

DSC 40A

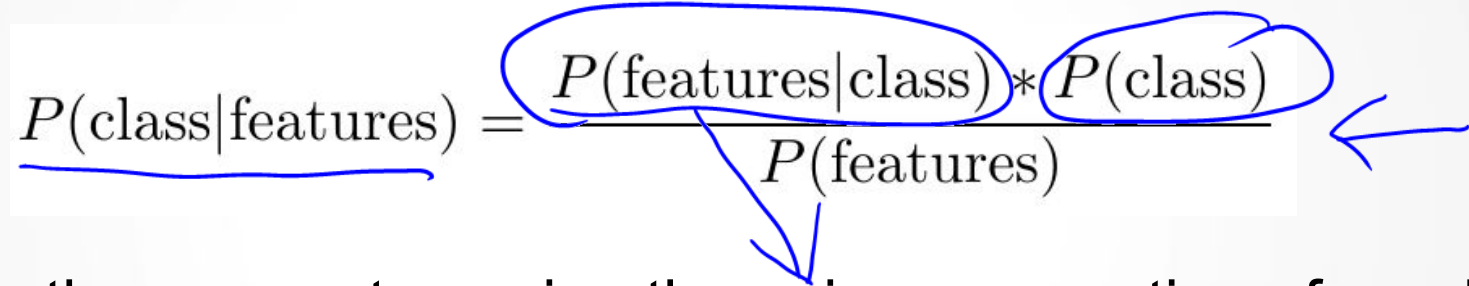
Theoretical Foundations of Data Science I

Last Time

- Classification is the problem of predicting a categorical variable.
- The Naive Bayes algorithm gives a strategy for classifying data according to its features.
- It is naive because it relies on an assumption of conditional independence of the features, for a given class.

Naive Bayes Algorithm

- Bayes' Theorem shows how to calculate $P(\text{class} | \text{features})$.

$$\underline{P(\text{class} | \text{features})} = \frac{P(\text{features} | \text{class}) * P(\text{class})}{P(\text{features})}$$



- Rewrite the numerator, using the naive assumption of conditional independence of features given the class.
- Estimate each term in the numerator based on the training data.
- Select class based on whichever has the larger numerator.

In This Video

- How can we use naive Bayes to classify text?

Bayes' Theorem for Text Classification

Text classification problems include:

- sentiment analysis
 - positive and negative customer reviews
- determining genre
 - news articles, blog posts, etc.
- email foldering
 - promotions tab in Gmail
-  **spam filtering**
 - **separating spam from ham (good, non-spam email)**

Features

Represent an email as a vector or array of features

$$(\underline{x}_1, \underline{x}_2, \underline{x}_3, \dots, \underline{x}_n)$$

where i is an index into a dictionary of n possible words, and

$x_i = 1$ if word i is present in the email
 $x_i = 0$ otherwise

} indicator
variable

Features

Called the “bag of words” model:

Ignores location of words within the email, and the frequency of words

Example:

dictionary

can

do

each

something

understand

when

“I do understand if
something comes up and
you can not do this
today.”

(1, 1, 0, 1, 1, 0)



usually $n = 10,000$ to
50,000 words in practice

Naive Bayes Spam Classifier

$$P(\text{class}|\text{features}) = \frac{P(\text{features}|\text{class}) * P(\text{class})}{P(\text{features})}$$

To classify an email, we will use Bayes' Theorem to calculate the probability of it belonging to each class:

$P(\text{spam} | \text{features})$ and $P(\text{ham} | \text{features})$

Then choose the class according to the **larger** of these two probabilities.

Naive Bayes Spam Classifier

$$P(\text{class}|\text{features}) = \frac{P(\text{features}|\text{class}) * P(\text{class})}{P(\text{features})}$$

Observe: the formulas for $P(\text{spam} | \text{features})$ and $P(\text{ham} | \text{features})$ have the same denominator, $P(\text{features})$.

We can find the larger of the two probabilities by just comparing numerators.

$P(\text{features} | \text{spam}) * P(\text{spam})$ vs. $P(\text{features} | \text{ham}) * P(\text{ham})$

Naive Bayes Spam Classifier

$$P(\text{class}|\text{features}) = \frac{P(\text{features}|\text{class}) * P(\text{class})}{P(\text{features})}$$

To use Bayes' Theorem, need to determine four quantities:

- i. $P(\text{features} \mid \text{spam})$
- ii. $P(\text{features} \mid \text{ham})$
- iii. $P(\text{spam})$
- iv. $P(\text{ham})$

Which of these probabilities should add to 1?

- A. i, ii
- B. iii, iv
- C. both A and B
- D. neither A nor B

Estimating Parameters with Training Data

parameter:

$P(\text{spam})$

estimate:

$\frac{\text{\# spam emails in training set}}{\text{size of training set}}$

parameter:

$P(\text{ham})$

estimate:

$\frac{\text{\# ham emails in training set}}{\text{size of training set}}$

$P(\text{features} \mid \text{spam})$

$P(\text{features} \mid \text{ham})$



harder to estimate

Assumption of Conditional Independence

To estimate $P(\text{features} \mid \text{spam})$ and $P(\text{features} \mid \text{ham})$, we assume that the probability of a word appearing in an email of a given class is not affected by other words in the email.

$$P(\underbrace{x_1=0}_{\text{blue}} \text{ and } \underbrace{x_2=1}_{\text{blue}} \text{ and } \underbrace{x_3=1}_{\text{blue}} \text{ and } \dots \mid \text{spam}) = \leftarrow \text{assumed equal}$$

$$\underbrace{P(x_1=0 \mid \text{spam})}_{\text{blue}} * \underbrace{P(x_2=1 \mid \text{spam})}_{\text{blue}} * \underbrace{P(x_3=1 \mid \text{spam})}_{\text{blue}} * \dots$$

Is this a reasonable assumption? NO, words have meaning, and certain words tend to appear together

Estimating Parameters with Training Data

$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots | \text{spam}) =$ ← **assumed equal**

$$P(x_1=0 | \text{spam}) * P(x_2=1 | \text{spam}) * P(x_3=1 | \text{spam}) * \dots$$

parameter: $P(x_1=0 | \text{spam})$

estimate: $\frac{\text{\# spam emails in training set not containing the first word in the dictionary}}{\text{\# spam emails in training set}}$

Estimating Parameters with Training Data

$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots | \text{spam}) =$ ← **assumed equal**

$$P(x_1=0 | \text{spam}) * P(x_2=1 | \text{spam}) * P(x_3=1 | \text{spam}) * \dots$$

parameter: $P(x_2=1 | \text{spam})$

estimate: $\frac{\text{\# spam emails in training set containing the second word in the dictionary}}{\text{\# spam emails in training set}}$

Estimating Parameters with Training Data

$$P(\underbrace{x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots}_{\text{assumed equal}} \mid \text{spam}) = \leftarrow$$

$$P(x_1=0 \mid \text{spam}) * P(x_2=1 \mid \text{spam}) * P(x_3=1 \mid \text{spam}) * \dots$$

parameter: $P(x_3=1 \mid \text{spam})$

estimate: $\frac{\text{\# spam emails in training set containing the third word in the dictionary}}{\text{\# spam emails in training set}}$

Can term-by-term estimate $P(\text{features} \mid \text{class})$

Naive Bayes Spam Classifier: Recap

Bayes' Theorem shows how to calculate $P(\text{spam} \mid \text{features})$ and $P(\text{ham} \mid \text{features})$.

$$P(\text{class} \mid \text{features}) = \frac{P(\text{features} \mid \text{class}) * P(\text{class})}{P(\text{features})}$$

Rewrite the numerator, using the naive assumption of conditional independence of words given the class.

Estimate each term in the numerator based on the training data.

Select class based on whichever has the larger numerator.

Modifications and Extensions

- features are pairs (or longer sequences) of words rather than individual words
 - better captures dependencies between words
 - less naive
 - much bigger feature space
 - n words $\rightarrow n^2$ pairs of words
- features are the number of occurrences of each word
 - captures low-frequency vs. high-frequency words
- smoothing
 - better handling of previously unseen words



Smoothing

$$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots | \text{spam}) = \\ P(x_1=0 | \text{spam}) * P(x_2=1 | \text{spam}) * P(x_3=1 | \text{spam}) * \dots$$

Dictionary

1. a
2. aardvark
3. abacus
4. abandon
5. abate
- ...

Suppose you are classifying an email containing the word “abacus,” which does not appear in any emails in your training data. For this new email’s features, what is **P(features | spam)** according to a naive Bayes classifier?

- A. $P(\text{features} | \text{spam}) = \text{undefined}$
- B. $P(\text{features} | \text{spam}) = 0$
- C. $P(\text{features} | \text{spam}) = 1/n$
- D. $P(\text{features} | \text{spam}) = 1$

Smoothing

$$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots | \text{ham}) = \frac{0}{\# \text{ham}} P(x_1=0 | \text{ham}) P(x_2=1 | \text{ham}) P(x_3=1 | \text{ham}) \dots$$

Dictionary

1. a
2. aardvark
3. abacus
4. abandon
5. abate
- ...

Suppose you are classifying an email containing the word “abacus,” which does not appear in any emails in your training data. For this new email’s features, what is $P(\text{features} | \text{ham})$ according to a naive Bayes classifier?

- A. $P(\text{features} | \text{ham}) = \text{undefined}$
- B. $P(\text{features} | \text{ham}) = 0$
- C. $P(\text{features} | \text{ham}) = 1/n$
- D. $P(\text{features} | \text{ham}) = 1$

Smoothing

$P(\text{features} \mid \text{spam}) = 0$ and $P(\text{features} \mid \text{ham}) = 0$

Tiebreaker: randomly select one of the classes?

Smoothing

$P(\text{features} \mid \text{spam}) = 0$ and $P(\text{features} \mid \text{ham}) = 0$

Tiebreaker: randomly select one of the classes?

Better solution: make sure probabilities can't be zero

Key idea: just because you've never seen something happen doesn't mean it's impossible

Estimating Parameters with Training Data

$\#spam \rightarrow \#spam + 1$
 $\#ham \rightarrow \#ham + 1$

parameter:

P(spam)

estimate:

+

parameter:

P(ham)

estimate:

||

|

Without Smoothing

$$\frac{\#spam}{\#spam + \#ham}$$

+

$$\frac{\#ham}{\#spam + \#ham}$$

$$\frac{\#spam + 1}{\#spam + \#ham}$$

With Smoothing

$$\frac{\#spam + 1}{\#spam + 1 + \#ham + 1}$$

+

$$\frac{\#ham + 1}{\#spam + 1 + \#ham + 1}$$

$$\frac{\#spam + 1 + \#ham + 1}{\#spam + 1 + \#ham + 1}$$

Estimating Parameters with Training Data

parameter: $P(x_i=1 \mid \text{spam})$

#spam cont. word $i \rightarrow \text{add } 1$
#spam not cont. word $i \rightarrow \text{add } 1$

estimate:

$$\frac{\text{\#spam containing word } i}{\text{\#spam containing word } i + \text{\#spam not containing word } i}$$

Without
Smoothing

$$\frac{(\text{\#spam containing word } i) + 1}{(\text{\#spam containing word } i) + 1 + (\text{\#spam not containing word } i) + 1}$$

With
Smoothing

Similarly for other parameters $P(x_i=0 \mid \text{spam})$, $P(x_i=1 \mid \text{ham})$, $P(x_i=0 \mid \text{ham})$.

Smoothing

$$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots | \text{ham}) =$$

$$P(x_1=0 | \text{ham}) * P(x_2=1 | \text{ham}) * P(x_3=1 | \text{ham}) * \dots$$

$$\frac{\# \text{ham containing abacus} + 1}{\# \text{ham containing abacus} + \# \text{ham NOT containing abacus} + 1} = \frac{0 + 1}{\# \text{ham} + 2}$$

Dictionary

1. a
2. aardvark
3. abacus
4. abandon
5. abate
- ...

Suppose you are classifying an email containing the word “abacus,” which does not appear in any emails in your training data. What is $P(x_3 = 1 | \text{ham})$ according to a naive Bayes classifier **with smoothing**?

- A. $P(x_3 = 1 | \text{ham}) = 0$
- B. $P(x_3 = 1 | \text{ham}) = 1/2$
- C. $P(x_3 = 1 | \text{ham}) = 1/(\text{total \#ham} + 1)$
- D. $P(x_3 = 1 | \text{ham}) = 1/(\text{total \#ham} + 2)$
- E. $P(x_3 = 1 | \text{ham}) = 1/(\text{total \#ham} + \text{total \#spam} + 2)$

Smoothing

Suppose your training set includes only six emails, all of which are ham. For a new email, how will we estimate $P(\text{ham})$

without smoothing?

~~A. $P(\text{ham}) = 0$~~

~~B. $P(\text{ham}) = 0$~~

C. $P(\text{ham}) = 1$

D. $P(\text{ham}) = 1$

E. $P(\text{ham}) = 1$

with smoothing?

$P(\text{ham}) = 1/8$

$P(\text{ham}) = 6/7$

$P(\text{ham}) = 1/8$

$P(\text{ham}) = 6/7$

$P(\text{ham}) = 7/8$

↓

$$\frac{\# \text{ham} + 1}{\# \text{ham} + 1 + \# \text{spam} + 1} = \frac{6 + 1}{6 + 1 + 0 + 1} = \frac{7}{8}$$

$$\frac{10,000 + 1}{10,000 + 2}$$

$$7/8$$

Summary

- The Naive Bayes algorithm is useful for text classification.
- The bag of words model treats each word in a large dictionary as a feature.
- Smoothing is one modification that allows for better predictions when there are words that have never been seen before.