Logistic Regression for Text Classification

CS 485, Fall 2024
Applications of Natural Language Processing

Brendan O'Connor
College of Information and Computer Sciences
University of Massachusetts Amherst

[With slides from Ari Kobren and SLP3]

BOW linear model for text classif.

- Problem: classify doc *d* into one of *k*∈*1..K* classes

- Parameters: For each class *k*, and word type *w*, there is a *word weight*

$\beta_{w,k}$

$\beta \in \mathbb{R}^{|V| \times K}$

	k=NEG	k=POS
dog	0.1	0.1
happy	-1.4	3.1
⋮		

- Representation: bag-of-words vector of doc *d*'s word counts

$x = (1, 4, 1, 0, 0 \ldots)$

$x_w \in \mathbb{R}^V$

- Prediction rule: choose class *y* with highest score

$$\text{score}(y) = \sum_{w \in V} x_w \beta_{w,k}$$

$$\hat{y} = \arg\max_{k \in 1..K} \text{score}(k)$$

Keyword count as linear model

- Problem: classify doc d into one of $k \in 1..K$ classes

(e.g. "the $\underbrace{\text{happy}}_{+}$ movie was $\underbrace{\text{great}}_{+}$ ...")

- Parameters: For each class k, and word type w, there is a *word weight*

$$L_{POS} = \{ \text{happy}, \text{great} \dots \}$$

$$L_{NEG} = \{ \dots \}$$

$$L_k = \{ \dots \}$$

$$\beta_{w,k} = 1\{ w \in L_k \}$$

- Representation: bag-of-words vector of doc d's word counts

$$x_w$$

- Prediction rule: choose class y with highest score

$$\Rightarrow \underbrace{\sum_{w \in V} x_w \beta_{w,k}} = \text{how many words from doc. are in lexicon "k"}$$

$$\Rightarrow \arg\max_k \text{score}(k) = \text{highest lex. count}$$

Naive Bayes as linear model

- Problem: classify doc *d* into one of $k \in 1..K$ classes

- Parameters: For each class *k*, and word type *w*, there is a *word weight*

- Representation: bag-of-words vector of doc *d*'s word counts

- Prediction rule: choose class *y* with highest score

Linear classification models

- The foundational model for machine learning-based NLP!

- Examples
  - The humble "keyword count" classifier (no ML)
  - Naive Bayes ("generative" ML)

- Today: **Logistic Regression**
  - a linear classification model, trained to be good at *prediction*
  - allows for *features*
  - used within more complex models (neural networks)

Motivation: feature engineering

- For Naive Bayes, we used counts of each word in the vocabulary (BOW representation). But why not also use....
  - Number of words from "CS485 Crowdsource Positive Lexicon"
  - ...from "CS485 Crowdsource Negative Lexicon" ... or another....
  - Phrases?
  - Words/phrases with negation markers?
  - Number of "!" occurrences?
  - or...?

- NB tends to work poorly when there are many potentially repetitive features (why?) → Cond. indep. assumption is wrong!

Var	Definition	Value in Fig. 5.2
x_1	count(positive lexicon words $\in$ doc)	3
x_2	count(negative lexicon words $\in$ doc)	2
x_3	$\begin{cases} 1 & \text{if “no”} \in \text{doc} \\ 0 & \text{otherwise} \end{cases}$	1
x_4	count(1st and 2nd pronouns $\in$ doc)	3
x_5	$\begin{cases} 1 & \text{if “!”} \in \text{doc} \\ 0 & \text{otherwise} \end{cases}$	0
x_6	ln(word count of doc)	$\ln(66) = 4.19$

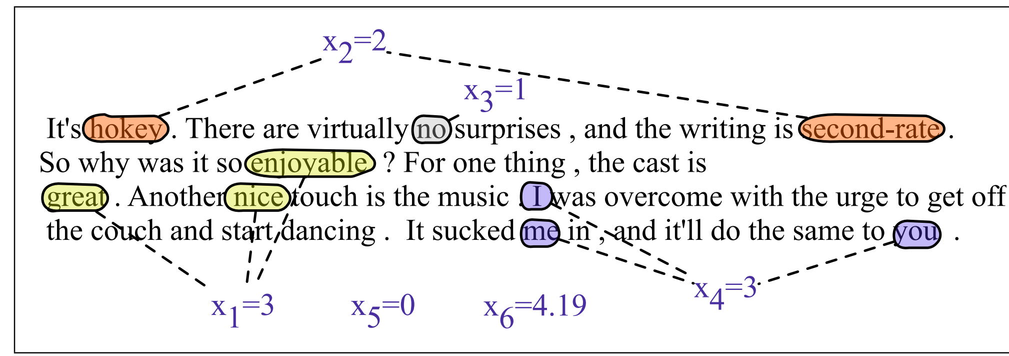

Figure 5.2 A sample mini test document showing the extracted features in the vector x.

- Logistic regression can accommodate ***any arbitrary features***
- <u>Feature engineering</u>: when you spend a lot of trying and testing new features. Very important!! This is a place to put linguistics in, or just common sense about your data.

Negation

Das, Sanjiv and Mike Chen. 2001. Yahoo! for Amazon: Extracting market sentiment from stock message boards. In Proceedings of the Asia Pacific Finance Association Annual Conference (APFA).
Bo Pang, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment Classification using Machine Learning Techniques. EMNLP-2002, 79—86.

Add NOT_ to every word between negation and following punctuation:

```
didn't like this movie , but I
```

⬇

```
didn't NOT_like NOT_this NOT_movie but I
```

[Slide: SLP3]

Classification: LogReg (I)

First, we'll discuss **how LogReg works.**

1 Binary LogReg

3 Learning

2 Multiclass LR

Then, **why** it's set up the way that it is.

Application: **spam filtering**

Classification: LogReg (I)

- compute **features** (xs)

$$x_i = (\text{count "nigerian", count "prince", count "nigerian prince"})$$

- given **weights** (betas)

$$\beta = (-1.0, \quad -1.0, \quad 4.0)$$

Classification: LogReg (II)

- Compute the **dot product**

$$\boxed{z} = x^T\beta = x'\beta = \sum_{j=1}^{3} x_j \beta_j$$

$$\underset{x}{[\;]}\;\underset{\beta}{[\;]} \Rightarrow \underset{x^T}{[\quad]}\;\underset{\beta}{[\;]} = \underset{[\;]}{z} \quad 1\times 1$$

- Compute the **logistic function** for the label probability

$$p(y=1|x) = \frac{e^z}{1+e^z} = \frac{1}{1+e^{-z}}$$

$$p(y=0|x) = 1 - p(y=1|x)$$

9/17/24 exercise. Name: ______________________

LogReg Exercise

features: (count "nigerian", count "prince", count "nigerian prince")

x = (1, 1, 1)

β = (-1.0, -1.0, 4.0)

$z = x^T\beta = x_1\beta_1 + x_2\beta_2 + x_3\beta_3 = 2$

P(y=1 | x) = $\frac{1}{1+e^{-z}} = \frac{1}{1+e^{-2}} = 0.88$

Classification: Dot Product

$N_{feat} = |V|$ for BOW

$$z = \sum_{j=1}^{\boxed{\text{Nfeat}}} \beta_j x_{ij}$$

Weighted sum of feature values

"Linear Model"

Why the **logistic function**?

$$\sigma(z) = \frac{1}{1+e^{-z}}$$

$$z \in \mathbb{R} = (-\infty, \infty)$$

$$p(y=1|x) \in [0, 1]$$

Inputs & Outputs

$z \to \infty$

$$\frac{1}{1+\underbrace{e^{-z}}_{\to 0}} \to 1$$

Logistic Function

$$P(z) = \frac{e^z}{e^z + 1} = \frac{1}{1 + e^{-z}}$$

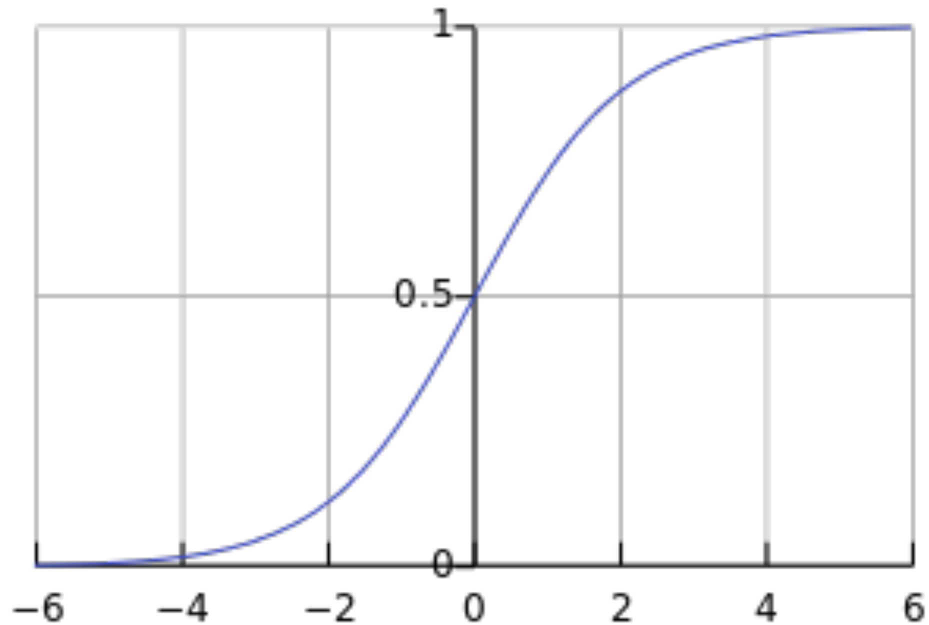

Multiclass Logistic Regression

“Softmax Reg”
“Mult. LR”
“LR”

- Generalize to K>2 classes
- Each class has its own weight vector (across all features; e.g. BOW counts)

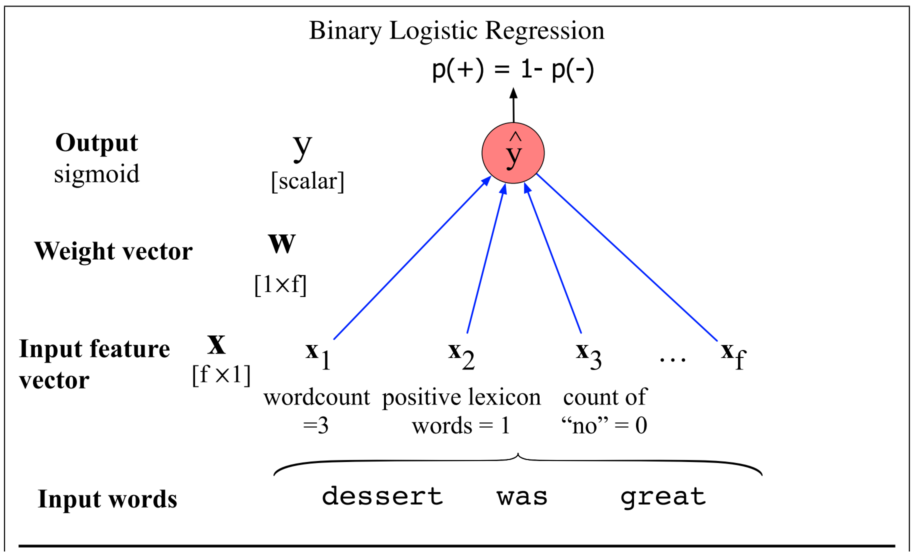

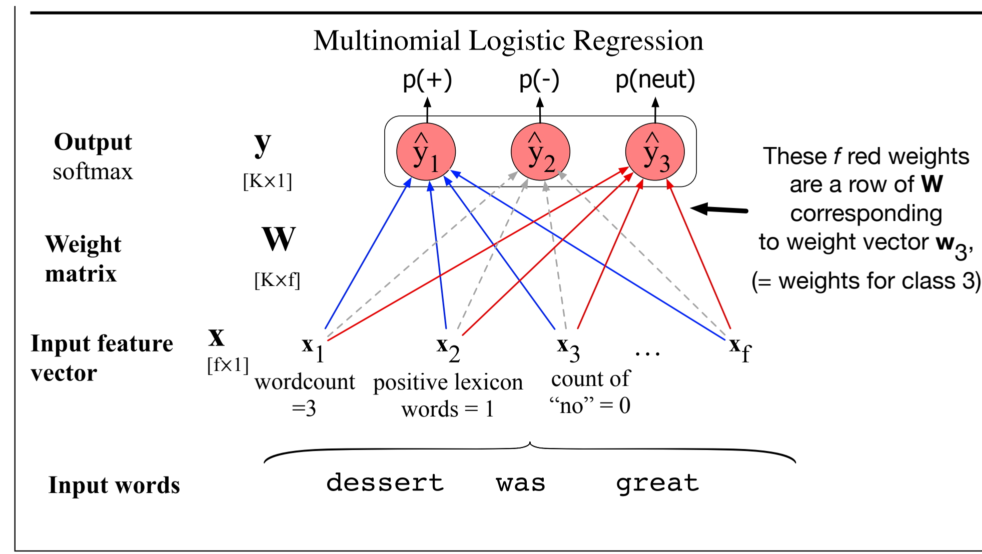

Multiclass Logistic Regression

- Weight vector for each class

$\beta_k \in \mathbb{R}^f$ e.g. $\beta_{k,f} = 3.2$

for $k = 1..K$

- Prediction: dot product for each class

$z_1 = \beta_1^T x = \sum_{j=1}^{f} \beta_{1,j} x_j$

$z_2 = \beta_2^T x$

$\vdots$

$z_K = \beta_K^T x$

- Predicted probabilities: apply the *softmax function* to normalize

$$p(y=k|x) = \frac{e^{z_k}}{\sum_{l=1}^{K} e^{z_l}}$$

Why the softmax function?

want $[p(y=1|x), p(y=2|x), \dots p(y=K|x)]$

① Non-neg ② Sum to 1

have $[z_1, z_2, \dots z_K]$

$$\Downarrow \quad \Downarrow$$

$$e^{z_1} \quad e^{z_2} \quad \dots \quad e^{z_K}$$

$$\Rightarrow \frac{e^{z_2}}{\sum_{\ell=1}^{K} e^{z_\ell}}$$

NB as Log-Linear Model

$$P(\text{spam}|D) \propto P(\text{spam}) \cdot \prod_{w_i \in D} P(w_i|\text{spam})$$

NB as Log-Linear Model

$$P(\text{spam}|D) \propto P(\text{spam}) \cdot \prod_{w_i \in D} P(w_i|\text{spam})$$

$$P(\text{spam}|D) \propto P(\text{spam}) + \prod_{w_i \in \text{Vocab}} \cdot P(w_i|\text{spam})^{x_i}$$

NB as Log-Linear Model

$$P(\text{spam}|D) \propto P(\text{spam}) \cdot \prod_{w_i \in D} P(w_i|\text{spam})$$

$$P(\text{spam}|D) \propto P(\text{spam}) \quad \cdot \prod_{w_i \in \text{Vocab}} \cdot P(w_i|\text{spam})^{x_i}$$

$$\log[P(\text{spam}|D)] \propto \log[P(\text{spam})] + \sum_{w_i \in \text{Vocab}} x_i \cdot \log[P(w_i|\text{spam})]$$

NB as log-linear model

$$P(\text{spam} \mid D) = \frac{1}{Z} P(\text{spam}) \prod_{t=1}^{\text{len}(D)} P(w_t \mid \text{spam})$$

$$P(\text{spam} \mid D) = \frac{1}{Z} P(\text{spam}) \prod_{w \in \mathcal{V}} P(w \mid \text{spam})^{x_w}$$

$$\log P(\text{spam} \mid D) = \log P(\text{spam}) + \sum_{w \in \mathcal{V}} x_w \log P(w \mid \text{spam}) - \log Z$$

NB as Log-Linear Model

In both NB and LogReg

we **compute the dot product!**

NB vs. LogReg

- Both compute the dot product

- **NB**: sum of log probs; **LogReg**: logistic fun.

Learning Weights

- **NB**: learn conditional probabilities separately via **counting**

- **LogReg**: learn weights **jointly**

Learning Weights

- given: a set of **feature vectors** and **labels**

- goal: learn the weights.

Learning Weights

$$\begin{matrix} x_{00} & x_{01} & \cdots & x_{0m} & y_0 \\ x_{10} & x_{11} & \cdots & x_{1m} & y_1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ x_{n0} & x_{n1} & \cdots & x_{nm} & y_n \end{matrix}$$

n examples; xs - features; ys - class

Learning Weights

We know:

$$g(z) = \frac{1}{1 + e^{-z}} \qquad P(y = 1 \mid x) = g\left(\sum_{j=1}^{\mathrm{Nfeat}} \beta_j x_{ij}\right)$$

So let's try to maximize probability of the entire dataset - **maximum likelihood estimation**

Learning Weights

So let's try to maximize probability of the entire dataset - **maximum likelihood estimation**

$$\beta^{MLE} = \arg\max_{\beta} \log P(y_0, \ldots, y_n | \mathbf{x_0}, \ldots, \mathbf{x_n}; \beta)$$

Gradient ascent/descent learning

$$\beta^{MLE} = \arg\max_{\beta} \log P(y_0, \ldots, y_n | \mathbf{x_0}, \ldots, \mathbf{x_n}; \beta)$$

- Follow direction of *steepest ascent*. Iterate: $\beta^{(new)} = \beta^{(old)} + \eta \frac{\partial \ell}{\partial \beta}$

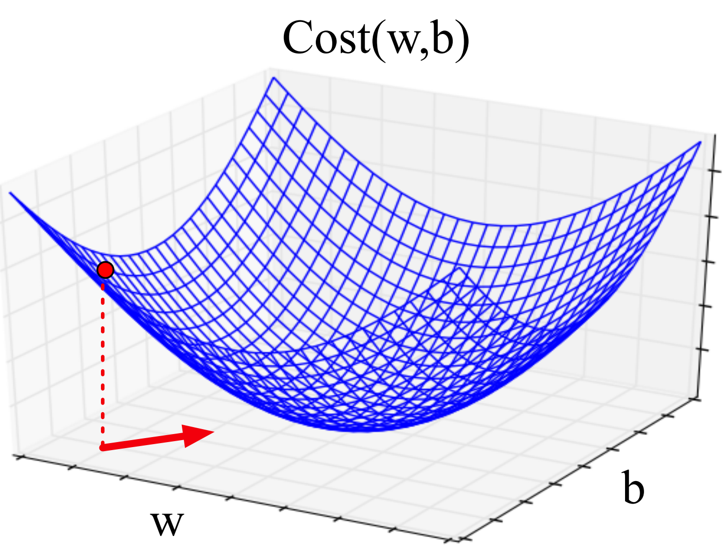

$\left(\frac{\partial \ell}{\partial \beta_1}, \ldots, \frac{\partial \ell}{\partial \beta_J} \right)$: Gradient vector (vector of per-element derivatives)

GD is a generic method for optimizing differentiable functions — widely used in machine learning!

Pros & Cons

- LogReg doesn’t assume independence
  - better calibrated probabilities

- NB is faster to train; less likely to overfit

NB & Log Reg

- Both are linear models:

$$z = \sum_{j=1}^{\mathrm{Nfeat}} \beta_j x_{ij}$$

- Training is different:
  - NB: weights trained independently
  - LogReg: weights trained jointly

Overfitting and generalization

- Overfitting: your model performs overly optimistically on training set, but generalizes poorly to other data (even from same distribution)
- To diagnose: separate training set vs. test set.
- How did we regularize Naive Bayes and language modeling?

- For logistic regression: L2 regularization for training

Regularization tradeoffs

- No regularization <---------------> Very strong regularization

Visualizing a classifier in feature space

"*Bias term*"

Feature vector $x = (1,\ \text{count ``happy''},\ \text{count ``hello''}, \ldots)$

Weights/parameters $\beta =$

50% prob where

$$\beta^{\mathsf{T}} x = 0$$

Predict y=1 when

$$\beta^{\mathsf{T}} x > 0$$

Predict y=0 when

$$\beta^{\mathsf{T}} x \leq 0$$

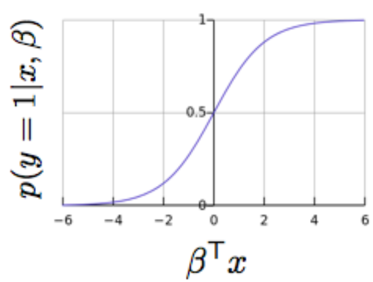

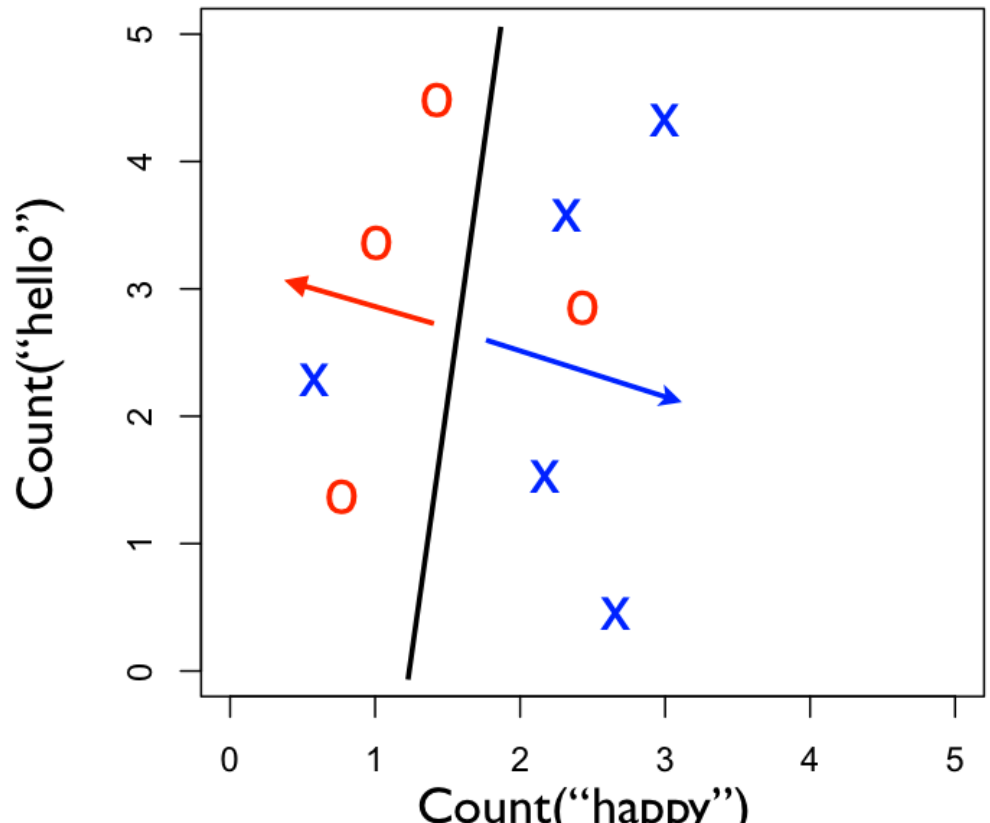

Logistic regression wrap-up

- Given you can extract features from your text, logistic regression is the best, easy-to-use, method
  - Logistic regression with BOW features is an excellent baseline method to try at first
  - Will be a foundation for more sophisticated models, later in course
- Always regularize your LR model
- We recommend using the implementation in scikit-learn
  - Useful: CountVectorizer to help make BOW count vectors

- Next: but where do the LABELS in supervised learning come from?