

# "Naive Bayes Classifier"

## Wikipedia

A **Bayes classifier** is a simple probabilistic classifier based on applying Bayes' theorem (from Bayesian statistics) with strong (naive) independence assumptions. A more descriptive term for the underlying probability model would be "independent feature model".

In simple terms, a naive Bayes classifier assumes that the presence (or absence) of a particular feature of a class is unrelated to the presence (or absence) of any other feature. For example, a fruit may be considered to be an apple if it is red, round, and about 4" in diameter. Even if these features depend on each other or upon the existence of the other features, a naive Bayes classifier considers all of these properties to independently contribute to the probability that this fruit is an apple.

Depending on the precise nature of the probability model, naive Bayes classifiers can be trained very efficiently in a supervised learning setting. In many practical applications, parameter estimation for naive Bayes models uses the method of maximum likelihood; in other words, one can work with the naive Bayes model without believing in Bayesian probability or using any Bayesian methods.

In spite of their naive design and apparently over-simplified assumptions, naive Bayes classifiers have worked quite well in many complex real-world situations. In 2004, analysis of the Bayesian classification problem has shown that there are some theoretical reasons for the apparently unreasonable efficacy of naive Bayes classifiers.<sup>[1]</sup> Still, a comprehensive comparison with other classification methods in 2006 showed that Bayes classification is outperformed by more current approaches, such as boosted trees or random forests.<sup>[2]</sup>

An advantage of the naive Bayes classifier is that it requires a small amount of training data to estimate the parameters (means and variances of the variables) necessary for classification. Because independent variables are assumed, only the variances of the variables for each class need to be determined and not the entire covariance matrix.

## The naive Bayes probabilistic model

Abstractly, the probability model for a classifier is a conditional model

$$p(C|F_1, \dots, F_n)$$

over a dependent class variable  $C$  with a small number of outcomes or *classes*, conditional on several feature variables  $F_1$  through  $F_n$ . The problem is that if the number of features  $n$  is large or when a feature can take on a large number of values, then basing such a model on probability tables is infeasible. We therefore reformulate the model to make it more tractable.

Using Bayes' theorem, we write

$$p(C|F_1, \dots, F_n) = \frac{p(C) p(F_1, \dots, F_n|C)}{p(F_1, \dots, F_n)}.$$

In plain English the above equation can be written as

$$\text{posterior} = \frac{\text{prior} \times \text{likelihood}}{\text{evidence}}.$$

In practice we are only interested in the numerator of that fraction, since the denominator does not depend on  $C$  and the values of the features  $F_i$  are given, so that the denominator is effectively constant. The numerator is equivalent to the joint probability model

$$p(C, F_1, \dots, F_n)$$

which can be rewritten as follows, using repeated applications of the definition of conditional probability:

$$p(C, F_1, \dots, F_n)$$

$$\begin{aligned}
&= p(C) p(F_1, \dots, F_n | C) \\
&= p(C) p(F_1 | C) p(F_2, \dots, F_n | C, F_1) \\
&= p(C) p(F_1 | C) p(F_2 | C, F_1) p(F_3, \dots, F_n | C, F_1, F_2) \\
&= p(C) p(F_1 | C) p(F_2 | C, F_1) p(F_3 | C, F_1, F_2) p(F_4, \dots, F_n | C, F_1, F_2, F_3) \\
&= p(C) p(F_1 | C) p(F_2 | C, F_1) p(F_3 | C, F_1, F_2) \dots p(F_n | C, F_1, F_2, F_3, \dots, F_{n-1}).
\end{aligned}$$

Now the "naive" conditional independence assumptions come into play: assume that each feature  $F_i$  is conditionally independent of every other feature  $F_j$  for  $j \neq i$ . This means that

$$p(F_i | C, F_j) = p(F_i | C)$$

for  $i \neq j$ , and so the joint model can be expressed as

$$\begin{aligned}
p(C, F_1, \dots, F_n) &= p(C) p(F_1 | C) p(F_2 | C) p(F_3 | C) \dots \\
&= p(C) \prod_{i=1}^n p(F_i | C).
\end{aligned}$$

This means that under the above independence assumptions, the conditional distribution over the class variable  $C$  can be expressed like this:

$$p(C | F_1, \dots, F_n) = \frac{1}{Z} p(C) \prod_{i=1}^n p(F_i | C)$$

where  $Z$  (the evidence) is a scaling factor dependent only on  $F_1, \dots, F_n$ , i.e., a constant if the values of the feature variables are known.

Models of this form are much more manageable, since they factor into a so-called *class prior*  $p(C)$  and independent probability distributions  $p(F_i | C)$ . If there are  $k$  classes and if a model for each  $p(F_i | C = c)$  can be expressed in terms of  $r$  parameters, then the corresponding naive Bayes model has  $(k - 1) + n r k$  parameters. In practice, often  $k = 2$  (binary classification) and  $r = 1$  (Bernoulli variables as features) are common, and so the total number of parameters of the naive Bayes model is  $2n + 1$ , where  $n$  is the number of binary features used for classification and prediction

## Parameter estimation

All model parameters (i.e., class priors and feature probability distributions) can be approximated with relative frequencies from the training set. These are maximum likelihood estimates of the probabilities. A class' prior may be calculated by assuming equiprobable classes (i.e., priors =  $1 / (\text{number of classes})$ ), or by calculating an estimate for the class probability from the training set (i.e., (prior for a given class) = (number of samples in the class) / (total number of samples)). To estimate the parameters for a feature's distribution, one must assume a distribution or generate nonparametric models for the features from the training set. <sup>[3]</sup> If one is dealing with continuous data, a typical assumption is that the continuous values associated with each class are distributed according to a Gaussian distribution.

For example, suppose the training data contains a continuous attribute,  $x$ . We first segment the data by the class, and then compute the mean and variance of  $x$  in each class. Let  $\mu_c$  be the mean of the values in  $x$  associated with class  $c$ , and let  $\sigma_c^2$  be the variance of the values in  $x$  associated with class  $c$ . Then, the probability of some value given a class,  $P(x = v | c)$ , can be computed by plugging  $v$  into the equation for a Normal distribution parameterized by  $\mu_c$  and  $\sigma_c^2$ . That is,

$$P(x = v | c) = \frac{1}{\sqrt{2\pi\sigma_c^2}} e^{-\frac{(v-\mu_c)^2}{2\sigma_c^2}}$$

Another common technique for handling continuous values is to use binning to discretize the values. In general, the distribution method is a better choice if there is a small amount of training data, or if the precise distribution of the data is known. The discretization method tends to do better if there is a large amount of training data because it will

learn to fit the distribution of the data. Since naive Bayes is typically used when a large amount of data is available (as more computationally expensive models can generally achieve better accuracy), the discretization method is generally preferred over the distribution method.

## Sample correction

If a given class and feature value never occur together in the training set then the frequency-based probability estimate will be zero. This is problematic since it will wipe out all information in the other probabilities when they are multiplied. It is therefore often desirable to incorporate a small-sample correction in all probability estimates such that no probability is ever set to be exactly zero.

## Constructing a classifier from the probability model

The discussion so far has derived the independent feature model, that is, the naive Bayes probability model. The naive Bayes classifier combines this model with a decision rule. One common rule is to pick the hypothesis that is most probable; this is known as the *maximum a posteriori* or *MAP* decision rule. The corresponding classifier is the function `classify` defined as follows:

$$\text{classify}(f_1, \dots, f_n) = \underset{c}{\operatorname{argmax}} p(C = c) \prod_{i=1}^n p(F_i = f_i | C = c).$$

## Discussion

Despite the fact that the far-reaching independence assumptions are often inaccurate, the naive Bayes classifier has several properties that make it surprisingly useful in practice. In particular, the decoupling of the class conditional feature distributions means that each distribution can be independently estimated as a one dimensional distribution. This in turn helps to alleviate problems stemming from the curse of dimensionality, such as the need for data sets that scale exponentially with the number of features. Like all probabilistic classifiers under the MAP decision rule, it arrives at the correct classification as long as the correct class is more probable than any other class; hence class probabilities do not have to be estimated very well. In other words, the overall classifier is robust enough to ignore serious deficiencies in its underlying naive probability model. Other reasons for the observed success of the naive Bayes classifier are discussed in the literature cited below.

## Examples

### Sex Classification

Problem: classify whether a given person is a male or a female based on the measured features. The features include height, weight, and foot size.

### Training

Example training set below.

sex	height (feet)	weight (lbs)	foot size(inches)
male	6	180	12
male	5.92 (5'11")	190	11
male	5.58 (5'7")	170	12
male	5.92 (5'11")	165	10
female	5	100	6
female	5.5 (5'6")	150	8
female	5.42 (5'5")	130	7
female	5.75 (5'9")	150	9

The classifier created from the training set using a Gaussian distribution assumption would be:

sex	mean (height)	variance (height)	mean (weight)	variance (weight)	mean (foot size)	variance (foot size)
male	5.855	3.5033e-02	176.25	1.2292e+02	11.25	9.1667e-01
female	5.4175	9.7225e-02	132.5	5.5833e+02	7.5	1.6667e+00

Let's say we have equiprobable classes so  $P(\text{male}) = P(\text{female}) = 0.5$ . There was no identified reason for making this assumption so it may have been a bad idea. If we determine  $P(C)$  based on frequency in the training set, we happen to get the same answer.

### Testing

Below is an example sample that we want to classify as a male or female.

sex	height (feet)	weight (lbs)	foot size(inches)
sample	6	130	8

We want to determine which posterior is greater, male or female.

$$\text{posterior (male)} = P(\text{male}) * P(\text{height}|\text{male}) * P(\text{weight}|\text{male}) * P(\text{foot size}|\text{male}) / \text{evidence}$$

$$\text{posterior (female)} = P(\text{female}) * P(\text{height}|\text{female}) * P(\text{weight}|\text{female}) * P(\text{foot size}|\text{female}) / \text{evidence}$$

The evidence (also termed normalizing constant) may be calculated since the sum of the posteriors equals one.

$$\text{evidence} = P(\text{male}) * P(\text{height}|\text{male}) * P(\text{weight}|\text{male}) * P(\text{foot size}|\text{male}) + P(\text{female}) * P(\text{height}|\text{female}) * P(\text{weight}|\text{female}) * P(\text{foot size}|\text{female})$$

The evidence may be ignored since it is a positive constant (Normal distributions are always positive). Now, let's determine the sex of the sample.

$P(\text{height}|\text{male}) = 1.5789$  (A probability density greater than 1 is OK. It is the area under the bell curve that is equal to 1.)

$$P(\text{weight}|\text{male}) = 5.9881e-06$$

$$P(\text{foot size}|\text{male}) = 1.3112e-3$$

$$\text{posterior numerator (male)} = 6.1984e-09$$

$$P(\text{height}|\text{female}) = 2.2346e-1$$

$$P(\text{weight}|\text{female}) = 1.6789e-2$$

$$P(\text{foot size}|\text{female}) = 2.8669e-1$$

$$\text{posterior numerator (female)} = 5.3778e-04$$

Since posterior numerator (female) > posterior numerator (male), the sample is female.

Try some more test data...

sample, 6, 200, 12

sample, 5.5, 200, 10

sample, 5.5, 120, 6

## Document Classification

Here is a worked example of naive Bayesian classification to the document classification problem. Consider the problem of classifying documents by their content, for example into spam and non-spam E-mails. Imagine that documents are drawn from a number of classes of documents which can be modelled as sets of words where the (independent) probability that the  $i$ -th word of a given document occurs in a document from class  $C$  can be written as

$$p(w_i|C)$$

(For this treatment, we simplify things further by assuming that words are randomly distributed in the document - that is, words are not dependent on the length of the document, position within the document with relation to other words, or other document-context.)

Then the probability that a given document  $D$  contains all of the words  $w_i$ , given a class  $C$ , is

$$p(D|C) = \prod_i p(w_i|C)$$

The question that we desire to answer is: "what is the probability that a given document  $D$  belongs to a given class  $C$ ?" In other words, what is  $p(C|D)$ ?

Now by definition

$$p(D|C) = \frac{p(D \cap C)}{p(C)}$$

and

$$p(C|D) = \frac{p(D \cap C)}{p(D)}$$

Bayes' theorem manipulates these into a statement of probability in terms of likelihood.

$$p(C|D) = \frac{p(C)}{p(D)} p(D|C)$$

Assume for the moment that there are only two mutually exclusive classes,  $S$  and  $\neg S$  (e.g. spam and not spam), such that every element (email) is in either one or the other;

$$p(D|S) = \prod_i p(w_i|S)$$

and

$$p(D|\neg S) = \prod_i p(w_i|\neg S)$$

Using the Bayesian result above, we can write:

$$p(S|D) = \frac{p(S)}{p(D)} \prod_i p(w_i|S)$$

$$p(\neg S|D) = \frac{p(\neg S)}{p(D)} \prod_i p(w_i|\neg S)$$

Dividing one by the other gives:

$$\frac{p(S|D)}{p(\neg S|D)} = \frac{p(S) \prod_i p(w_i|S)}{p(\neg S) \prod_i p(w_i|\neg S)}$$

Which can be re-factored as:

$$\frac{p(S|D)}{p(\neg S|D)} = \frac{p(S)}{p(\neg S)} \prod_i \frac{p(w_i|S)}{p(w_i|\neg S)}$$

Thus, the probability ratio  $p(S|D) / p(\neg S|D)$  can be expressed in terms of a series of likelihood ratios. The actual probability  $p(S|D)$  can be easily computed from  $\log(p(S|D) / p(\neg S|D))$  based on the observation that  $p(S|D) + p(\neg S|D) = 1$ .

Taking the logarithm of all these ratios, we have:

$$\ln \frac{p(S|D)}{p(\neg S|D)} = \ln \frac{p(S)}{p(\neg S)} + \sum_i \ln \frac{p(w_i|S)}{p(w_i|\neg S)}$$

(This technique of "log-likelihood ratios" is a common technique in statistics. In the case of two mutually exclusive alternatives (such as this example), the conversion of a log-likelihood ratio to a probability takes the form of a sigmoid curve: see logit for details.)

Finally, the document can be classified as follows. It is spam if  $p(S|D) > p(\neg S|D)$  (i.e. that  $\ln \frac{p(S|D)}{p(\neg S|D)} > 0$ ), otherwise it is not spam.

## References

- [1] Harry Zhang "The Optimality of Naive Bayes". FLAIRS2004 conference. (*available online: PDF* (<http://www.cs.unb.ca/profs/hzhang/publications/FLAIRS04ZhangH.pdf>))
- [2] Caruana, R. and Niculescu-Mizil, A.: "An empirical comparison of supervised learning algorithms". Proceedings of the 23rd international conference on Machine learning, 2006. (*available online PDF* (<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.122.5901&rep=rep1&type=pdf>))
- [3] George H. John and Pat Langley (1995). Estimating Continuous Distributions in Bayesian Classifiers. Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence. pp. 338-345. Morgan Kaufmann, San Mateo.

## External links

- Benchmark results of Naive Bayes implementations (<http://tunedit.org/results?d=UCI/&a=bayes>)
- Hierarchical Naive Bayes Classifiers for uncertain data (<http://www.biomedcentral.com/1471-2105/7/514>) (an extension of the Naive Bayes classifier).
- winnowTag (<http://winnowtag.org>) Create tags that run Bayesian classifiers on over 1/2 million items updated from 7,500 feeds.
- Online application of a Naive Bayes classifier (<http://www.convo.co.uk/x02/>) (emotion modelling), with a full explanation.
- dANN: Naive Classifier - Explanation and Example ([http://wiki.syncleus.com/index.php/DANN:Naive\\_Classifier](http://wiki.syncleus.com/index.php/DANN:Naive_Classifier))
- BayesNews - Bayesian RSS Reader (<http://www.bayesnews.xpg.com.br/>) (useful for personal news clipping).
- Naive Bayes in PMML (<http://www.dmg.org/v4-0/NaiveBayes.html>)

### Software

- IMSL Numerical Libraries Collections of math and statistical algorithms available in C/C++, Fortran, Java and C#/.NET. Data mining routines in the IMSL Libraries include a Naive Bayes classifier.
- Orange, a free data mining software suite, module orngBayes (<http://www.ailab.si/orange/doc/modules/orngBayes.htm>)
- Winnow content recommendation (<http://doc.winnowtag.org/open-source>) Open source Naive Bayes text classifier works with very small training and unbalanced training sets. High performance, C, any Unix.

Source URL: [http://www.en.wikipedia.org/wiki/Naive\\_Bayes\\_classifier](http://www.en.wikipedia.org/wiki/Naive_Bayes_classifier)  
 Saylor URL: <http://www.saylor.org/courses/cs408>

Attributed to [Wikipedia]



- Naive Bayes implementation in Visual Basic ([http://paul.luminos.nl/documents/show\\_document.php?d=198](http://paul.luminos.nl/documents/show_document.php?d=198)) (includes executable and source code).
- An interactive Microsoft Excel spreadsheet Naive Bayes implementation ([http://downloads.sourceforge.net/naivebayesclass/NaiveBayesDemo.xls?use\\_mirror=osdn](http://downloads.sourceforge.net/naivebayesclass/NaiveBayesDemo.xls?use_mirror=osdn)) using VBA (requires enabled macros) with viewable source code.
- jBNC - Bayesian Network Classifier Toolbox (<http://jbnc.sourceforge.net/>)
- POPFile (<http://popfile.sourceforge.net/>) Perl-based email proxy system classifies email into user-defined "buckets", including spam.
- Statistical Pattern Recognition Toolbox for Matlab (<http://cmp.felk.cvut.cz/cmp/software/stprtool/>).
- Document Classification Using Naive Bayes Classifier with Perl ([http://www.lwebzem.com/cgi-bin/res/naive\\_bayes\\_tm\\_classifier.cgi](http://www.lwebzem.com/cgi-bin/res/naive_bayes_tm_classifier.cgi)).
- sux0r (<http://www.sux0r.org>) An Open Source Content management system with a focus on Naive Bayesian categorization and probabilistic content.
- ifile (<http://people.csail.mit.edu/jrennie/ifile/>) - the first freely available (Naive) Bayesian mail/spam filter
- NClassifier (<http://nclassifier.sourceforge.net/>) - NClassifier is a .NET library that supports text classification and text summarization. It is a port of the Nick Lothian's popular Java text classification engine, Classifier4J.
- Classifier4J (<http://classifier4j.sourceforge.net/>) - Classifier4J is a Java library designed to do text classification. It comes with an implementation of a Bayesian classifier, and now has some other features, including a text summary facility.
- MALLET (<http://mallet.cs.umass.edu/>) - A Java package for document classification and other natural language processing tasks.
- nBayes (<http://nbayes.codeplex.com>) - nBayes is an open source .NET library written in C#
- Apache Mahout (<http://mahout.apache.org>) - Machine learning package offered by Apache Open Source
- Weka (<http://www.cs.waikato.ac.nz/ml/weka/>) - Popular machine learning package

#### Publications

- Domingos, Pedro & Michael Pazzani (1997) "On the optimality of the simple Bayesian classifier under zero-one loss". *Machine Learning*, 29:103–137. (also online at *CiteSeer* (<http://citeseer.ist.psu.edu/>): (<http://citeseer.ist.psu.edu/domingos97optimality.html>))
- Rish, Irina. (2001). "An empirical study of the naive Bayes classifier". IJCAI 2001 Workshop on Empirical Methods in Artificial Intelligence. (available online: PDF (<http://www.research.ibm.com/people/r/rish/papers/RC22230.pdf>), PostScript (<http://www.research.ibm.com/people/r/rish/papers/ijcai-ws.ps>))
- Hand, DJ, & Yu, K. (2001). "Idiot's Bayes - not so stupid after all?" *International Statistical Review*. Vol 69 part 3, pages 385-399. ISSN 0306-7734.
- Webb, G. I., J. Boughton, and Z. Wang (2005). Not So Naive Bayes: Aggregating One-Dependence Estimators (<http://www.springerlink.com/content/u8w306673m1p866k/>). *Machine Learning* 58(1). Netherlands: Springer, pages 5-24.
- Mozina M, Demsar J, Kattan M, & Zupan B. (2004). "Nomograms for Visualization of Naive Bayesian Classifier". In Proc. of PKDD-2004, pages 337-348. (available online: PDF (<http://www.ailab.si/blaz/papers/2004-PKDD.pdf>))
- Maron, M. E. (1961). "Automatic Indexing: An Experimental Inquiry." *Journal of the ACM (JACM)* 8(3):404–417. (available online: PDF (<http://delivery.acm.org/10.1145/330000/321084/p404-maron.pdf?key1=321084&key2=9636178211&coll=GUIDE&dl=ACM&CFID=56729577&CFTOKEN=37855803>))
- Minsky, M. (1961). "Steps toward Artificial Intelligence." *Proceedings of the IRE* 49(1):8-30.
- McCallum, A. and Nigam K. "A Comparison of Event Models for Naive Bayes Text Classification". In AAAI/ICML-98 Workshop on Learning for Text Categorization, pp. 41–48. Technical Report WS-98-05. AAAI Press. 1998. (available online: PDF (<http://www.kamalnigam.com/papers/multinomial-aaaiws98.pdf>))

- Rennie J, Shih L, Teevan J, and Karger D. Tackling The Poor Assumptions of Naive Bayes Classifiers. In Proceedings of the Twentieth International Conference on Machine Learning (ICML). 2003. (*available online: PDF* (<http://people.csail.mit.edu/~jrennie/papers/icml03-nb.pdf>))

# Article Sources and Contributors

**Naive Bayes classifier** *Source:* <http://en.wikipedia.org/w/index.php?oldid=411315852> *Contributors:* Akella, AllenDowney, Anders gorm, Anna Lincoln, Arauzo, Arichnad, Bewildebeast, Bkell, BlueNovember, Bovineone, Btyner, Caesura, Cagri, Calimo, Can't sleep, clown will eat me, Cancan101, Classifier1234, Coffee2theorems, ComodiCast, Cyp, David Eppstein, Dchwalisz, Ddxc, Den fjättrade ankan, Dianegarey, Dkemper, Doobliebop, Dstanfor, EverGreg, Feady2007, Gene s, Geoffrey I Webb, Gifflite, Headlessplatter, HebrewHammerTime, Herlocker, Hike395, Icestorm815, Intgr, Jason Davies, JimD, Jklin, John Vandenberg, Johnnyw, Joseagonzalez, Jrennie, Justin W Smith, KKramer, Karada, Karipuf, Kotsiantis, Macrakis, Maghnus, MarkSweep, Mat1971, Mbusux, Mebden, Melcombe, Memming, Michael Hardy, Mitar, MrOllie, Mwojnars, Neilc, Neshatian, Newtman, Oleg Alexandrov, Olivier, OrangeDog, Orphan Wiki, PerVognsen, Peterjoel, Proffviktor, Prolog, RJASE1, RPHv, RedWolf, Rich Farmbrough, Rickyphyllis, Ringger, Rjwilmsi, Saurabh911, Sderose, Splatty, Stimpny, Sunsetsky, Svourdroculed, The Almighty Bob, The Anome, Tobym, Tommy2010, Toreau, Tremilux, Trevor MacInnis, Troos, Twexcom, User A1, Vera Rita, WMod-NS, WhisperToMe, Wile E. Heresiarch, Wingiii, Zeno Gantner, 188 anonymous edits

## License

---

Creative Commons Attribution-Share Alike 3.0 Unported  
<http://creativecommons.org/licenses/by-sa/3.0/>