# Discrete Features I <br> Exploring Independence and Modeling 

Matthew Stone<br>CS 520, Spring 2000<br>Lecture 7

## Bayesian Decision Theory DISCRETE FEATURES

- Finite set of $c$ states of nature
- Measurement is a discrete feature vector
- E.g., $\mathbf{x} \in\{0,1\}^{k}$
- $k$-dimensional binary feature space
- Provides setting for many key algorithms
- Markov models, belief nets, etc.


## Bayes Decision Rule

- Infer most likely state given measurement
- Using Bayes formula, here:

$$
P\left(\omega_{i} \mid \mathbf{x}\right)=\frac{P\left(\mathbf{x} \mid \omega_{i}\right) P\left(\omega_{i}\right)}{P(\mathbf{x})}
$$

## Bayes Decision Rule

- Again we have "curse of dimensionality" - Need $c 2^{k}$ numbers to specify distribution
- Worse -
- To estimate parameters $P\left(\mathbf{x} \mid \omega_{i}\right)$ with expected accuracy $1 / \varepsilon$
- Need $c \varepsilon^{2} 2^{k}$ training samples


## Possible Solution: Modeling

- Provide a specification outlining sparse relationships among features and classes


## Model Zero Naïve Bayes Classification

- Features are independent given the class
- "Model" requires ck parameters:
- Likelihoods $p_{i j}:=P\left(x_{j}=1 \mid \omega_{i}\right)$
- We get

$$
P\left(\mathbf{x} \mid \omega_{i}\right)=\prod_{j=1}^{k} p_{i j}^{x_{j}}\left(1-p_{i j}\right)^{1-x_{j}}
$$

## Naïve Bayes Classification CONTINUED

- Use usual discriminant function

$$
\begin{aligned}
& \quad g_{i}(\mathbf{x})=\ln P\left(\omega_{i}\right)+\ln P\left(\mathbf{x} \mid \omega_{i}\right) \\
& - \text { i.e: } \\
& \quad g_{i}(\mathbf{x})=\ln P\left(\omega_{i}\right)+\sum_{j=1}^{k}\left[x_{j} \ln p_{i j}+\left(1-x_{j}\right) \ln \left(1-p_{i j}\right)\right] \\
& - \text { i.e: } \\
& g_{i}(\mathbf{x})=\mathbf{w}_{i}^{\top} \mathbf{x}+\theta \\
& w_{i j}=\ln p_{i j}-\ln \left(1-p_{i j}\right) \\
& \theta=\ln P\left(\omega_{i}\right)+\sum_{i=1}^{k} \ln \left(1-p_{i j}\right)
\end{aligned}
$$

## Visualization

- Decision surfaces are hyperplanes


Solutions to equations

$$
g_{h}(\mathbf{x})=g_{j}(\mathbf{x})
$$

slice hypercube

## Case Study

- Text classification
- Assign a natural language document to a predefined category based on content


## Text Classification Examples

- Index medical journal article
- Catalogue book for library
- Fit web page into Yahoo! Hierarchy
- Filter news feed for personal interest
- Automatically delete spam email


## Formalizing Text Classification

- States of nature
- c states representing the different possible categories for documents, e.g.
- Email: $\omega_{1}$-interesting
$\omega_{2}$ - spam


## Formalizing Text Classification

- Measurement x for any document
- $k$-component binary feature vector
- Pick $k$ useful English words
- "useful" means occurs often with good correlation to some class
-Set $x_{i}=1$ if word $i$ occurs in document


## Sparse Data

(Aside)

- English has tens of thousands of words
- But narrow to 1 K that best discriminate
- Still 10300 feature vectors
- much larger than, e.g., go search space
- no hope of describing $\mathrm{P}(\mathbf{x} \mid \mathrm{c})$ without a model


## Naïve Bayes Model

- Assume features are independent
- Take maximum likelihood estimate for

$$
p_{i j}:=P\left(x_{j}=1 \mid \omega_{i}\right)
$$

- That's just
\# of docs in class $\omega_{i}$ containing term $j$
$\#$ of docs in class $\omega_{i}$


## Naïve Bayes Model (CONTINUED)

- Given measurement $x$, Bayes formula has

$$
P\left(\omega_{i} \mid \mathbf{x}\right)=\frac{\left\lfloor\prod_{j} P\left(x_{j}=1 \mid \omega_{i}\right) J P\left(\omega_{i}\right)\right.}{P(\mathbf{x})}
$$

- So compute
$P\left(\omega_{i} \mid \mathbf{x}\right) \propto \frac{\# \text { in class } \omega_{i}}{\# \text { of docs }} \prod_{j} \frac{\# \text { in class } \omega_{i} \text { containing term } j}{\# \text { in class } \omega_{i}}$


## Common Pitfall

- With parameters set by MLE, you could easily end up with all posteriors zero
- To see how, suppose:
- Each feature occurs with some high probability $p$ in a single class and some very low probability elsewhere: $1 / \varepsilon$
- You'd want some $2 \varepsilon$ samples per class, but you can't get that many - only $n$
- MLE estimate of $P\left(x_{j}=1 \mid \omega_{i}\right)$ is often zero


## Common Pitfall

- With parameters set by MLE, you could easily end up with all posteriors zero
- To see how, suppose:
- The number $k$ of features associated with each measurement is large
- You expect a rare feature to occur on a test guy with probability roughly $k / \varepsilon$
- Get rare feature not seen on any trainer from this category $-(k \varepsilon-n) / \varepsilon^{2}$ of the time


## Sparse Data Requires Smoothing

- Redistribute probability mass
- from what you saw
- to what you didn't see
- since you know other things can happen


## Simple Smoothing: Deleted Estimation

- Key question is often
- How often do you expect features in test data that never occur in training?
- Deleted estimation finds this
- by splitting training data
- and answering question empirically


## Deleted Estimation (CONTINUED)

- Take first half
- $N_{0}^{1}$ - how many features don't occur there
$-C_{0}^{12}$ - how many of these occur in half two
- Take second half
- $N_{0}^{2}$ - how many features don't occur there
$-C_{0}^{21}$ - how many of these occur in half two


## Deleted Estimation (CONTINUED)

- This gives evidence about how often new things happen

$$
r_{0}=\frac{C_{0}^{12}+C_{0}^{21}}{N_{0}^{1}+N_{0}^{2}}
$$

- Smoothed value replaces MLE estimate
- Similar smoothed values required for other counts, to ensure probabilities sum to one

