DSC 40A

Theoretical Foundations of Data Science I

* fill out SET
* HW due Wednesday
* final logistics will be sent Monday

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.

Last Time

- Classification is the problem of predicting a categorical variable.
- The Naïve Bayes algorithm gives a strategy for classifying data according to its features.
- It is naïve 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(class | features).

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

- 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.

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?

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

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

prior $\frac{7}{11}$

conditional independence

$P(\text{firm}|\text{ripe}) \cdot P(\text{gb}|\text{ripe}) \cdot P(\text{zutano}|\text{ripe})$

$= \frac{1}{7} \quad \frac{3}{7} \quad \frac{2}{7}$

$P(\text{ripe}|\text{feats}) \propto \frac{1 \cdot 3 \cdot 2}{7 \cdot 7 \cdot 7} \cdot \frac{7}{11} = \frac{6}{539}$

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?

common mistake

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

Assuming conditional independence of features given the class, calculate P(firm, green-black, Zutano | unripe).

A. 0
B. 1/4
C. 3/16
D. 1 - (1/7*3/7*2/7)

$P(\text{firm}, \text{gb}, \text{zutano} | \text{ripe})$

$\neq 1 - P(\text{firm gb zutano} | \text{ripe})$

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?

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

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

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

$P(\text{unripe}) = \frac{4}{11} \Rightarrow \text{numerator}: \frac{3}{16} \cdot \frac{4}{11} = \frac{3}{44} > \frac{6}{539}$

Naïve Bayes Algorithm

- Bayes' Theorem shows how to calculate P(class | features).

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

- 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.

Agenda

- How can we use naïve Bayes to classify text?

Naïve 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

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

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

100k – 500k

Example:

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

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

(0 1 0 1 1 0)
1 2 3 4 5 6

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

Naïve Bayes Spam Classifier

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

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

$$P(\text{spam} \mid \text{features}) \overset{?}{<} \text{and} \overset{?}{>} P(\text{ham} \mid \text{features})$$

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

Naïve Bayes Spam Classifier

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

Observe: the formulas for P(spam | features) and P(ham | features) have the same denominator, P(features).

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

P(features | spam)**P*(spam) vs. *P*(features | ham)**P*(ham)

Naive Bayes Spam Classifier

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

B A

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

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

$P(A|B) + P(A|\bar{B}) \neq 1$ (not necessary)

i. *P*(features | spam)

ii. *P*(features | ham)

iii. *P*(spam)

iv. *P*(ham)

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

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

$\Rightarrow$ counting

parameter: $P(\text{ham})$ **estimate:** $\frac{\#\text{ ham emails in training set}}{\text{size of training set}}$

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

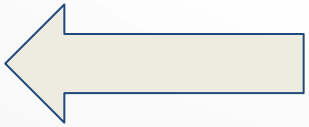

harder to estimate

Assumption of Conditional Independence

To estimate P(features | spam) and P(features | 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.

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

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

Is this a reasonable assumption? No, this is naive. Even given ham/spam we have word co-occurences with different probabilities.

Estimating Parameters with Training Data

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

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

counting

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

word 1 doesn't appear in 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

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

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

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

prob. of word 2 appearing in spam email

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

$P(x_1=0 \text{ and } x_2=1 \text{ and } x_3=1 \text{ and } \ldots \mid \text{spam}) =$ ← assumed equal

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

parameter: $P(x_3=1 \mid \text{spam})$ word 3 appeared in spam

estimate:

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

Can term-by-term estimate P(features|class)

Naïve Bayes Spam Classifier: Recap

Bayes' Theorem shows how to calculate P(spam | features) and P(ham | features).

conditional indep. prior

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

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

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

Select class based on whichever has the larger numerator.

Modifications and Extensions

(n-grams)

- 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 → 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

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

$\frac{0}{\#\text{spam}} = 0$

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(features | spam) = undefined
B. P(features | spam) = 0
C. P(features | spam) = 1/n
D. P(features | spam) = 1

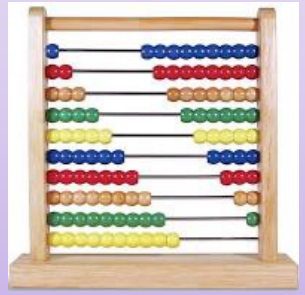

Smoothing

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

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

0

#ham

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 | ham)** according to a naïve Bayes classifier?

A. P(features | ham) = undefined
B. P(features | ham) = 0
C. P(features | ham) = 1/n
D. P(features | ham) = 1

Smoothing

If any $P(\text{feat}_i \mid \text{class}) = 0 \Rightarrow P(\text{features} \mid \text{class}) = 0$

P(features | spam) = 0 and *P*(features | ham) = 0

Tiebreaker: randomly select one of the classes?

Smoothing

P(features | spam) = 0 and P(features | 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 ⟹ make it a small prob event

Estimating Parameters with Training Data

#spam → #spam+1

ham → #ham+1

smoothing

prob. of class

		Without Smoothing	With Smoothing
parameter:	P(spam)	estimate: $\frac{\#spam}{\#spam + \#ham}$	$\frac{\#spam + 1}{\#spam + 1 + \#ham + 1}$
+		+	+
parameter:	P(ham)	estimate: $\frac{\#ham}{\#spam + \#ham}$	$\frac{\#ham + 1}{\#spam + 1 + \#ham + 1}$
=	1	=	=
		1	1

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

Estimating Parameters with Training Data

smoothing conditional prob

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

estimate:

Without Smoothing

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

0 if word never appeared

With Smoothing

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

$$= \frac{\#(\text{spam containing word i}) + 1}{\#\text{spam} + 2}$$

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

$$\frac{\#\text{ham with abacus} + 1}{\#\text{ham with abacus} + 1 + \#\text{ham without abacus} + 1} = \frac{1}{\#\text{ham} + 2}$$

(low probability)

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

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

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 | ham)** according to a naïve Bayes classifier ***with smoothing***?

A. P(x$_3$= 1 | ham) = 0
B. P(x$_3$= 1 | ham) = 1/2
C. P(x$_3$= 1 | ham) = 1/(total #ham +1)
D. P(x$_3$= 1 | ham) = 1/(total #ham + 2)
E. P(x$_3$= 1 | ham) = 1/(total #ham + total #spam+ 2)

conditioned on ham

Smoothing

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

	without smoothing?	with smoothing?
A.	P(ham) = 0	P(ham) = 1/8
B.	P(ham) = 0	P(ham) = 6/7
C.	P(ham) = 1	P(ham) = 1/8
D.	P(ham) = 1	P(ham) = 6/7
E.	P(ham) = 1	P(ham) = 7/8

consider
#ham = 6k

$\frac{6k+1}{6k+2} \approx \frac{6k}{6k} \approx 1 > \frac{7}{8}$

almost p = 1

why?
we have more training data!

all my training were ham

$P(\text{ham}) = \frac{\#\text{ham}+1}{\#\text{ham}+1+\#\text{spam}+1} = \frac{6+1}{6+2} = \frac{7}{8}$

Smoothing

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

	without smoothing?	**with smoothing?**
A.	P(ham) = 0	P(ham) = 1/8
B.	P(ham) = 0	P(ham) = 6/7
C.	P(ham) = 1	P(ham) = 1/8
D.	P(ham) = 1	P(ham) = 6/7
E.	P(ham) = 1	P(ham) = 7/8

Is it still true that *P*(ham) + *P*(spam) = 1 with smoothing?

$$P(spam) = \frac{\#spam + 1}{\#spam + 1 + \#ham + 1} = \frac{1}{\#ham + 2} = \frac{1}{8} \qquad \frac{1}{8} + \frac{7}{8} = 1 \checkmark$$

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.