Tutorial for AI

Thursday, September 4, 2014 3:03 PM

moahmed@cs.ubc.ca

use to learn y

$[x^1, x^2 \cdots x^D] \rightarrow y$ discrete

given features, how to find label

features: ex. a word (world cup) to guess what kind of document (in document, or count in document)

$p(y \mid [x^1 \cdots x^D])$ probability of y given the features

pick the one with the maximum probability

features:	"world cup"	"Germany"	"Brazil"
binary:	1	0	1

↑ word not in document

(Y) document sports?

mention "world cup winner" — probability: highest "germany" (since they won)

(Y) tourism

"visit" — probability: any nation

$$p(y \mid \underbrace{[x^1 \cdots x^D]}_{x}) = \frac{p(x \mid y)\,p(y)}{p(x)} \quad \Leftarrow \text{bayes rule}$$

$$= \frac{p([x^1 x^2 \cdots x^D] \mid y)\,p(y)}{p([x^1 x^2 \cdots x^D])}$$

features are independent given the labels

feature: words

label: "sports document"

so, given

features are independent given the document

given this assumption of independence,

$$= \frac{\prod_{i=1}^{D} p(x^i \mid y)\,p(y)}{p(x)}$$

probability of word in document

how can we calculate $p(y)$?

feature → class

estimate the probability of y

ratio; count compared to total

$y = 1 = \frac{\text{\#} = 1}{\text{total \#}}$

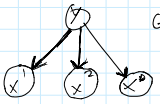

Given y, these features are independent

or do dependency on previous word

$p(x^2 \mid x^1)$.. etc

avoid $p = 0$ with y having no occurrences?

prior (our beliefs about the documents)

then adjust by ratio from training data

add certain value, normalize.

5.2 Bayes rule for medical diagnosis

After your yearly checkup, the doctor has bad news and good news. The bad news is that you tested positive for a serious disease, and that the test is 99% accurate (i.e., the probability of testing positive given that you have the disease is 0.99, as is the probability of testing negative given that you don't have the disease). The good news is that this is a rare disease, striking only one in 10,000 people. What are the chances that you actually have the disease?

$p(\text{test} = \text{true} \mid \text{Disease}) = 0.99$

$p(\text{test} = \text{true} \mid \text{No disease}) = 0.99$

$p(\text{Dis}) = 10^{-4}$

$$p(\text{Dis} \mid \text{test} = \text{pos}) = \frac{p(\text{test} = \text{true}) * p(\text{Dis})}{p(\text{test} = \text{pos})}$$

numerator: pos, dis / pos, no dis