

DSC 40A

Theoretical Foundations of Data Science I

Announcements

- Homework 7 released today and due 12/6.
- Next Friday review session for final exam. 4-6pm

‡ Edstem

* Please fill out SETs!

Question

Answer at q.dsc40a.com

Remember, you can always ask questions at
q.dsc40a.com!

If the direct link doesn't work, click the "Lecture Questions" link in the top right corner of dsc40a.com.

Agenda

- Naïve Bayes Classifier
- Text Classifier

MODIFIED BAYES' THEOREM:

$$P(H|X) = P(H) \times \left(1 + P(C) \times \left(\frac{P(x|H)}{P(x)} - 1 \right) \right)$$

H: HYPOTHESIS

x: OBSERVATION

P(H): PRIOR PROBABILITY THAT H IS TRUE

P(x): PRIOR PROBABILITY OF OBSERVING x

P(C): PROBABILITY THAT YOU'RE USING
BAYESIAN STATISTICS CORRECTLY

Source: xkcd

Avocado Ripeness

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a firm green-black Zutano avocado. Based on this data, would you predict that your avocado is ripe or unripe?

Problem: We have not seen an avocado with all these features. Both probabilities will be undefined.

$P(\text{ripe} \mid \text{firm, green-black, Zutano})$

$P(\text{unripe} \mid \text{firm, green-black, Zutano})$

Avocado Ripeness

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a firm green-black Zutano avocado. Based on this data, would you predict that your avocado is ripe or unripe?

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

Solution: Use Bayes' Theorem, plus a simplifying assumption, to calculate the two numerators.

Avocado Ripeness

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a firm green-black Zutano avocado. Based on this data, would you predict that your avocado is ripe or unripe?

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

Simplifying assumption: Within a given class, the features are independent.

$$P(\text{firm, green-black, Zutano} | \text{ripe}) = P(\text{firm} | \text{ripe}) * P(\text{green-black} | \text{ripe}) * P(\text{Zutano} | \text{ripe})$$

Conditional Independence

- Recall that A and B are independent if

$$P(A \text{ and } B) = P(A) * P(B)$$

- A and B are conditionally independent given C if

$$P(\underbrace{(A \text{ and } B)}_{\text{features}} | \underbrace{C}_{\text{class}}) = P(\underbrace{A}_{\text{feat 1}} | \underbrace{C}_{\text{class}}) * P(\underbrace{B}_{\text{feat 2}} | \underbrace{C}_{\text{class}})$$

- Given that C occurs, this says that A and B are independent of one another.

Avocado Ripeness

$$P(\text{ripe}) = \frac{7}{11}$$

$$P(\text{ripe} | \text{features}) \propto \frac{7}{11} \cdot \frac{6}{7^3} = \frac{6}{49 \cdot 11} = \frac{6}{539}$$

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a **firm** **green-black** **Zutano** avocado. Based on this data, would you predict that your avocado is ripe or unripe?

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

$$P(\text{color, softness, variety} | \text{ripeness}) = P(\text{color} | \text{ripeness}) * P(\text{softness} | \text{ripeness}) * P(\text{variety} | \text{ripeness})$$

$$P(\text{greenblack} | \text{ripe}) \quad P(\text{firm} | \text{ripe}) \quad P(\text{zutano} | \text{ripe})$$

$$\frac{1}{7}$$

$$\frac{3}{7}$$

$$\frac{2}{7} = \frac{6}{7^3}$$

Avocado Ripeness

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a firm green-black Zutano avocado. Based on this data, would you predict that your avocado is ripe or unripe?

Assuming conditional independence of features given the class, calculate $P(\text{firm, green-black, Zutano} \mid \text{unripe})$.

~~A.~~ 0

B. 1/4

C. 3/16

~~D.~~ $1 - (1/7 * 3/7 * 2/7)$

$$P(A|B) \neq 1 - P(A|\bar{B}) \quad \times$$

$$P(A|B) = 1 - P(\bar{A}|B) \quad \checkmark$$

$$1 - P(\text{greenblack, firm, zutano} | \text{ripe})$$

Avocado Ripeness

Color	Softness	Variety	Ripeness
bright green	firm	Zutano	unripe
green-black	medium	Hass	ripe
purple-black	firm	Hass	ripe
green-black	medium	Hass	unripe
purple-black	soft	Hass	ripe
bright green	firm	Zutano	unripe
green-black	soft	Zutano	ripe
purple-black	soft	Hass	ripe
green-black	soft	Zutano	ripe
green-black	firm	Hass	unripe
purple-black	medium	Hass	ripe

You have a firm green-black Zutano avocado. Based on this data, would you predict that your avocado is ripe or unripe?

$$P(\text{unripe}|\text{feats}) \propto \frac{3}{16} \cdot \frac{4}{11} = \frac{3}{44} = \frac{6}{88} > \frac{6}{539}$$

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

↳ cond. ind. assump.

$$P(\text{firm}|\text{unripe}) \cdot P(\text{greenblack}|\text{unripe}) \cdot P(\text{Zutano}|\text{unripe})$$

$$= \frac{3}{4} \cdot \frac{2}{4} \cdot \frac{2}{4} = \frac{3}{16}$$

$$P(\text{unripe}) = \frac{4}{11}$$

Naïve Bayes Algorithm

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

$$P(\text{class}|\text{features}) = \frac{P(\text{features}|\text{class}) * P(\text{class})}{P(\text{features})} \quad \leftarrow \text{don't need to calculate}$$

- Rewrite the numerator, using the naïve 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.

Naive Bayes Classifier for Text Classification



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

$$(x_1, x_2, x_3, \dots, x_n)$$

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

“ I do understand if something
comes up and you cannot do
this today ”

- 1 can
- 2 do
- 3 each
- 4 something
- 5 understand
- 6 when

(0 1 0 1 1 0)

↑ no “can” ↑ appeared twice



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

Naïve Bayes Spam Classifier

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

ham spam bag of words

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

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

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

Naïve Bayes Spam Classifier

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

Shortcut

← don't calculate

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

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

$$P(A|B) + P(\bar{A}|B) = \frac{P(A \cap B)}{P(B)} + \frac{P(\bar{A} \cap B)}{P(B)} = \frac{P(S \cap B)}{P(B)} = \frac{P(B)}{P(B)} = 1$$

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

$$P(A|B) + P(A|\bar{B}) \neq 1$$

i. $P(\text{features} | \text{spam})$

ii. $P(\text{features} | \text{ham})$

iii. $P(\text{spam})$

iv. $P(\text{ham})$

$$P(B) + P(\bar{B}) = 1$$

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:

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

counting

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.

$$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \dots \mid \text{spam}) = \leftarrow \text{assumed equal}$$
$$P(x_1, x_2, x_3 \mid \text{class}) = P(x_1 \mid \text{class}) \cdot P(x_2 \mid \text{class}) \cdot P(x_3 \mid \text{class})$$
$$P(x_1=0 \mid \text{spam}) * P(x_2=1 \mid \text{spam}) * P(x_3=1 \mid \text{spam}) * \dots$$

Is this a reasonable assumption? NO! words appear together based on the meaning of sentence \Rightarrow words are correlated

Estimating Parameters with Training Data

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

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

parameter:

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

*first word wasn't in the email
in the dictionary*

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

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

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

↳ word 2 in dict, did appear in email

parameter:

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

estimate:

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

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

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

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

Naïve 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.

*by counting per feat
conditioned on
class*

Modifications and Extensions

- features are pairs (or longer sequences) of words rather than individual words
 - better captures dependencies between words
 - less naïve
 - 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

(*n*-grams)

relative importance

Smoothing

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


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

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

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 naïve 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

Without Smoothing

With Smoothing

parameter:

P(spam)

estimate:

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

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

+

+

+

parameter:

P(ham)

estimate:

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

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

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

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

=
1

Estimating Parameters with Training Data

spam containing word $i \rightarrow$ # spam containing word $i+1$

parameter:

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

estimate:

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

Without Smoothing

$$\frac{(\text{\#spam containing word } i) + 1}{(\text{\#spam containing word } i) + 1 + (\text{\#spam not containing word } i) + 1}$$
$$= \frac{\text{\#spam containing word } i + 1}{\text{\#spam} + 2}$$

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

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

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 naïve 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? with smoothing?

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

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

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

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

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

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

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

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

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

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

$$P(\text{ham}) = \frac{\# \text{ham} + 1}{\# \text{ham} + 1 + \# \text{spam} + 1} = \frac{\# \text{ham} + 1}{\# \text{ham} + 2} = 1 \quad \frac{7}{8} < \frac{10k + 1}{10k + 2} \approx \frac{10k}{10k} = 1$$

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? with smoothing?

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

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

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

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

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

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

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

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

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

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

$$P(\text{spam}) = \frac{0+1}{\#_{\text{ham}}+2} \Rightarrow \frac{1}{8} \quad \frac{1}{8} + \frac{7}{8} = 1$$

Is it still true that $P(\text{ham}) + P(\text{spam}) = 1$ with smoothing?



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.