# 433-352 Data on the Web 

Semester 2, 2007<br>Lecturer: Tim Baldwin



## Lecture 8

## Text Categorisation (2)

## Bayesian Methods

- Learning and classification methods based on probability theory
- Build a generative model that approximates how data is produced
- Categorisation produces a posterior probability distribution over the possible categories given a description of an instance


## Bayes' Rule

$$
\begin{gathered}
P(C, X)=P(C \mid X) P(X)=P(X \mid C) P(C) \\
P(C \mid X)=\frac{P(X \mid C) P(C)}{P(X)}
\end{gathered}
$$

## Naive Bayes (NB) Classifiers

- Task: classify an instance $D=\left\langle x_{1}, x_{2}, \ldots, x_{n}\right\rangle$ according to one of the classes $c_{j} \in C$

$$
\begin{aligned}
c & =\operatorname{argmax}_{c_{j} \in C} P\left(c_{j} \mid x_{1}, x_{2}, \ldots, x_{n}\right) \\
& =\operatorname{argmax}_{c_{j} \in C} \frac{P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c_{j}\right) P\left(c_{j}\right)}{P\left(x_{1}, x_{2}, \ldots, x_{n}\right)} \\
& =\operatorname{argmax}_{c_{j} \in C} P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c_{j}\right) P\left(c_{j}\right)
\end{aligned}
$$

## Simplifying Assumption

- $P\left(c_{j}\right)$
* can be estimated from the frequency of classes in the training examples [maximum likelihood estimate]
- $P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c_{j}\right)$
* $O\left(|X|^{n}|C|\right)$ parameters (cannot be estimated in practice)
- Naive Bayes Conditional Independence Assumption:
* assume that the probability of observing the conjunction of attributes is equal to the product of the individual probabilities $P\left(x_{i} \mid c_{j}\right)$ [hence "naive"]


## The Final NB Formulation

- Applying the conditional independence assumption:

$$
\begin{aligned}
c & =\operatorname{argmax}_{c_{j} \in C} P\left(x_{1}, x_{2}, \ldots, x_{n} \mid c_{j}\right) P\left(c_{j}\right) \\
& =\operatorname{argmax}_{c_{j} \in C} P\left(c_{j}\right) \prod_{i} P\left(x_{i} \mid c_{j}\right)
\end{aligned}
$$

## Multivariate Binomial NB

- Represent each word as a binary feature (= DF model)
- Represent a document according to the word types it contains
- No indication of how often a given word occurs in a given document
- "Bag of word types" document model


## Multivariate Binomial NB: Mechanics

$$
P\left(D \mid c_{i}\right)=\prod_{j=1}^{|\mathscr{V}|}\left(B_{j} P\left(t_{j} \mid c_{i}\right)+\left(1-B_{j}\right)\left(1-P\left(t_{j} \mid c_{i}\right)\right)\right)
$$

where $B_{j} \in\{0,1\}$ indicates the presence or absence of the $j$ th term in $D, \mathscr{V}$ is the set of all terms, and

$$
P\left(t \mid c_{i}\right)=\frac{1+\sum_{k=1}^{|\mathscr{D}|} B_{k} P\left(c_{i} \mid D_{k}\right)}{2+\sum_{k=1}^{|\mathscr{O}|} P\left(c_{i} \mid D_{k}\right)}
$$

## Multivariate Binomial NB: Example

- Test document:
we few, we happy few, we band of brothers
- Test document representation:
$\langle 0, \quad 0, \ldots, 1, \ldots 0, \ldots 1, \ldots, 1, \ldots, 1, \ldots 0, \ldots, 1, \ldots, 0\rangle$
aarvark aback band betwixt brothers few happy thee we zymogen
- Shakespeare training document set:
then happy i, that love and am beloved

if we shadows have offended
$\langle 0, \quad 0, \ldots, 0, \ldots \quad 0, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 1, \ldots, 0\rangle$ aarvark aback band betwixt brothers few happy thee we zymogen
- Beatles training document set:


## we can work it out

$$
\left.\left.\begin{array}{c}
\langle 0,
\end{array} 0, \ldots, 0, \ldots \quad 0, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 1, \ldots, \quad 0\right\rangle\right\rangle
$$

sgt pepper's lonely hearts club band

$$
\begin{gathered}
\langle 0, \quad 0, \ldots, 1, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0\rangle \\
\text { aarvark aback } \quad \text { band betwixt brothers few happy thee we zymogen }
\end{gathered}
$$

- $P($ we $\mid$ Shakespeare $)=\frac{1+(0 \times 1+1 \times 1+1 \times 0+0 \times 0)}{2+(1+1+0+0)}=\frac{1}{2}$ $P($ we $\mid$ Beatles $)=\frac{1+(0 \times 0+1 \times 0+1 \times 1+0 \times 1)}{2+(0+0+1+1)}=\frac{1}{2}$
$P($ band $\mid$ Shakespeare $)=\frac{1+(0 \times 1+0 \times 1+0 \times 0+1 \times 0)}{2+(1+1+0+0)}=\frac{1}{4}$
$P($ band $\mid$ Beatles $)=\frac{1+(0 \times 0+0 \times 0+0 \times 1+1 \times 1)}{2+(0+0+1+1)}=\frac{1}{2}$
$P($ happy $\mid$ Shakespeare $)=\frac{1+(1 \times 1+0 \times 1+0 \times 0+0 \times 0)}{2+(1+1+0+0)}=\frac{1}{2}$
$P($ happy $\mid$ Beatles $)=\frac{1+(1 \times 0+0 \times 0+0 \times 0+0 \times 0)}{2+(0+0+1+1)}=\frac{1}{4}$
- $P(D \mid$ Shakespeare $)=\left(\left(0 \times \frac{1}{4}+(1-0) \times \frac{3}{4}\right) \times\left(0 \times \frac{1}{4}+(1-0) \times \frac{3}{4}\right) \times \ldots \times(1 \times\right.$ $\left.\frac{1}{4}+(1-1) \times \frac{3}{4}\right) \times \ldots \times\left(0 \times \frac{1}{4}+(1-0) \times \frac{3}{4}\right) \times \ldots \times\left(1 \times \frac{1}{4}+(1-1) \times \frac{3}{4}\right) \times \ldots \times$ $\left(1 \times \frac{1}{4}+(1-1) \times \frac{3}{4}\right) \times \ldots \times\left(1 \times \frac{1}{2}+(1-1) \times \frac{1}{2}\right) \times \ldots \times\left(0 \times \frac{1}{4}+(1-0) \times \frac{3}{4}\right) \times$ $\left.\ldots \times\left(1 \times \frac{1}{2}+(1-1) \times \frac{1}{2}\right) \times \ldots \times\left(0 \times \frac{1}{4}+(1-0) \times \frac{3}{4}\right)\right)$


## Multinomial NB

- Represent each word as an integer
- Represent a document according to the word tokens it contains
- Optionally include a term for $P\left(L=l_{D} \mid c_{i}\right)$ (to normalise for document length)
- "Bag of word tokens" document model
- Assumes that (a) the position of a word in the document and (b) the context of a word are irrelevant in classification


## Multinomial NB: Mechanics

$$
P\left(D \mid c_{i}\right)=\prod_{j=1}^{|\mathcal{V}|} \frac{P\left(t_{j} \mid c_{i}\right)^{N_{D, t_{j}}}}{N_{D, t_{j}}!}
$$

where $N_{D, t_{j}}$ is the frequency of the $j$ th term in $D, \mathscr{V}$ is the set of all terms, $l_{D}$ is the length of $D$, and

$$
P\left(t \mid c_{i}\right)=\frac{1+\sum_{k=1}^{|\mathscr{D}|} N_{k, t} P\left(c_{i} \mid D_{k}\right)}{|\mathscr{V}|+\sum_{j=1}^{|\mathscr{V}|} \sum_{k=1}^{|\mathscr{D}|} N_{k, t_{j}} P\left(c_{i} \mid D_{k}\right)}
$$

## Multinomial NB: Example

- Test document:
we few, we happy few, we band of brothers
- Test document representation:

aarvark aback band betwixt brothers few happy thee we zymogen
- Assume $|\mathscr{V}|=100$
- Shakespeare training document set:
then happy i, that love and am beloved

if we shadows have offended
$\langle 0, \quad 0, \ldots, 0, \ldots \quad 0, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 1, \ldots, 0\rangle$ aarvark aback band betwixt brothers few happy thee we zymogen
- Beatles training document set:


## we can work it out

$$
\left.\left.\begin{array}{c}
\langle 0,
\end{array} 0, \ldots, 0, \ldots \quad 0, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 1, \ldots, \quad 0\right\rangle\right\rangle
$$

sgt pepper's lonely hearts club band

$$
\begin{gathered}
\langle 0, \quad 0, \ldots, 1, \ldots \quad 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0, \ldots, 0\rangle \\
\text { aarvark aback } \quad \text { band betwixt brothers few happy thee we zymogen }
\end{gathered}
$$

- $P($ we $\mid$ Shakespeare $)=\frac{1+(0 \times 1+1 \times 1+1 \times 0+0 \times 0)}{100+(8+5)}=\frac{2}{113}$
$P($ we $\mid$ Beatles $)=\frac{1+(0 \times 0+1 \times 0+1 \times 1+0 \times 1)}{100+(5+6)}=\frac{2}{111}$
$P($ band $\mid$ Shakespeare $)=\frac{1+(0 \times 1+0 \times 1+0 \times 0+1 \times 0)}{100+(8+5)}=\frac{1}{113}$
$P($ band $\mid$ Beatles $)=\frac{1+(0 \times 0+0 \times 0+0 \times 1+1 \times 1)}{100+(5+6)}=\frac{2}{111}$
$P($ happy $\mid$ Shakespeare $)=\frac{1+(1 \times 1+0 \times 1+0 \times 0+0 \times 0)}{100+(8+5)}=\frac{2}{113}$
$P($ happy $\mid$ Beatles $)=\frac{1+(1 \times 0+0 \times 0+0 \times 0+0 \times 0)}{100+(5+6)}=\frac{1}{111}$
- $P(D \mid$ Shakespeare $)=\frac{\frac{1}{133}^{0}}{0!} \times \frac{\frac{1}{113}^{0}}{0!} \times \ldots \times \frac{\frac{11}{13}^{1}}{1!} \times \ldots \times \frac{\frac{11}{0}^{0}}{0!} \times \ldots \times \frac{\frac{1}{113}^{1}}{1!} \times \ldots \times \frac{\frac{1}{113}^{2}}{2!} \times$
$\ldots \times \frac{\frac{2}{13}^{1}}{1!} \times \ldots \times \frac{\frac{1}{13}^{0}}{0!} \times \ldots \times \frac{\frac{1}{113}^{3}}{3!} \times \ldots \times \frac{\frac{1}{113}^{113}}{0!} \times$


## Evaluation

- In order to evaluate which of the two models performs best, what vocabulary size to use, etc, we require "held-out" test data
- Evaluation usually takes the form of simple classification accuracy
- Average results over multiple "splits" of training and test data for best results


## Results over the Yahoo Science Dataset



## Theoretical Properties of NB Models

- Multiclass classification method
- Parametric
only have to store attribute-value counts/probabilities for each class, not the actual instances
- Incremental
easy to add extra data to the classifier on the fly
- Simple ( $\rightarrow$ fast)


## Practical Properties of NB Models

- Surprising accuracy over text categorisation tasks
- Highly robust over irrelevant features
- Very good at balancing up lots of "marginally relevant" features
- Actual posterior probability estimates tend to be awry, but as a classification task, we are only interested in the relative values
- Multinomial model tends to outperform binomial
- Feature selection more important for binomial model than multinomial model - why?


## Real World Practicalities

## Extra Features in Text Categorisation

- There's lots more to web text categorisation than words:
* metadata
* domain of source page
* page structure
* link structure
* diachronic stability of page
* balance of different content types
* relative use of different HTML attributes
* well-formedness of HTML :


## Multi-topic Documents

- Realistically, it is possible for a document to belong to multiple categories
- Ways to model this:
* thresholding
* multiclass categories
* probabilistic class assignment


## Summary

- How do Bayesian methods differ from NN methods?
- What are the simplifying assumptions in the NB method?
- What are the two basic variants of the Naive Bayes algorithm, and what are the strengths and weaknesses of each?


## References

Chakrabarti, Soumen. 2003. Mining the Web: Discovering Knowledge from Hypertext Data. San Francisco, USA: Morgan Kaufmann.
Jackson, Peter, and Isabelle Moulinier. 2002. Natural Language Processing for Online Applications: Text Retrieval, Extraction and Categorization. Amsterdam, Netherlands: John Benjamins.
McCallum, Andrew, and Kamal Nigam. 1998. A comparison of event models for Naive Bayes text classification. In Proceedings of the AAAI-98 Workshop on Learning for Text Categorization, Madison, USA.

