

Naive Bayes

Hiroshi Shimodaira*

January-March 2020

In the previous chapter we introduced the use of Bayes' Theorem for pattern classification. For a test vector \mathbf{x} , we estimate the posterior probability $P(C_k|\mathbf{x})$ for each class k . To classify \mathbf{x} we choose the class with the largest estimated posterior probability. To do this we re-express the posterior probabilities using Bayes' Theorem:

$$P(C_k|\mathbf{x}) = \frac{P(\mathbf{x}|C_k)P(C_k)}{P(\mathbf{x})} \\ \propto P(\mathbf{x}|C_k)P(C_k).$$

Thus for each class we need to provide an estimate of the likelihood $P(\mathbf{x}|C_k)$ and the prior $P(C_k)$.

To estimate the likelihood we need a statistical *model* which can provide a likelihood estimate. In the 'fish' example, the likelihood was estimated using a histogram for each class: the likelihood of a fish with length x coming from class k was estimated as the relative frequency of fish of length x estimated from the training data for class k .

Imagine if in addition to the length we also had some other information such as weight and circumference. In this case, the input feature vector for each fish would contain 3 elements: (length, weight, circumference). If we have 20 possible values for each feature, then the number of bins in a histogram of these vectors would be $20^3 = 8000$. In this case 100 examples per class would mean that it is not possible to observe examples for the vast majority of the bins. We would need many more examples to obtain reasonable estimates based on relative frequencies. As the number of feature dimensions increases, the total number of possible feature vectors increases *exponentially*.

This severe difficulty that arises from moving to spaces of higher dimension was termed the *curse of dimensionality* by Richard Bellman in the 1950s. This is a critical problem: it becomes impossible to reliably estimate histograms once we have more than a few dimensions, even if we have millions of examples.

In this chapter we shall look at an approach to the problem called the Naive Bayes approximation. After showing how it works in a small, fictitious example, we'll go on to see how this approach may be applied to the problem of text classification, which uses high-dimensional feature vectors (e.g., 10^4 - to 10^7 -dimensional).

*©2014-2020 University of Edinburgh. All rights reserved. This note is heavily based on notes inherited from Steve Renals and Iain Murray.

6.1 The Naive Bayes assumption

One way to deal with the curse of dimensionality is to assume that the different feature dimensions are *independent*.¹ If we are using histograms of relative frequencies to estimate likelihoods, then this means that we construct, for each class, D 1-dimensional histograms, rather than a single D -dimensional histogram. If we assume that each dimension can take M values, we now only need to estimate $D \times M$ relative frequencies, rather than M^D relative frequencies!

Consider a D -dimensional feature vector $\mathbf{x} = (x_1, x_2, \dots, x_D)$. We write the probability of a data point \mathbf{x} given class k as a joint distribution of the D components of \mathbf{x} (Equation (6.1)). We can re-express any joint distribution as a product of conditional distributions using successive applications of the product rule (Equation (6.2)). The result is known as the chain rule of probability, which in this case is:

$$P(\mathbf{x}|C_k) = P(x_1, x_2, \dots, x_D|C_k) \quad (6.1)$$

$$= P(x_1|x_2, \dots, x_D, C_k)P(x_2|x_3, \dots, x_D, C_k) \dots P(x_{D-1}|x_D, C_k)P(x_D|C_k). \quad (6.2)$$

This decomposition in itself doesn't address the curse of dimensionality, since the first term on the right-hand side of Equation (6.2) is conditioned on $(D-1)$ terms, the second on $(D-2)$ terms, and so on.

However we can simplify things if we *naively* assume that the individual feature dimensions (x_1, x_2, \dots, x_D) are independent, that is:

$$P(x_1|x_2, \dots, x_D, C_k) = P(x_1|C_k)$$

$$P(x_2|x_3, \dots, x_D, C_k) = P(x_2|C_k)$$

$$\vdots$$

This is called the *Naive Bayes* assumption. The assumption is drastic and rarely true: for example, in the above case Naive Bayes states that the length of a male fish is independent of its weight and circumference. However making this approximation allows us to have a much simpler form for the likelihood, since Equation (6.2) is simplified to:

$$P(\mathbf{x}|C_k) = P(x_1|C_k)P(x_2|C_k) \dots P(x_D|C_k) = \prod_{d=1}^D P(x_d|C_k). \quad (6.3)$$

We have approximated the probability of a D -dimensional feature vector as a product of D probabilities of the 1-dimensional feature vectors.

Using this assumption, we can express Bayes' theorem as follows:

$$P(C_k|\mathbf{x}) = \frac{P(\mathbf{x}|C_k)P(C_k)}{P(\mathbf{x})} \quad (6.4)$$

$$= \frac{\prod_{d=1}^D P(x_d|C_k)P(C_k)}{\prod_{d=1}^D P(x_d)} \propto P(C_k) \prod_{d=1}^D P(x_d|C_k) \quad (6.5)$$

6.2 Example

The following (fictitious) example comes from the book *Data Mining* by Witten and Frank.

¹We normally assume that features are 'conditionally independent' – a relaxed version of independence, in which features are mutually independent in the class.

Consider a game which is played or not depending on the weather conditions: outlook (sunny / overcast / rainy), temperature (hot / mild / cool), humidity (high / normal) and windy (true / false). The input variable is 4-dimensional, with 2 or 3 values per dimension. There are 36 possible combinations ($3 \times 3 \times 2 \times 2$).

We have an input training set of 14 examples, shown in Table 6.1. As is usually the case in machine learning, there are fewer training examples than possible settings of the input variables; most possible conditions are not directly observed. We can tabulate the data in terms of the frequencies for each of the four input variables (Table 6.2), and in terms of the relative frequencies (Table 6.3). The relative frequencies can be used as probability estimates, for example $P(T = h \mid \text{Play} = Y) = 2/9$.

Outlook	Temperature	Humidity	Windy	Play
sunny	hot	high	false	NO
sunny	hot	high	true	NO
overcast	hot	high	false	YES
rainy	mild	high	false	YES
rainy	cool	normal	false	YES
rainy	cool	normal	true	NO
overcast	cool	normal	true	YES
sunny	mild	high	false	NO
sunny	cool	normal	false	YES
rainy	mild	normal	false	YES
sunny	mild	normal	true	YES
overcast	mild	high	true	YES
overcast	hot	normal	false	YES
rainy	mild	high	true	NO

Table 6.1: Training data for the weather example

Outlook	Y	N	Temperature	Y	N	Humidity	Y	N	Windy	Y	N
sunny	2	3	hot	2	2	high	3	4	false	6	2
overcast	4	0	mild	4	2	normal	6	1	true	3	3
rainy	3	2	cool	3	1						

Table 6.2: Play counts for different weather conditions.

Outlook	Y	N	Temperature	Y	N	Humidity	Y	N	Windy	Y	N
sunny	2/9	3/5	hot	2/9	2/5	high	3/9	4/5	false	6/9	2/5
overcast	4/9	0/5	mild	4/9	2/5	normal	6/9	1/5	true	3/9	3/5
rainy	3/9	2/5	cool	3/9	1/5						

Table 6.3: Play relative frequencies for different weather conditions. There was play on 9/14 cases.

We are given the following test example:

Outlook	Temp.	Humidity	Windy	Play
sunny	cool	high	true	?

This example's feature vector, \mathbf{x} , was not observed in the training set. We can generalise from the other examples by using Naive Bayes:

$$\begin{aligned}
 P(\text{play} = Y \mid \mathbf{x}) &\propto P(\text{play} = Y) \cdot P(O = s \mid \text{play} = Y) \cdot P(T = c \mid \text{play} = Y) \\
 &\quad \cdot P(H = h \mid \text{play} = Y) \cdot P(W = t \mid \text{play} = Y) \\
 &= \frac{9}{14} \cdot \frac{2}{9} \cdot \frac{3}{9} \cdot \frac{3}{9} \cdot \frac{3}{9} \\
 &\approx 0.0053 \\
 P(\text{play} = N \mid \mathbf{x}) &\propto P(\text{play} = N) \cdot P(O = s \mid \text{play} = N) \cdot P(T = c \mid \text{play} = N) \\
 &\quad \cdot P(H = h \mid \text{play} = N) \cdot P(W = t \mid \text{play} = N) \\
 &= \frac{5}{14} \cdot \frac{3}{5} \cdot \frac{1}{5} \cdot \frac{4}{5} \cdot \frac{3}{5} \\
 &\approx 0.0206
 \end{aligned}$$

And the ratio of posterior probabilities is:

$$\frac{P(\text{play} = Y \mid \mathbf{x})}{P(\text{play} = N \mid \mathbf{x})} = \frac{0.0053}{0.0206} \approx 0.26$$

Hence we classify \mathbf{x} as $\text{play} = N$.

6.3 Conclusion

In this chapter we have introduced the Naive Bayes approximation. When we have a multidimensional feature vector, Naive Bayes approximated the likelihoods by considering each feature dimension to be independent:

$$P(x_1, x_2, \dots, x_D \mid C) \approx P(x_1 \mid C) \cdot P(x_2 \mid C) \cdots P(x_D \mid C),$$

which approximates a D -dimensional distribution as D 1-dimensional distributions. If each dimension can take M different values, we need only estimate $M \times D$ probabilities, rather than M^D probabilities.

Exercises

1. Considering the example shown in Section 6.2, predict how likely a game is played for each of the following conditions:
 - (a) No weather conditions are known.
 - (b) Outlook = rainy.
 - (c) Outlook = rainy, Temperature = hot, Humidity = normal, and Windy = false.