and a final one of rank 4.

**d** Our work on the linear algebra showed us how to take a matrix $M$, say of rank $n$, some number $k \leq n$, and give the "closest possible" matrix to $M$ of rank $n$. Describe how we do this in terms of the eigenvalues, eigenvectors, diagonal matrices, etc. Then say why it is of interest. [Most of your work in this last part will be re-reading stuff from your notes and various handouts. The point is to write it in your own words, and thus to review the whole thing.]

## 2 A short problem: probability and entropy review

Let's consider a shuffled deck of cards. (Formally, we are considering a probability space of all 52! possible arrangements of the 52 cards.) Let $X$ be a random variable on this space that tells the first card. So $X$ has 52 possible values, and all of them are equally likely. Let $Y$ be a random variable on this space that tells the second card. Again, $Y$ has 52 possible values, all with the same probability.

1. What is meant by the random variable $X, Y$?
2. What is the entropy of $X$? [That is, calculate this. You may leave logarithms in your answer.]
3. In terms of the notation for entropy, what is $H_X(Y)$? [Again, calculate it.]

4. One of the axioms of entropy tells us that $H(X,Y) = H(X) + H_X(Y)$. You calculated the two parts of the right-hand side in previous parts of this problem. Now calculate $H(X,Y)$ using your answer to (a), and check out (using whatever facts about logarithms you like) that in fact, the equation $H(X,Y) = H(X) + H_X(Y)$ really does hold.

## 3 Hopfield Nets and Linear Algebra

Suppose that we have two vectors (of 1's and −1's) of length 6, say

$$(a_1, a_2, a_3, a_4, a_5, a_6) \quad \text{and} \quad (b_1, b_2, b_3, b_4, b_5, b_6).$$

Suppose also that these are orthogonal (that is, their dot product is zero).

Consider the Hopfield net that we get from these two vectors. That is, we use our two vectors to set the weights of a net by the method from the first day of class. As we know, the resulting Hopfield net is *not* guaranteed to have the given vectors as local minima of energy.

1. Are the two given vectors above fixed points of the net (running it sequentially)? Are they local minima of energy? Why? (If you say yes, please give a reason.)

2. If the given vectors are linearly independent but not orthogonal, what if anything can we say?

[Hints: this one takes a little algebra, so don't be afraid if that's what you're getting into. You will need to know that $a_1^2 = a_2^2 = \cdots = a_6^2 = 1$, since all the $a$'s are 1 or $-1$.]

## 4 Satisfiable sentences in logic

This problem is partly a pure logic problem, and partly a motivation for the work in Problem 5 below.

A logic sentence $S$ is *satisfiable* if there is a way to assign truth values to the basic letters that occur in $S$ in such a way that $S$ comes out *true* according to the truth-values.

For example, consider

$$(P \vee Q) \wedge (\neg P \vee Q).$$

We can make this true by setting $P = true$ and $Q = true$. This is called a *satisfying assignment* of the original sentence. (Another satisfying assignment would be $P = false$ and $Q = true$.)

The sentence $P \wedge \neg P$ is not satisfiable.

The best way to tell if a sentence is satisfiable or not is probably to try out all the possible truth assignments to the variables. (You can also try reasoning it out.) There are various other proceedures that people sometimes present. But for small sentences, you might just as well try out everything.

1. Is $(P \to Q) \wedge (\neg P \wedge \neg Q)$ satisfiable or not?

2. Is $(\neg P \vee R) \wedge (\neg Q \vee R) \wedge (P \vee Q)$ satisfiable or not?

## 5 Markov Chains: a review

Suppose we start with three urns, each with one ball. We repeatedly select an urn at random and do the following. If the urn is empty, we stop the whole process. Otherwise, we pick one ball from the urn at random and put it in one of the other two urns. All choices in this problem are made with equal likelihood. We keep going like this until we select an empty urn.

Your problem: first, formulate this as a Markov chain. You need to decide what the states are and what the transitions between the states are.

Second: given that we start in the state with one ball in each urn, how many steps do we expect it to take before we end? You may decide yourself on what you call a "step" here, but one natural choice would be to count the number of times we select an urn. Whatever you decide, please make it clear what your choice is. Also, please be sure to justify your answer.

[Suggestion: if you make a Markov chain with lots of states, it will be too hard to get the expected number by hand. You may of course use a computer program if you like. You also can make your life easier by getting a Markov chain with a small number of states. It is possible in this problem to do this. Finally, if you are stuck on the actual calculations, you certainly will get partial credit for a clear explanation of what steps you would follow to get the actual answer.]

## 6 Bayesian Nets and Logic

Consider the Bayesian Net below:

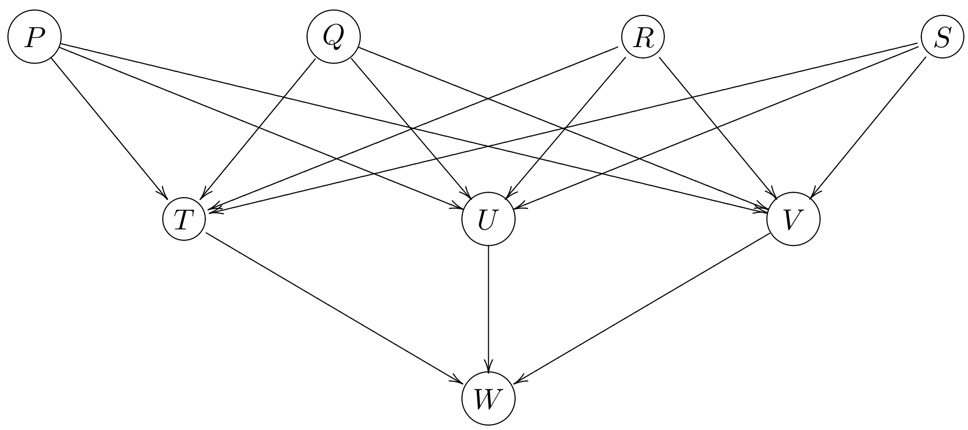


The value spaces of all the variables are $\{t, f\}$. These stand for "true" and "false". I will present the tables, but if you want to see where they come from you should look down at the logic sentence in (1) below.

The tables for the Bayesian net are: $\Pr(P = t) = .5$, $\Pr(P = f) = .5$; $\Pr(Q = t) = .5$,

$\Pr(Q = f) = .5$; similarly for $R$ and $S$;

$$\begin{aligned}
\Pr(T=t|P=t,Q=t,R=t,S=t) &= 1\\
\Pr(T=t|P=t,Q=t,R=t,S=f) &= 1\\
\Pr(T=t|P=t,Q=t,R=f,S=t) &= 1\\
\Pr(T=t|P=t,Q=t,R=f,S=f) &= 1\\
\Pr(T=t|P=t,Q=f,R=t,S=t) &= 1\\
\Pr(T=t|P=t,Q=f,R=t,S=f) &= 1\\
\Pr(T=t|P=t,Q=f,R=f,S=t) &= 1\\
\Pr(T=t|P=t,Q=f,R=f,S=f) &= 1\\
\Pr(T=t|P=f,Q=t,R=t,S=t) &= 1\\
\Pr(T=t|P=f,Q=t,R=t,S=f) &= 1\\
\Pr(T=t|P=f,Q=t,R=f,S=t) &= 1\\
\Pr(T=t|P=f,Q=t,R=f,S=f) &= 1\\
\Pr(T=t|P=f,Q=f,R=t,S=t) &= 1\\
\Pr(T=t|P=f,Q=f,R=t,S=f) &= 1\\
\Pr(T=t|P=f,Q=f,R=f,S=t) &= 0\\
\Pr(T=t|P=f,Q=f,R=f,S=f) &= 0
\end{aligned}$$

$$\begin{aligned}
\Pr(U=t|P=t,Q=t,R=t,S=t) &= 1\\
\Pr(U=t|P=t,Q=t,R=t,S=f) &= 0\\
\Pr(U=t|P=t,Q=t,R=f,S=t) &= 1\\
\Pr(U=t|P=t,Q=t,R=f,S=f) &= 1\\
\Pr(U=t|P=t,Q=f,R=t,S=t) &= 1\\
\Pr(U=t|P=t,Q=f,R=t,S=f) &= 0\\
\Pr(U=t|P=t,Q=f,R=f,S=t) &= 1\\
\Pr(U=t|P=t,Q=f,R=f,S=f) &= 1\\
\Pr(U=t|P=f,Q=t,R=t,S=t) &= 1\\
\Pr(U=t|P=f,Q=t,R=t,S=f) &= 1\\
\Pr(U=t|P=f,Q=t,R=f,S=t) &= 1\\
\Pr(U=t|P=f,Q=t,R=f,S=f) &= 1\\
\Pr(U=t|P=f,Q=f,R=t,S=t) &= 1\\
\Pr(U=t|P=f,Q=f,R=t,S=f) &= 1\\
\Pr(U=t|P=f,Q=f,R=f,S=t) &= 1\\
\Pr(U=t|P=f,Q=f,R=f,S=f) &= 1
\end{aligned}$$

$$\begin{aligned}
\Pr(V=t|P=t,Q=t,R=t,S=t) &= 1\\
\Pr(V=t|P=t,Q=t,R=t,S=f) &= 1\\
\Pr(V=t|P=t,Q=t,R=f,S=t) &= 1\\
\Pr(V=t|P=t,Q=t,R=f,S=f) &= 1\\
\Pr(V=t|P=t,Q=f,R=t,S=t) &= 1\\
\Pr(V=t|P=t,Q=f,R=t,S=f) &= 1\\
\Pr(V=t|P=t,Q=f,R=f,S=t) &= 1\\
\Pr(V=t|P=t,Q=f,R=f,S=f) &= 1\\
\Pr(V=t|P=f,Q=t,R=t,S=t) &= 0\\
\Pr(V=t|P=f,Q=t,R=t,S=f) &= 0\\
\Pr(V=t|P=f,Q=t,R=f,S=t) &= 1\\
\Pr(V=t|P=f,Q=t,R=f,S=f) &= 1\\
\Pr(V=t|P=f,Q=f,R=t,S=t) &= 1\\
\Pr(V=t|P=f,Q=f,R=t,S=f) &= 1\\
\Pr(V=t|P=f,Q=f,R=f,S=t) &= 1\\
\Pr(V=t|P=f,Q=f,R=f,S=f) &= 1
\end{aligned}$$

$$\begin{aligned}
\Pr(W=t|T=t,U=t,V=t) &= 1\\
\Pr(W=t|T=t,U=t,V=f) &= 0\\
\Pr(W=t|T=t,U=f,V=t) &= 0\\
\Pr(W=t|T=t,U=f,V=f) &= 0\\
\Pr(W=t|T=f,U=t,V=t) &= 0\\
\Pr(W=t|T=f,U=t,V=f) &= 0\\
\Pr(W=t|T=f,U=f,V=t) &= 0\\
\Pr(W=t|T=f,U=f,V=f) &= 0
\end{aligned}$$

Recall the difference between the specification of a Bayesian net and an actual net. The actual net has a probability space $\mathcal{S}$ and random variables over $\mathcal{S}$ corresponding to the nodes of our dag.

1. In the actual net, what is the probability

$$\Pr(P=t,Q=f,R=t,S=f,T=t,U=f,V=t,W=f)?$$

2. In the actual net, what is the probability

$$\Pr(P=t,Q=f,R=f,S=f,T=t,U=f,V=t,W=f)?$$

3. Consider again the logic formula

$$(P \vee Q \vee R) \wedge (\neg P \vee \neg R \vee S) \wedge (P \vee \neg Q \vee \neg R). \tag{1}$$

   Show that (1) has a satisfying assignment if and only if in the space $\mathcal{S}$, $\Pr(W = t) > 0$.

   Please do not try to calculate out the huge space $\mathcal{S}$ and all the probabilities. Also, please do not bother to try to find a satisfying assignment for (1). The point here is to find a short *reason* without doing the drudgery.

The point of this problem is that it shows that a natural logic problem can be simulated on a Bayesian net. So in a sense, Bayesian nets are a *more powerful* formalism than basic logic. To put things another way: it is not believed to be tractable to decide whether a sentence like (1) has a satisfying assignment. There is no known way that does better than to simply try out all the possibilities. If (and this is a big *IF*) we had a program that could process Bayesian nets quickly and read off the actual net probabilities from the specifications, we *would* have a way to test for the existence of satisfying assignments. (This is just what the last part of the problem suggests.) So knowing this fact, people tend to believe that there are no tractable algorithms that will do tell us everything we could want to know about all possible Bayesian nets.

[This problem might scare you because it is so long. Please don't let it do that! If you review the way that specifications of Bayesian nets give us actual nets, this problem will not be so hard. The answer is fairly short.]

## 7 Translation from English to first-order logic

Translate the following sentences into logic notation, using the following dictionary:

| | |
|---|---|
| $L(x, y)$ | $x$ loves $y$ |
| $S(x, y)$ | $x$ sees $y$ |
| $j$ | John |
| $m$ | Mary |

1. John sees Mary.
2. John sees himself.
3. Someone sees himself.
4. Everyone is seen by either John or Mary.
5. Someone sees someone other than him/herself.
6. Someone who sees him/herself loves John.
7. John sees everyone who Mary sees.
8. If anyone sees John, then anyone sees John. [The reading I have in mind here is where John is either very-well-hidden, or else in obvious view. If you have trouble with this, ask me to say it out loud. The point of this problem is to highlight the difficulties surrounding words like *any*.]

## 8 The Three Wisemen

This is based a problem popularized by John McCarthy, one of the founders of Artificial Intelligence.

Three "wisemen" are imprisoned by a king. They are lead blindfolded to a room where there are three seats facing forward, and the wisemen are arranged as follows:

$A \rightarrow$ $\qquad\qquad$ $B \rightarrow$ $\qquad\qquad$ $C \rightarrow$

Then the king takes out some hats colored white or black. He places them on the wisemen, and removes the blindfolds. He announces to the group that everyone is wearing a white or black hat. The only thing the wisemen can see is the seat(s) indicated by the arrows. They can't see themselves. They have no idea how many white or black hats were used.

We assume that $A$ can see both $B$ and $C$, and also that $C$ knows that the one immediately behind him/her is $B$.

1. Write out the Kripke model for this situation in full. Use atomic sentences $w$ (for *white*) and $b$ (for *black* Call the model overall $W$.

2. Write three sentences in English that you think are true (intuitively) about the situation at this point. These should be sentences which translate to the formal logic that we have studied. Then check that your formal sentences really are true in the Kripke model that you have given.

3. The king then announces "At least one of you has a white hat."

   Write out the Kripke model for the situation now. Call your model $X$. Be sure to say how your model relates to the work we did in class on public announcements. That, is what is the relation between $W$ and $X$?

4. The king tells everyone that they should think about whether they know the color of their hats. They have to announce their knowledge at the end of each minute until someone knows. Being "wise," they do a perfect job of logical reasoning.

   Is it possible for the hats to be arranged so that $C$ is the first person to know his/her hat color? (This seems counterintuitive, at least on first glance, since $C$ starts out by seeing less than everyone else.) If so, how? And what does this have to do with our logic and representation?