#### Naïve Bayes Classifiers

Doug Downey Northwestern EECS 395/495 Fall 2011

## Naïve Bayes Classifiers

- Combines all ideas we've covered
  - Conditional Independence
  - Bayes' Rule
  - Statistical Estimation
  - Machine Learning
- ... in a simple, yet accurate classifier
  - Classifier: Function  $f(\mathbf{x})$  from  $\mathbf{X} = \{\langle x_1, \dots, x_d \rangle\}$  to *Class*
  - E.g., X = {<GRE, GPA, Letters>}, Class = {yes, no, wait}

## Probability => Classification (1 of 2)

- Classification Task:
  - Learn function  $f(\mathbf{x})$  from  $\mathbf{X} = \{\langle x_1, ..., x_d \rangle\}$  to *Class*
  - Given: Examples D={(x, y)}
- Probabilistic Approach
  - Learn P(Class = y | X = x) from D
  - Given **x**, pick the maximally probable y

# Probability => Classification (2 of 2)

- More formally
  - $f(\mathbf{x}) = \arg \max_{y} P(Class = y | \mathbf{X} = \mathbf{x}, \boldsymbol{\theta}_{MAP})$
  - $\theta_{\text{MAP}}$  : MAP parameters, learned from data
    - That is, parameters of P(Class = y | X = x)
  - ...we'll focus on using MAP estimate, but can also use ML or Bayesian
- Predict next coin flip? Instance of this problem
  - X = null
  - Given D= hhht...tht, estimate P( $\theta \mid D$ ), find MAP
  - Predict *Class* = heads iff  $\theta_{MAP}$  >  $\frac{1}{2}$

### **Example: Text Classification**

Dear Sir/Madam,

We are pleased to inform you of the result of the Lottery Winners International programs held on the 30/8/2004. Your e-mail address attached to ticket number: EL-23133 with serial Number: EL-123542, batch number: 8/163/EL-35, lottery Ref number: EL-9318 and drew lucky numbers 7-1-8-36-4-22 which consequently won in the 1st category, you have therefore been approved for a lump sum pay out of US\$1,500,000.00 (One Million, Five Hundred Thousand United States dollars)

• SPAM

NOT SPAM?

### Representation

- **X** = document
- Estimate P(Class = {spam, non-spam} | X)
- Question: how to represent **X**?
  - One dimension for each possible e-mail, i.e. possible permutation of words?
    - No.
  - Lots of possibilities, common choice: "bag of words"

Dear Sir/Madam,		Sir	1
We are pleased to inform you of the result of the Lottery Winners International programs held on the 30/8/2004. Your e-mail address attached to ticket number: EL-23133 with serial Number: EL-		Lottery	10
123542, batch number: 8/163/EL-35, lottery Ref number: EL-9318 and drew lucky numbers 7-1-8-36-4-22 which consequently won in	>	Dollars	7
the 1st category, you have therefore been approved for a lump sum pay out of US\$1,500,000.00 (One Million, Five Hundred Thousand		With	38
United States dollars) 		•••	

# Bag of Words

- Ignores Word Order, i.e.
  - No emphasis on title
  - No compositional meaning ("Cold War" -> "cold" and "war")
  - Etc.
  - But, massively reduces dimensionality/complexity
- Still and all...
  - Recording presence or absence of a 100,000-word vocab entails 2^100,000 distinct vectors

# Naïve Bayes Classifiers

- *P*(*Class* | *X*) for |Val(*X*)| = 2^100,000 requires
   2^100,000 parameters
  - Problematic.
- Bayes' Rule:
   *P*(*Class* | *X*) = P(*X* | *Class*) P(*Class*) / P(*X*)
- Assume presence of word *i* is independent of all other words given *Class*:
   *P*(*Class* | *X*) = Π<sub>i</sub> P(*w<sub>i</sub>* | *Class*) P(*Class*) / P(*X*)
- Now only 200,001 parameters for *P*(*Class* | **X**)



### Naïve Bayes Assumption

- Features are conditionally independent given class
  - Not P("Republican", "Democrat") = P("Republican")P("Democrat")
     but instead
    - P("Republican", "Democrat" | Class = Politics) =
      P("Republican" | Class = Politics)P("Democrat" | Class = Politics)
- Still, an absurd assumption
  - ("Lottery", "Winner"  $\perp$  SPAM)? ("lunch", "noon"  $\perp$  Not SPAM)?
- But: offers massive tractability advantages and works quite well in practice
  - Lesson: Overly strong independence assumptions sometimes allow you to build an accurate model where you otherwise couldn't

### Getting the parameters from data

- Parameters  $\theta = \langle \theta_{ij} = P(w_i | Class = j) \rangle$
- Maximum Likelihood: Estimate P(w<sub>i</sub> | Class = j) from
   D by counting
  - Fraction of documents in class *j* containing word *i*
  - But if word *i* never occurs in class *j* ?
- Commonly used MAP estimate:
  - $\frac{(\# \text{ docs in class } j \text{ with word } i) + 1}{(\# \text{ docs in class } j) + |V|}$

#### Caveats

- Naïve Bayes effective as a *classifier*
- Not as effective in producing probability estimates
   Π<sub>i</sub> P(w<sub>i</sub> | Class) pushes estimates toward 0 or 1
- In practice, numerical underflow is typical at classification time
  - Compare sum of logs instead of product

I'm at work, neighbor John calls to say my alarm is ringing, but neighbor Mary doesn't call. Sometimes it's set off by minor earthquakes. Is there a burglar?

Variables: Burglar, Earthquake, Alarm, JohnCalls, MaryCalls