CS 530 — Principles of AI
Sample Problems

Problem 1. The goal of this problem is to familiarize yourself with Belief Networks as flexible specifications of probability distributions. The following network is a clustering model:

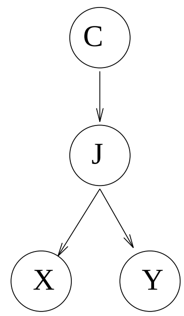

C is a variable for the underlying category of an object; X and Y are observable object features (assume they are discrete), and J is a hidden variable representing which cluster the object belongs to.

1a. What probability distributions do you have to specify to determine a joint distribution on X, Y, J and C in this network?

1b. Suppose that the distribution of objects is as follows. A quarter of all objects are in class $C = 1$ and have $X =$true and $Y =$false. A quarter of all objects are in class $C = 1$ and have $X =$false and $Y =$true. A quarter are in class $C = 2$ and have $X =$true and $Y =$true. A quarter are in class $C = 2$ and have $X =$false and $Y =$false. Specify parameters for this network (as outlined problem 1a) to capture this distribution exactly.

1c. Can the belief network shown below describe the distribution of problem 1b exactly? Why or why not?

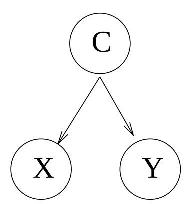

1d. Can the belief network shown below describe the distribution of problem 1b exactly? Why or why not?

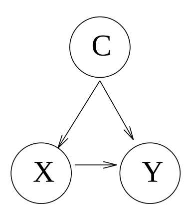

Problem 3. Here is an influence diagram (the extension of Belief Networks that describes choice as well as uncertainty).

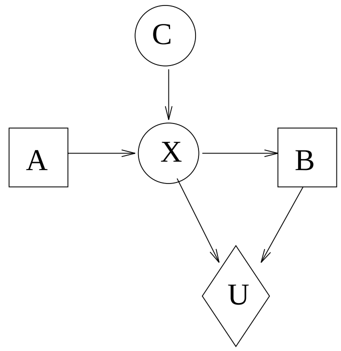

In summary, this diagram describes the following situation. The world starts out in a hidden state represented by the variable C, which is true or false. Without knowing C, the agent has to make a choice of action A, either a or a'. A and C probabilistically determine an effect X, which is true or false. The agent then observes X and makes a choice at B of b or b'. The utility of the outcome depends on X and B.

Here are parameters of the model which yield a complete specification of the agent's decision-making.

$P(C)$ 0.6

$P(X|A,C)$

A	C	$P(X=T|A,C)$
a	T	0.9
a	F	0.4
a'	T	0.3
a'	F	0.3

$U(X,B)$

X	B	$U(X,B)$
T	b	0.5
T	b'	0.3
F	b	0.4
F	b'	0.7

Problem 3a. Formulate an equivalent tree decision model.

Problem 3b. What is the optimal policy in the tree decision model? What is its utility?

Problem 3c. Explain how the independence assumptions of the influence diagram allow you to specify the optimal policy more compactly.

Problem 4. The first three problems look at taking *actions* that respond to the *inferred state* of the world. As usual, the state of the world can be thought of categorically; the current category C takes the value $c1$ or $c2$. Suppose that $P(C = c1)$ is one third and $P(C = c2)$ is two thirds.

As usual, in making decisions, we just get some measurement x which gives noisy evidence about the true category. Suppose our our measurement x is such that $P(x|C = c1)$ is three quarters and $P(x|C = c2)$ is one quarter. Use Bayes's theorem to calculate the posterior distribution $P(C|x)$.

Problem 5. Suppose that the variable C in this case represents the category of an object that the agent is interacting with. The agent can take any of three actions A in response to the object: it can *save* it (s), it can *discard* it (d) or it can ask for *help* with it (h).

Suppose that the agent's decisions lead to outcomes as follows. The agent is used for quality control at the AA Widget factory. AA is already prepared to pay \$.40 per widget for shipping to stores. We measure utility by dollar improvement per widget from this \$-.40 baseline.

- When the agent decides to save the widget, it ships. $C = c1$ are good widgets, on which AA recoups \$1. (So the profit is \$.60, but that's not what we're measuring here.) $C = c2$ are faulty widgets, which never sell so AA recoups none of its costs.
- When the agent decides to discard the widget, AA keeps the \$.40 shipping cost.
- When the agent asks for help, the human identifies the category of the widget correctly. AA recoups \$1 on good widgets and keeps the \$.40 shipping cost on bad widgets. However, AA has to pay the person q for their labor.

Write out the agent's utility function $U(A,C)$.

Problem 6. For which values of q in Problem 5 should the agent ask for a human's help with the object you studied in Problem 1?

Problem 7. Draw an influence diagram corresponding to this decision problem.

Note. The moral of this problem is that Bayesian inference gives an agent a precise estimate of its uncertainty between two real-world situations. The agent can therefore optimize how it responds—or doesn't respond!—to this uncertainty based on the expected but uncertain outcome.

Problem 8. This block of problems compares two methodologies for modeling and classifying sequences that we studied this semester.

- The first strategy is to define *features* of sequences—predefined predicates that are true or false of each sequence—and treat the feature values as observations of a sequence. (Think of features for word-occurrence in text classification.) You do inference with a model where features give noisy evidence about category directly.

- The second strategy is to model the sequence as a whole as unfolding in time, using a Hidden Markov model. (We studied a variety of problems such as speech recognition from this point of view.)

Suppose you can predefine a specific collection of subsequences that give good evidence about the category of a sequence ahead of time. (For example, a sequence might very strongly tend to belong to category $c1$ when it contains observations *agdsahsd* in order somewhere within it.) Explain one way you could exploit this background knowledge in using a feature-based strategy in learning and reasoning about sequences.

Problem 9. What would be a difficulty with exploiting this kind of information in HMMs?

Problem 10. Suppose you discover that sequences of each category tend to have long subsequences in which some observations appear frequently (e.g., x and y) alternating with long sequences in which other observations appear frequently (e.g., z and w). Explain one way you could exploit this background knowledge in using HMMs in learning and reasoning about sequences.

Problem 11. What would be a difficulty with exploiting this kind of information with a feature-based sequence classifier?

Note. The moral of this problem is that you need to think carefully about the statistical properties of your observations, not just about their structure, in order to choose the best model to describe them.

Problem 12. Consider the following abstract decision-making scenario. An agent starts with perfect knowledge of the state of the world C. The agent must choose an action A based on C. However, action A takes time to execute; perhaps A requires resources that become available slowly, or perhaps A is delayed in its effects. Thus, when A happens, the state of the world has changed in a probabilistic way from C to D, so we have a model $P(D|C)$. A and D together determine the utility $U(A,D)$ that the agent achieves. What is the mathematical expression in terms of these parameters that determines which action A the agent should choose?

Problem 13. Draw an influence diagram corresponding to this situation.

Note. Evidently, in an influence model that describes change over time, correct reasoning must involve anticipating the future as the model describes it. Of course this is true of all models for decisions.

Problem 14. In perception, we frequently find a situation that can be represented as a Belief network of the following sort:

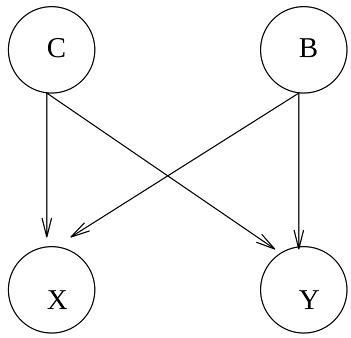

Observable measurements X and Y reflect both the true state of an object C and the background conditions B under which the measurements are taken. For example, in vision, the positions of features in an image depend on the real-world positions of objects and on the state of a camera (e.g., zoom).

What probability distributions are required in the Belief network above, assuming all variables are discrete?

Problem 15. What information is required if all the variables are continuous?

Problem 16. In this model, assuming all the variables are discrete, categorization continues to depend on the joint distribution of the category and the measurements $P(X,Y,C)$, where X and Y are of course given. Write a rule that determines this probability (as a function of C) from these parameters.

Problem 17. Write out the analogous rule for estimation that applies when all the variables are continuous?

Problem 18. Refer to your answer to Problem 16 or 17 to describe briefly and qualitatively how more accurate prior knowledge of background conditions B can lead to a better estimate of the state of the world C from your observations X and Y.

Note. Of course, since this is a Bayesian framework, you can also infer the background conditions B if you have strong expectations about C!

Problem 19. You have designed a classifier that tries to infer a user's state in using a word-processing system (ie., what command are they about to use? etc.) You started from a baseline system whose features included the time since the last keystroke, the pattern of recent mouse movements and selections, and the user's recent navigation through the menu. You added features corresponding to words that the user has just typed. It seems to infer the user's state more accurately now. In fact, your specific hypothesis is that most of the improvement is associated with uses of a few important words, which can be identified based on the extreme values in their class-conditional probability tables.

Describe the conditions for a possible experiment that might provide evidence for this hypothesis. Explain how you would go about making sure that the results of experiment would not be confounded with ceiling or floor effects, or with sampling biases.