

Logistic Regression for Text Classification

CS 485, Fall 2023

Applications of Natural Language Processing

https://people.cs.umass.edu/~brenocon/cs485_f23/

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 $k \in 1..K$ 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

Keyword count 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

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**
 - probabilistic model directly geared for *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 Crowdsourced Positive Lexicon"
 - ...from "CS485 Crowdsourced Negative Lexicon" ... or another...
 - Phrases?
 - Words/phrases with negation markers?
 - Number of "!" occurrences?
 - or...?



| 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 | $\ln(\text{word count of doc})$ | $\ln(66) = 4.19$ |

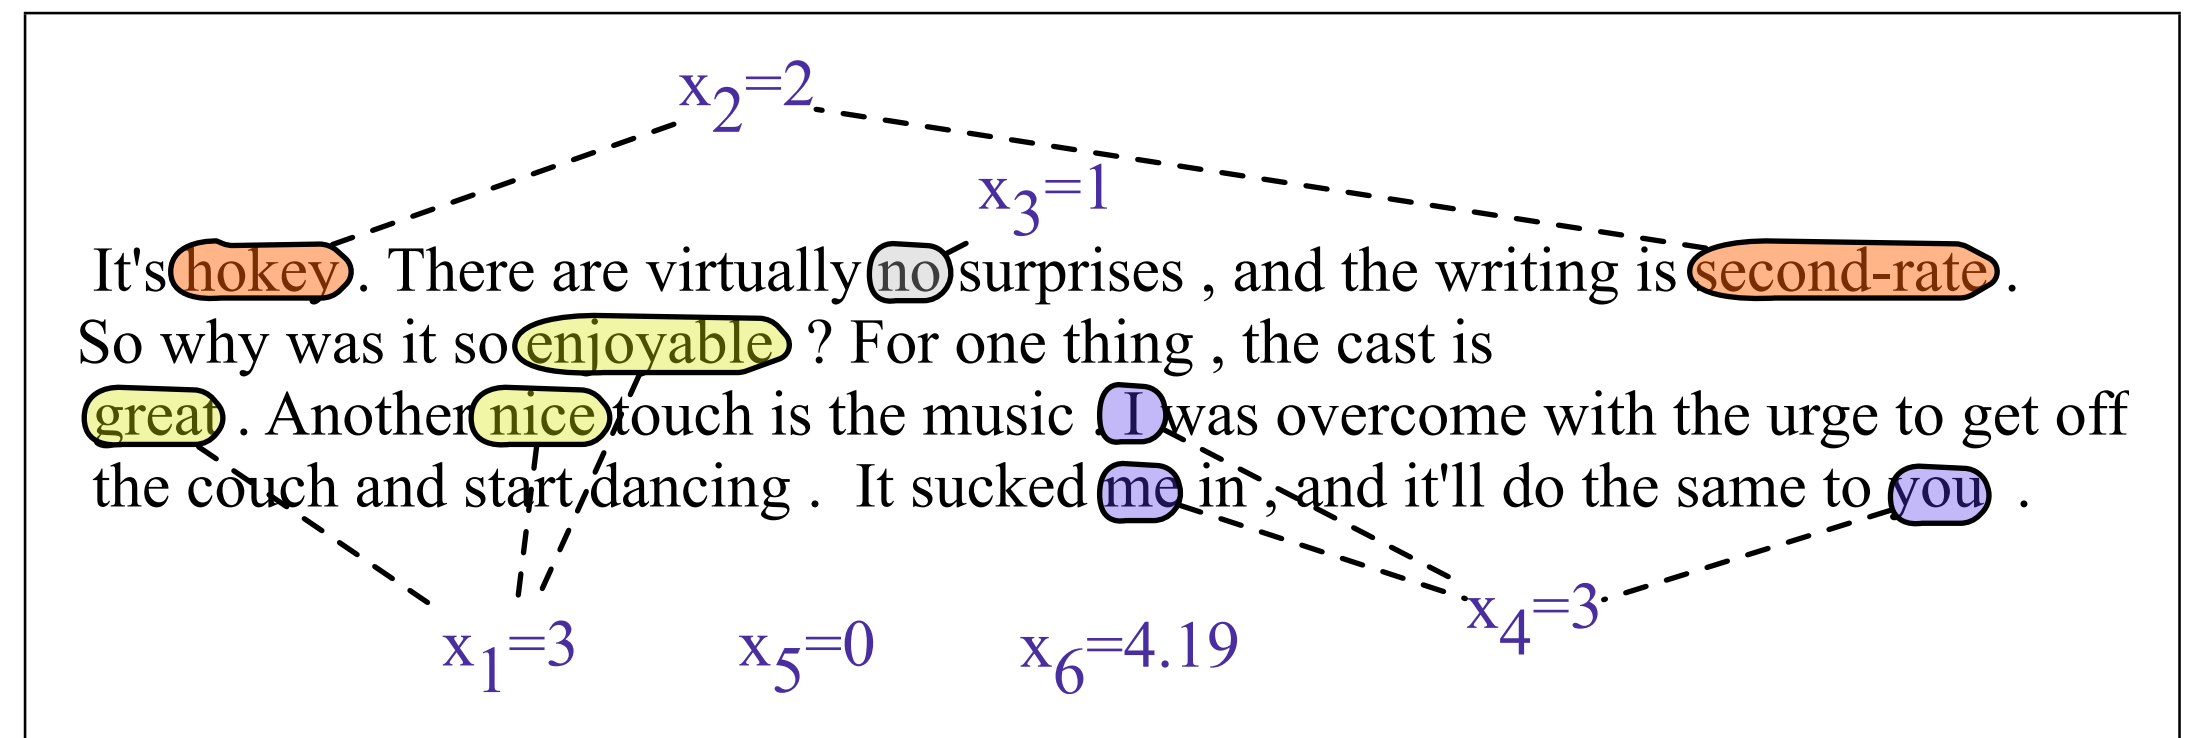


Figure 5.2 A sample mini test document showing the extracted features in the vector x .

- Logistic regression can accommodate **any arbitrary features**
- Feature engineering: 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.**

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

Application: **spam filtering**

Classification: LogReg (I)

- compute **features** (x s)

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

- given **weights** (betas)

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

LogReg Exercise

features: (count “nigerian”, count “prince”, count “nigerian prince”)

$$\mathbf{x} = (1, 1, 1)$$

$$\boldsymbol{\beta} = (-1.0, -1.0, 4.0)$$

$$P(y=1 \mid \mathbf{x}) =$$


Classification: Dot Product

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

Why the **logistic function**?

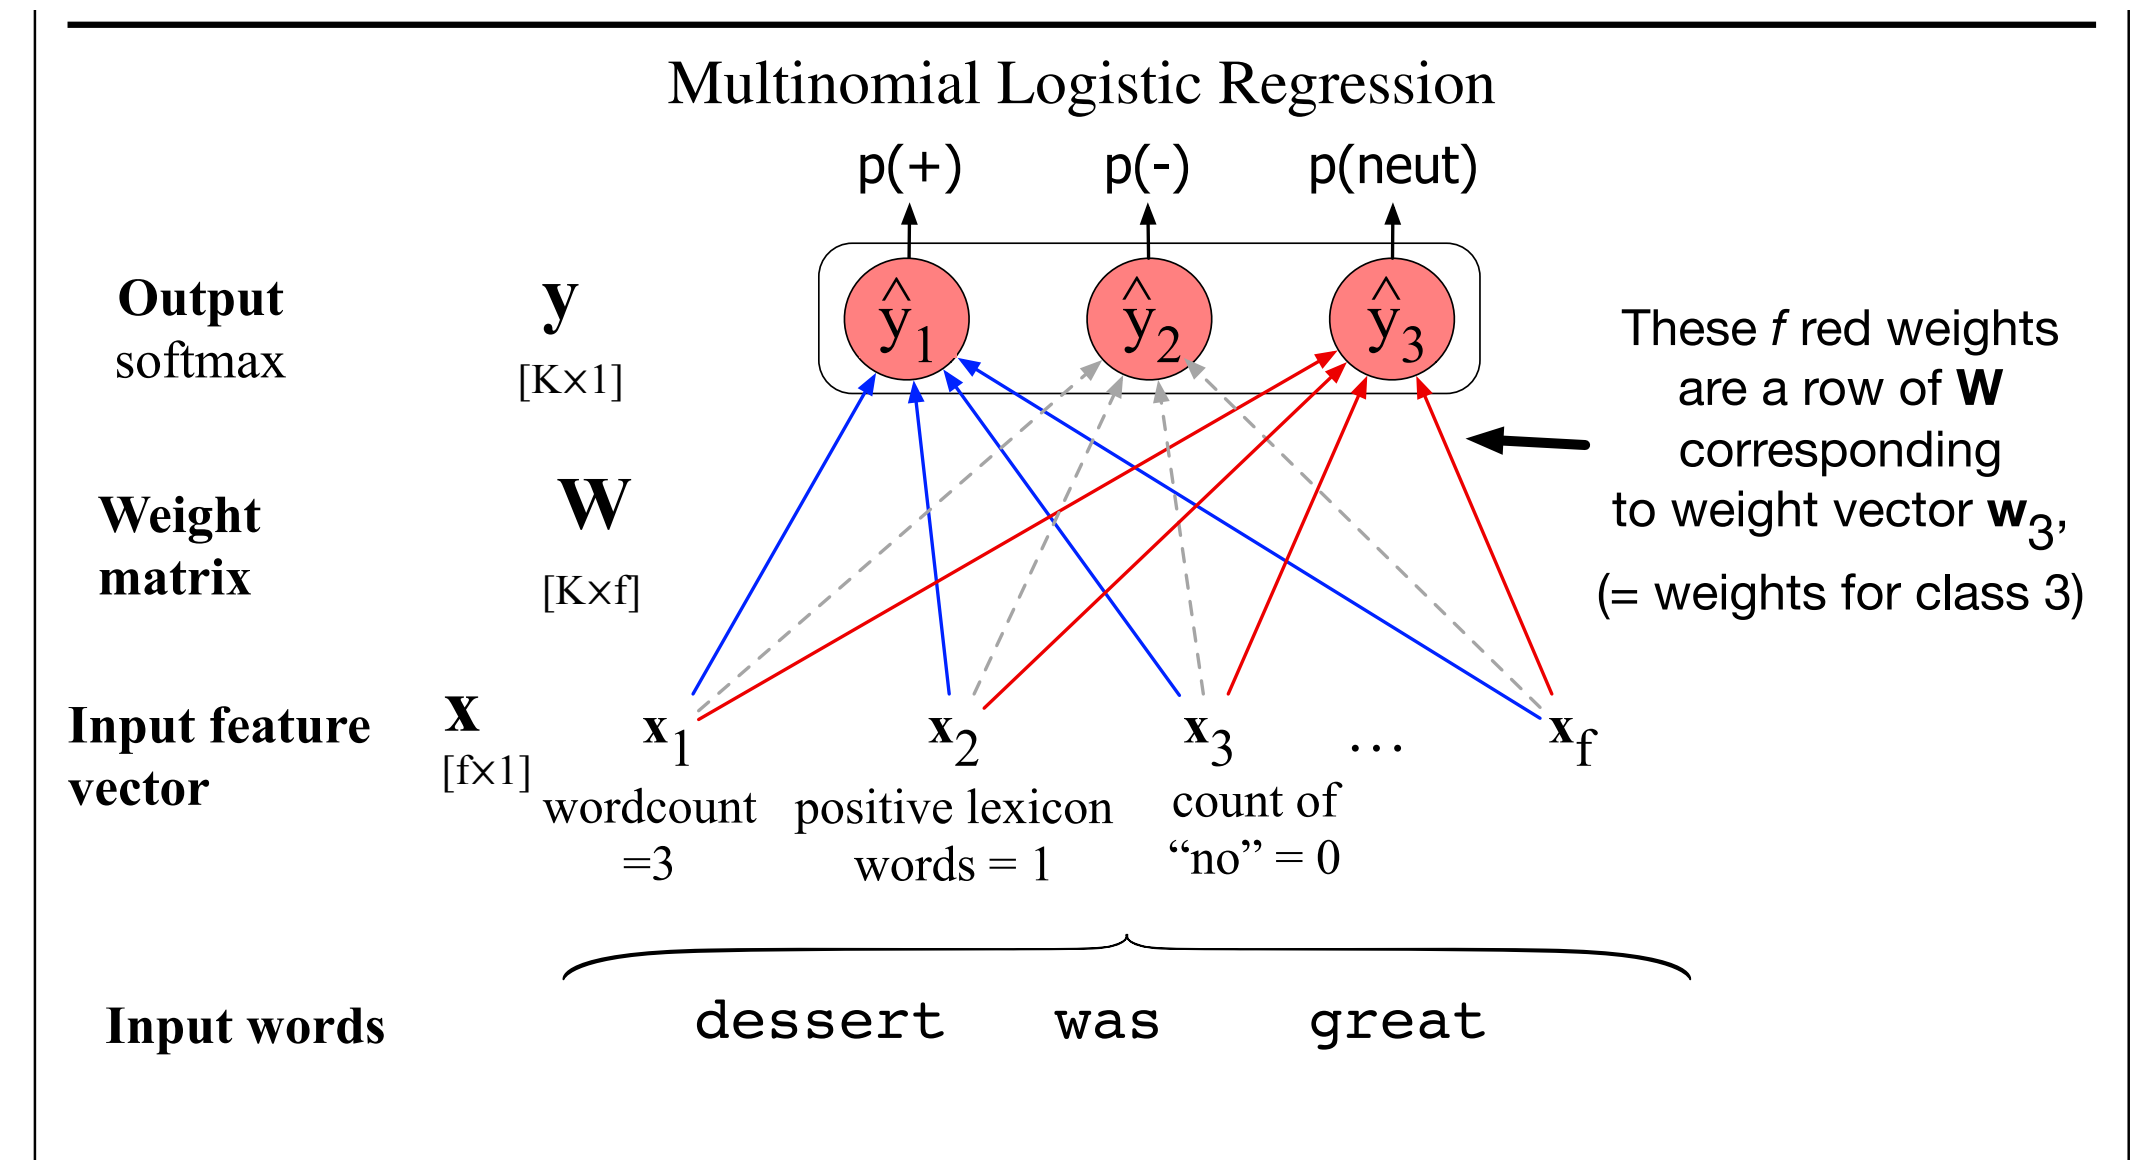
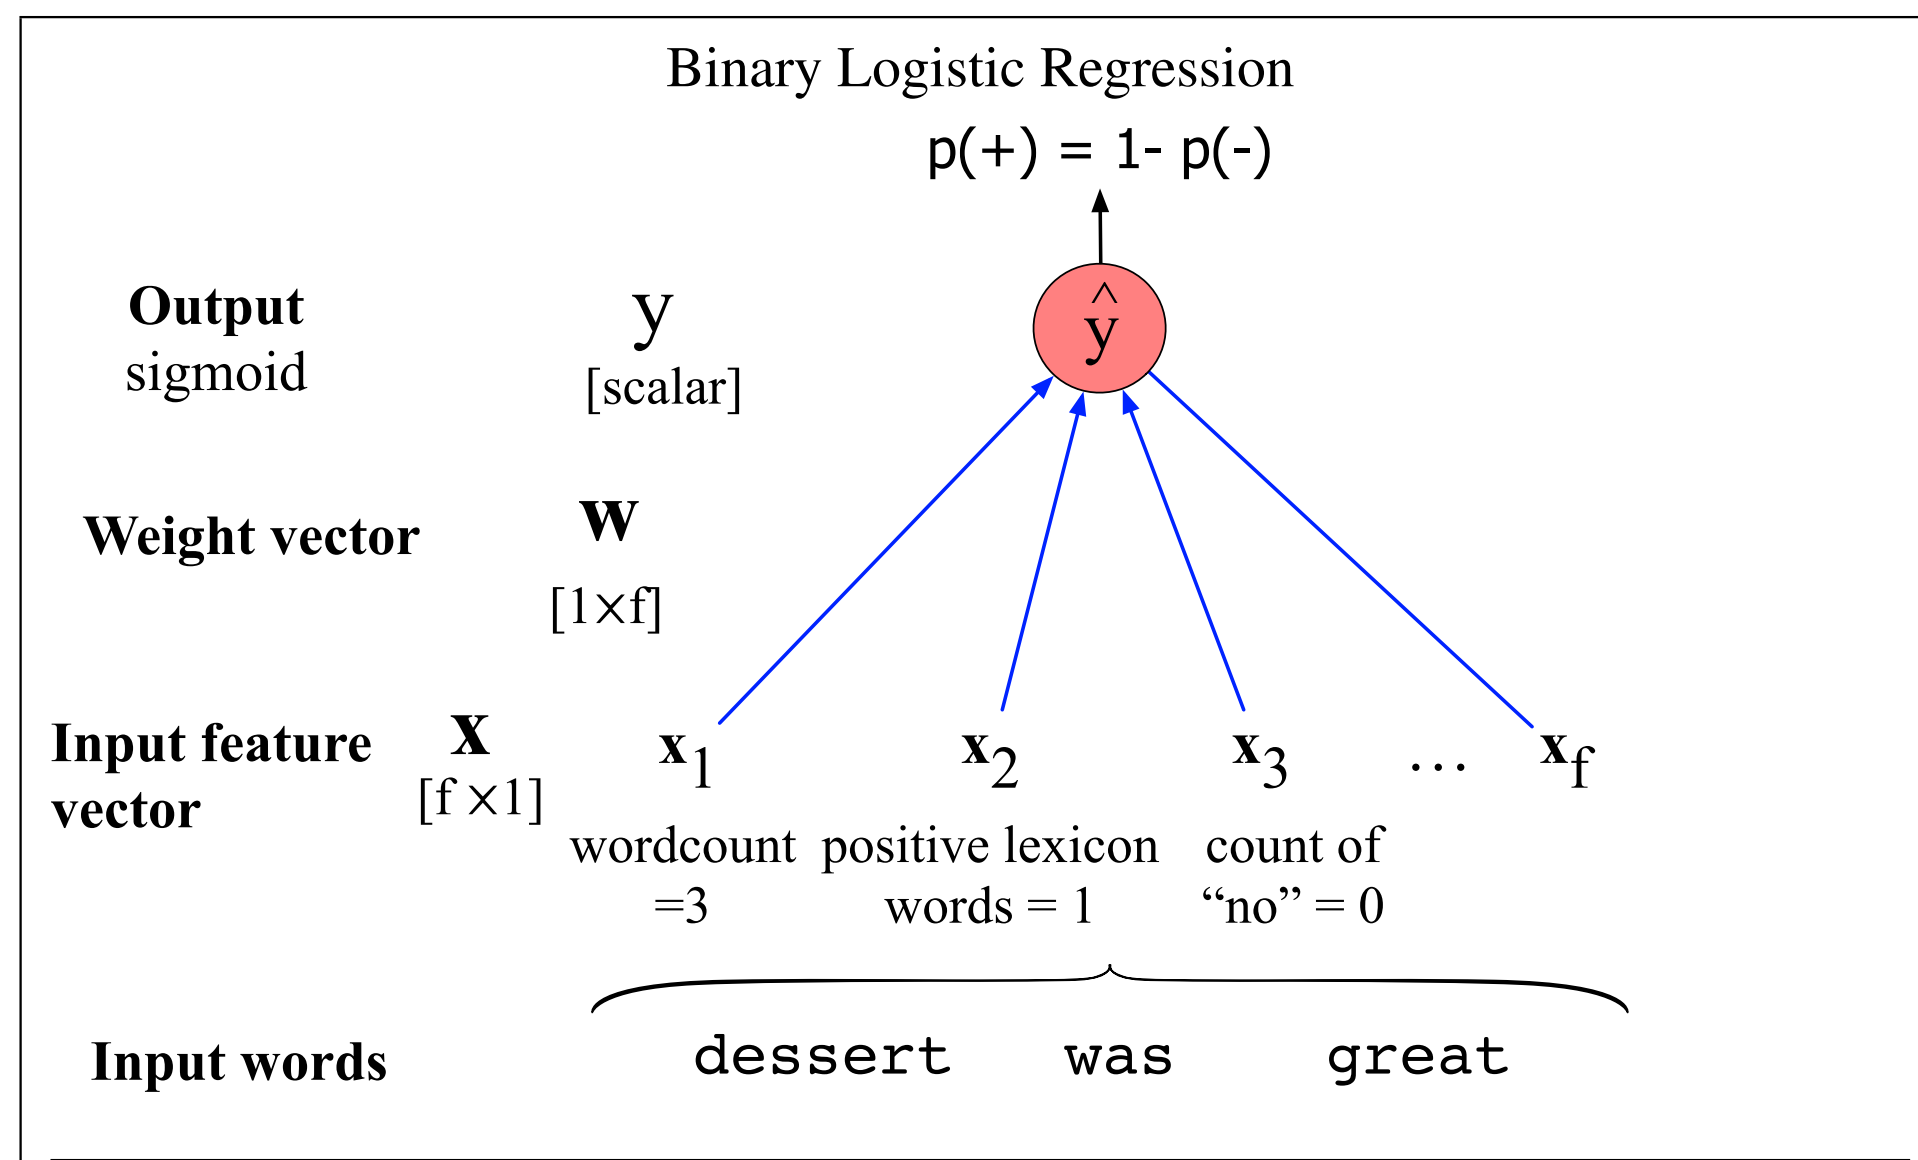
Logistic Function

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



Multiclass Logistic Regression

- 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
- Prediction: dot product for each class
- Predicted probabilities: apply the *softmax function* to normalize

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

| | | | | |
|----------|----------|----------|----------|----------|
| x_{00} | x_{01} | \dots | x_{0m} | y_0 |
| x_{10} | x_{11} | \dots | x_{1m} | y_1 |
| \vdots | \vdots | \ddots | \vdots | \vdots |
| x_{n0} | x_{n1} | \dots | x_{nm} | y_n |

n examples; xs - features; ys - class

Learning Weights

We know:

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

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

Learning Weights

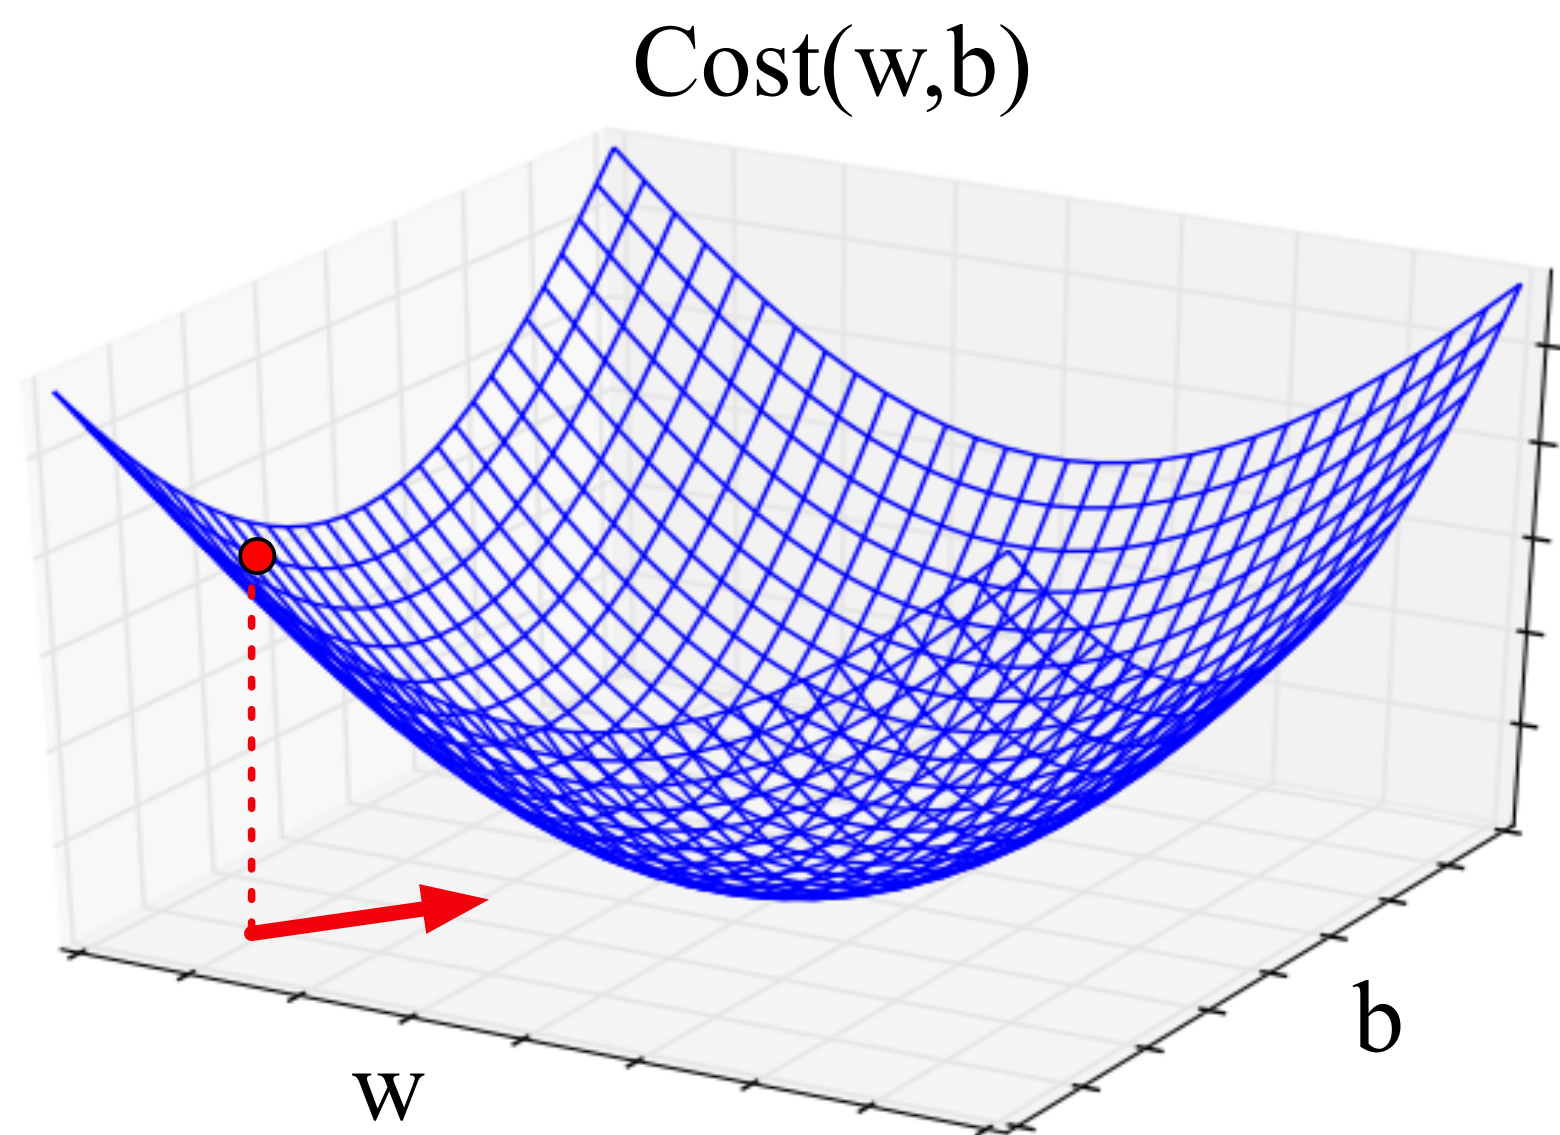
So let's try to maximize probability of the entire dataset - **maximum likelihood estimation**

$$\beta^{MLE} = \arg \max_{\beta} \log P(y_0, \dots, y_n | \mathbf{x}_0, \dots, \mathbf{x}_n; \beta)$$

Gradient ascent/descent learning

$$\beta^{MLE} = \arg \max_{\beta} \log P(y_0, \dots, y_n | \mathbf{x}_0, \dots, \mathbf{x}_n; \beta)$$

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



$\left(\frac{\partial \ell}{\partial \beta_1}, \dots, \frac{\partial \ell}{\partial \beta_J} \right)$: 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:

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

- 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



Feature vector $x = (1, \text{count "happy", count "hello", ...})$
Weights/parameters $\beta =$

"Bias term"
↓

50% prob where
 $\beta^T x = 0$

Predict $y=1$ when
 $\beta^T x > 0$

Predict $y=0$ when
 $\beta^T x \leq 0$



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?