A Template for Bayesian Decision Analysis

Bayesian probability theory is concerned primarily with assessing and incorporating the value of information (both prior and new) into the probability of future events. Bayesian methods are used extensively both in economics and in decision theory as a powerful tool for incorporating rational learning and valuing information.

Simply, Bayesian probability always assigns a probability of an event *given* some other event (conditional probabilities). For example, in valuing information, it would be important to know what the probability of an event happening was given the availability of the new information. Generally, the greater the change in the probability, the more value the information has.

First, consider Bayes' Theorem mathematically. We will consider two sets of events, $\Theta = \{\theta\}$ and $\Psi = \{\psi_1, \psi_2, \ldots, \psi_n\}$. Additionally, it is assumed that:

[1] $$P(\theta) \neq 0$$

[2] $$\sum_{i=1}^{n} \psi_i = 1$$

[3] $$P(\psi_i \cup \psi_j) = 0 \quad \forall\, i \neq j$$

If these three conditions are true, then the conditional probability $P(\psi_j | \theta)$ of ψ_j given that θ has been realized is determined by:

[4] $$P(\psi_j | \theta) = \frac{P(\psi_j) P(\theta | \psi_j)}{\Sigma_{i=1}^{n} P(\psi_i) P(\theta | \psi_i)}$$

[4] is known as Bayes' Theorem. [4] can accommodate n discrete events. In fact, Bayesian probability can be formulated to incorporate a continuum of events. We will not consider the continuous case or even the n discrete case above - they are beyond the scope of this class.

We will now consider the details of applying Bayes' Theorem in the context you will most likely find it in this class. These problems can often be phrased as assessing the probability of a hypothesis (H) given an observation (O). It is important that you are familiar with solving these problems by way of the 2×2 tables. We will explore each issue in turn.

Stated as an inference about a hypothesis given an observation, Bayes' Theorem can be represented according to the following:

[5] $$P(H | O) = \frac{P(H) P(O | H)}{P(H) P(O | H) + P(\bar{H}) P(O | \bar{H})}$$

H denotes the hypothesis, O denotes the observation, and the bar over the variables indicates the complement of that probability as represented by the logical "not" or one minus the probability of the event.

From [5], we may derive the following equalities:

$$P(\bar{H}|O)=\frac{P(\bar{H})P(O|\bar{H})}{P(H)P(O|H)+P(\bar{H})P(O|\bar{H})}$$

[6a, b, c] $$P(H|\bar{O})=\frac{P(H)P(\bar{O}|H)}{P(H)P(\bar{O}|H)+P(\bar{H})P(\bar{O}|\bar{H})}$$

$$P(\bar{H}|\bar{O})=\frac{P(\bar{H})P(\bar{O}|\bar{H})}{P(H)P(\bar{O}|H)+P(\bar{H})P(\bar{O}|\bar{H})}$$

Note that $P(H|O)+P(\bar{H}|O)=1$. This implies that $P(H\cup\bar{H}|O)=1$ because $P(H)+P(\bar{H})=1$ by the definition of a probability.

With the above formulas, you can calculate the solution to any Bayesian decision problem. However, it is often more insightful to consider the probabilities in a 2×2 table. Before we turn to the table, consider for a moment the applied meanings of the above conditional (or revised) probabilities. Both $P(H|O)$ and $P(\bar{H}|\bar{O})$ represent accuracy. They are the probabilities of a correct prediction. For example, the probability that some test or hypothesis is true given that it is observed to be true or new information arrives which implies that it is true. Usually, you want large numbers here. Also keep in mind that these probabilities will almost certainly *not* sum to one. They are not complementary.

On the other hand, $P(\bar{H}|O)$ and $P(H|\bar{O})$, are, respectively, false negatives and false positives. The first, a false negative incorrectly predicts that the hypothesis is false given that the observed information says it is true. The second, the false positive, says that the hypothesis is true given that it is in fact false. Usually, you want very small numbers here - particularly when dealing with medical decision making. Often the psychological costs to patients of false positives and false negatives is enormous.

For example, consider a real problem: HIV testing. You know the frequency of an HIV+ diagnosis, $P(HIV)$. You also know the accuracy of a test. Given that a patient is HIV+, the probability that a test will correctly confirm (indicated by the quotes) the presence of HIV is given by $P("HIV"|HIV)$. We will denote the prediction of the test by putting the statement in quotes: the test “says” HIV+, for example. Given that a patient is *not* HIV+, a test will correctly indicate HIV- status with the probability $P("\overline{HIV}"|\overline{HIV})$. What you may be interested in knowing, however, is the probability that you actually are HIV+ given that the test said you were HIV+ or $P(HIV|"HIV")$. Assume the following information:

Probability of Being HIV+	$P(HIV)$	0.0076
Probability of Correct Positive	$P("HIV"|HIV)$	0.976
Probability of Correct Negative	$P("\overline{HIV}"|\overline{HIV})$	0.995

This information can then be represented in a table. The table plays an important role in understanding the relationships between the probabilities. Essentially, we are trying to solve the following problem:

[7] $$P(HIV|"HIV")=\frac{P(HIV)P("HIV"|HIV)}{P(HIV)P("HIV"|HIV)+P(\overline{HIV})P("HIV"|\overline{HIV})}$$

	"HIV"	"$\overline{\text{HIV}}$"	
HIV	$P(\text{"HIV"} \mid HIV)$	$P(\text{"}\overline{HIV}\text{"} \mid HIV)$	$P(HIV)$
$\overline{\text{HIV}}$	$P(\text{"HIV"} \mid \overline{HIV})$	$P(\text{"}\overline{HIV}\text{"} \mid \overline{HIV})$	$P(\overline{HIV})$
	$P(\text{"HIV"})$	$P(\text{"}\overline{HIV}\text{"})$	1.0000

Because we are dealing with the complete set of exclusive states, the rows must each sum to one. A person will either be HIV+ or not. The columns, however, are unlikely to sum to one at this point. We don't yet know the specificity of the tests.

We are given the diagonal elements in the above table. That is, we know $P(\text{"HIV"} \mid HIV)$ and $P(\text{"}\overline{HIV}\text{"} \mid \overline{HIV})$ are 0.976 and 0.995 respectively. We also know that $P(HIV) = 0.0076$. Therefore, we also know that $P(\overline{HIV}) = 1 - P(HIV) = 0.9924$. Similarly, we can find the values for the off-diagonal elements in the table. By doing so, we can obtain the completed table. Remember that the off-diagonal elements are the false negatives and false positives. These values should be small. Finally, the $P(\text{"HIV"})$ and $P(\text{"}\overline{HIV}\text{"})$ are simply the sums of their respective columns.

	"HIV"	"$\overline{\text{HIV}}$"	
HIV	0.0074	0.0002	0.0076
$\overline{\text{HIV}}$	0.0050	0.9874	0.9924
	0.0124	0.9876	1.0000

The values inside the table are obtained by multiplying the values in the first table by their row sums to attain proportions. For example, $0.0074 = 0.976 \times 0.0076$. Now, to solve the problem, we need to know the probability that a person is HIV+ *given* that the test indicated that they are. This means we are concerned only with all the states in the first column (where the test returns "HIV"). We now know that probability:

$$P(HIV \mid \text{"HIV"}) = \frac{P(HIV \cap \text{"HIV"})}{P(\text{"HIV"})} = \frac{0.0074}{0.0124} = 0.5992$$

This indicates that despite seemingly high accuracy levels (97.6% and 99.5%), the probability that a person actually is HIV+ given that this test says they are is only about 60%. On the other hand, the probability that you are not HIV+ *given* that the test says you are not HIV+ is very high:

$$P(\overline{HIV} \mid \text{"}\overline{HIV}\text{"}) = \frac{P(\overline{HIV} \cap \text{"}\overline{HIV}\text{"})}{P(\text{"}\overline{HIV}\text{"})} = \frac{0.9874}{0.9876} = 0.9998$$

The off-diagonal elements (false negatives and false positives) may be found similarly. And, in fact, they are the complements of the above values:

$$P(\overline{HIV} \mid \text{"HIV"}) = 1 - P(HIV \mid \text{"HIV"}) = 0.4008$$
$$P(HIV \mid \text{"}\overline{HIV}\text{"}) = 1 - P(\overline{HIV} \mid \text{"}\overline{HIV}\text{"}) = 0.0002$$

Finally, the values can also be found in a more straight-forward manner by working explicitly with Bayes' formula. Observe the following final example:

$$P(H|O) = \frac{P(H)P(O|H)}{P(H)P(O|H) + P(\bar{H})P(O|\bar{H})}$$

$$P(HIV|"HIV") = \frac{P(HIV)P("HIV"|HIV)}{P(HIV)P("HIV"|HIV) + P(\overline{HIV})P("HIV"|\overline{HIV})}$$

$$= \frac{(0.0076)(0.976)}{[(0.0076)(0.976) + (0.005)(0.9924)]} = \frac{0.0074}{0.0074 + 0.0050} = 0.5992$$

David Rode
Decision Analysis and Decision Support Systems
Department of Social and Decision Sciences
Carnegie Mellon University
May 27, 1997