

# CS-GY 6923: Lecture 4

## Naive Bayes Classifier, Bayesian Linear Regression + Regularization

---

NYU Tandon School of Engineering, Prof. Christopher Musco

$$(x, y) \rightarrow Pr(x, y)$$

Specific approach to machine learning:

1. Design parametric probabilistic model (usual pretty simple) that plausibly could have generated our data  $(x_1, y_1), \dots, (x_n, y_n)$ .

2. Learn unknown parameters of the model based on observed training data.

3. Given new input  $x$ , predict the label  $y$  which is most likely under the now fixed probabilistic model.

**Two strategies covered:** Maximum a posteriori (MAP) estimation and maximum likelihood estimation (MLE).

$$Pr(x, \odot)$$

# NAIVE BAYES CLASSIFIER

$\begin{pmatrix} 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ X \end{pmatrix}$

$\rightarrow$

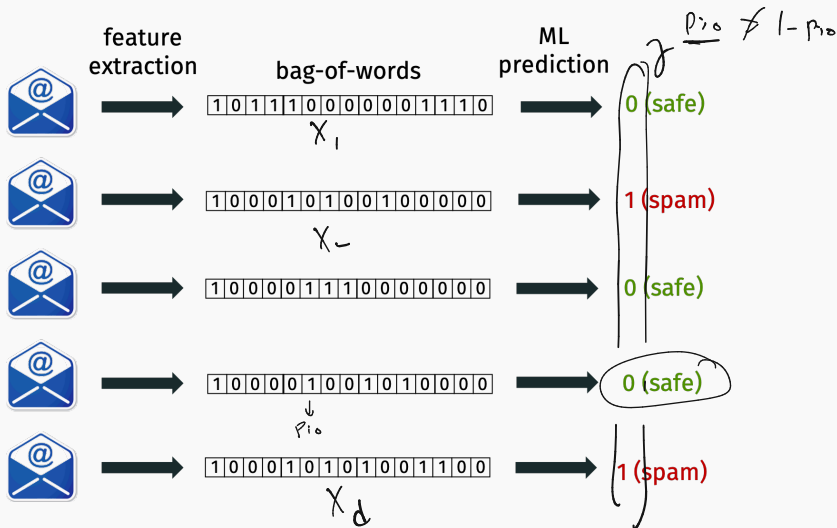
$\begin{pmatrix} 0 \\ 1 \\ 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{pmatrix} y$

Goal:

- Build a probabilistic model for a binary classification problem.
- Estimate parameters of the model.
- From the model derive a Maximum a posteriori (MAP) estimation classification rule for future predictions.  
Specifically we will see the **Naive Bayes Classifier**.

$x_1$	$x_2$	$x_3$	...	...

# SPAM PREDICTION



Both target labels and data vectors are binary.

## PROBABILISTIC MODEL FOR EMAIL


Probabilistic model for (bag-of-words, label) pair  $(\mathbf{x}, y)$ :

- Set  $y = 0$  with probability  $p_0$ ,  $y = 1$  with probability  $p_1 = 1 - p_0$ .
  - $p_0$  is probability an email is not spam (e.g. 99%).
  - $p_1$  is probability an email is spam (e.g. 1%).
- If  $y = 0$ , for each  $i$ , set  $x_i = 1$  with prob.  $p_{i0}$ .
- If  $y = 1$ , for each  $i$ , set  $x_i = 1$  with prob.  $p_{i1}$ .

Unknown model parameters:  $\lambda$  [ \_ \_ o \_ \_ \_ \_ \_ ]

$i \rightarrow$  "hello"

- $p_0, p_1$ ,
- $p_{10}, p_{20}, \dots, p_{n0}$ , one for each of the  $n$  vocabulary words.
- $p_{11}, p_{21}, \dots, p_{n1}$ , one for each of the  $n$  vocabulary words.



### How to set parameters:

- Set  $\underline{p_0}$  and  $\underline{p_1}$  to the empirical fraction of not spam/spam emails.
- For each word  $i$ , set  $\underline{p_{i0}}$  to the empirical probability word  $i$  appears in a non-spam email.
- For each word  $i$ , set  $p_{i1}$  to the empirical probability word  $i$  appears in a spam email.

Estimating these parameters from previous data examples is the only “training” we will do.

DONE WITH MODELING  
ON TO PREDICTION

## PROBABILITY REVIEW

- Probability:  $p(x)$  - the probability event  $x$  happens.  $p(x), p(y)$
- Joint probability:  $p(x,y)$  - the probability that event  $x$  and event  $y$  happen.  $p(x,y)$
- Conditional Probability  $p(x|y)$  - the probability  $x$  happens given that  $y$  happens.

$$p(x,y) = p(x|y) \cdot p(y)$$

$$p(x|y) = \frac{p(x,y)}{p(y)} = \frac{2/6}{1/2} = \frac{4}{6} = \frac{2}{3}$$

$X$  = event that dice rolls up odd

$$p(x) = 1/2$$

$Y$  = event that dice value  $\leq 3$

$$p(x|y) = \frac{2}{3}$$



## BAYES THEOREM/RULE

$$P(x|y) = \frac{P(x, y)}{P(y)} \quad P(y|x) = \frac{P(x, y)}{P(x)}$$

$$\underline{\underline{P(x|y) = \frac{P(y|x)P(x)}{P(y)}}}}$$

Proof:

$$\frac{P(x, y)}{P(y)} = \frac{P(x, y)}{\cancel{P(x)} \cdot P(y)}$$

$$P(x, y) = P(x, y)$$

↓

$$P(x) P(y|x) = P(y) \cdot P(x|y)$$



## CLASSIFICATION RULE

Given unlabeled input  $(\mathbf{w}, \_)$ , choose the label  $y \in \{0, 1\}$  which is most likely given the data. Recall  $\mathbf{w} = [0, 0, 1, \dots, 1, 0]$ .

Classification rule: maximum a posterior (MAP) estimate.

Step 1. Compute:

•  $p(y = 0 | \mathbf{w})$ : prob.  $y = 0$  given observed data vector  $\mathbf{w}$ .

•  $p(y = 1 | \mathbf{w})$ : prob.  $y = 1$  given observed data vector  $\mathbf{w}$ .

Step 2. Output: 0 or 1 depending on which probability is larger.

$p(y = 0 | \mathbf{w})$  and  $p(y = 1 | \mathbf{w})$  are called **posterior** probabilities.

prior probability  $\rightarrow p_0 \rightarrow p(y=0)$   
 $\rightarrow p_1 \rightarrow p(y=1)$

## EVALUATING THE POSTERIOR

How to compute the posterior? **Bayes rule!**

$$p(y=0 | \mathbf{w}) = \frac{p(\mathbf{w} | y=0)p(y=0)}{p(\mathbf{w})} \quad (1)$$

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}} \quad (2)$$

- **Prior:** Probability in class 0 prior to seeing any data.
- **Posterior:** Probability in class 0 after seeing the data.

## EVALUATING THE POSTERIOR

Goal is to determine which is larger:

$$\begin{aligned} p(y = 0 | \mathbf{w}) &= \frac{p(\mathbf{w} | y = 0)p(y = 0)}{p(\mathbf{w})} \\ p(y = 1 | \mathbf{w}) &= \frac{p(\mathbf{w} | y = 1)p(y = 1)}{p(\mathbf{w})} \end{aligned} \quad \text{vs.}$$

*Handwritten notes: "likelihood" with an arrow pointing to the likelihood terms in the numerators. The terms  $p(\mathbf{w})$  in the denominators are crossed out with lines.*

- We can ignore the evidence  $p(\mathbf{w})$  since it is the same for both sides!
- $p(y = 0)$  and  $p(y = 1)$  already known (computed from training data). These are our computed parameters  $p_0, p_1$ .
- $p(\mathbf{w} | y = 0)$  = ?  $p(\mathbf{w} | y = 1)$  = ?

## EVALUATING THE POSTERIOR

$$pr(w_i=0 | y=0) = 1 - \underline{pr(w_i=1 | y=0)} = \underline{1 - p_{i0}}$$

Consider the example  $w = [0, 1, 1, 0, 0, 0, 1, 0]$ .

Recall that, under our model, index  $i$  is 1 with probability  $p_{i0}$  if we are not spam, and 1 with probability  $p_{i1}$  if we are spam.

$$\begin{aligned} \underline{p(w | y=0)} &= pr(w_1=0 | y=0) pr(w_2=1 | y=0) \dots pr(w_n=0 | y=0) \\ &= \underline{(1-p_{10})} (p_{20}) (p_{30}) \dots (1-p_{n0}) \end{aligned}$$

$$\underline{p(w | y=1)} = (1-p_{11}) (p_{21}) (p_{31}) \dots (1-p_{n1})$$

$$(1-p_{10}) p_{20} p_{30} (1-p_{40}) (1-p_{50}) \dots$$

## Final Naive Bayes Classifier

Training/Modeling: Use existing data to compute:

- $p_0 = p(y = 0), p_1 = p(y = 1)$   $\arg \max(a, b)$
- For all  $i$  compute:  $\arg \max(\log(a), \log(b))$ 
  - $p_{i0} = p(x_i = 1 | y = 0)$  and  $(1 - p_{i0}) = p(x_i = 0 | y = 0)$
  - $\bar{p}_{i1} = p(x_i = 1 | y = 1)$  and  $(1 - \bar{p}_{i1}) = p(x_i = 0 | y = 1)$

Prediction:

- For new input  $w$ :
  - Compute  $p(w | y = 0) = \prod_i p(w_i | y = 0)$   $\rightarrow$  likelihood
  - Compute  $p(w | y = 1) = \prod_i p(w_i | y = 1)$
- Return


$$\arg \max [p(w | y = 0) \cdot p(y = 0), p(w | y = 1) \cdot p(y = 1)]$$

OTHER APPLICATIONS OF  
THE BAYESIAN PERSPECTIVE

# BAYESIAN REGRESSION

The Bayesian view offers an interesting alternative perspective on many machine learning techniques.

$$y = \beta + \eta$$




**Example:** Linear Regression.

**Probabilistic model:**

$$y = \langle x, \beta \rangle + \eta$$

where the  $\eta$  drawn from  $N(0, \sigma^2)$  is **random Gaussian noise**.



$$\Pr(\eta = z) \sim e^{-z^2}$$


The symbol  $\sim$  means “is proportional to”.



## GAUSSIAN DISTRIBUTION REFRESHER

Names for same thing: Normal distribution, Gaussian distribution, bell curve.

Parameterized by mean  $\mu$  and variance  $\sigma^2$ .




$\eta$  is a continuous random variable, so it has a probability density function  $p(\eta)$  with  $\int_{-\infty}^{\infty} p(\eta) d\eta = 1$

$$p(\eta) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\eta-\mu}{\sigma}\right)^2}$$

$\psi$   
normalizing

## GAUSSIAN DISTRIBUTION REFRESHER

The important thing to remember is that the the PDF falls off exponentially as we move further from the mean.




The normalizing constant in front  $1/2$ , etc. don't matter so much.

## QUICK CHECK

Example: Linear Regression.

Probabilistic model:



$$\underline{y = \langle x, \beta \rangle + \eta}$$

where the  $\eta$  drawn from  $N(0, \sigma^2)$  is **random Gaussian noise**.

The noise is independent for different inputs  $x_1, \dots, x_n$ .

If we knew  $\beta$  what is the maximum a posteriori (MAP) estimate for  $y$  given new data observation  $x$ ?

i.e. what value of  $y$  maximizes  $\Pr(\underline{y} \mid \underline{x})$ ?

$$y = \underbrace{\langle x, \beta \rangle}_{\text{known}} + \underbrace{\mu}_{\substack{\rightarrow \text{mean of } \eta \\ = 0}}$$

## How should we select $\beta$ for our model?

Also use a Bayesian approach!

Choose  $\beta$  to maximize:

$$\text{posterior} = \Pr(\beta | X, y) = \frac{\Pr(X, y | \beta) \Pr(\beta)}{\Pr(X, y)} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

$\downarrow$   
 training set

In this case, we don't have a prior – no values of  $\beta$  are inherently more likely than others.


Choose  $\beta$  to maximize just the likelihood:

$$\frac{\Pr(X, y | \beta) \Pr(\beta)}{\Pr(X, y)} = \frac{\text{likelihood} \times \text{prior}}{\text{evidence}}$$

This is called the maximum likelihood estimate.

# MAXIMUM LIKELIHOOD ESTIMATE

Data:



**Model:**  $y_i = \langle x_i, \beta \rangle + \eta_i$  where  $p(\eta_i = z) \sim e^{-z^2/2\sigma^2}$  and  $\eta_1, \dots, \eta_n$  are independent.

$$\Pr(y | \beta) \sim \Pr(x_1, y_1 | \beta) \cdot \Pr(x_2, y_2 | \beta) \cdots \Pr(x_n, y_n | \beta)$$

$$\downarrow \qquad \qquad \qquad \downarrow$$

$$e^{-\frac{1}{2} (y_1 - \langle x_1, \beta \rangle)^2 / 26^2} \cdots e^{-\frac{1}{2} (y_n - \langle x_n, \beta \rangle)^2 / 26^2}$$

$$\mu = 0$$

Easier to work with the log likelihood:

$$\arg \max_{\beta} \Pr(X, y | \beta) = \arg \max_{\beta} \prod_{i=1}^n e^{-(y_i - \langle x_i, \beta \rangle)^2 / 2\sigma^2}$$


$$= \arg \max_{\beta} \log \left( \prod_{i=1}^n e^{-(y_i - \langle x_i, \beta \rangle)^2 / 2\sigma^2} \right)$$

$$= \arg \max_{\beta} \sum_{i=1}^n \underbrace{-\frac{(y_i - \langle x_i, \beta \rangle)^2}{2\sigma^2}}_{\rightarrow \log_2(e^{-(y_i - \langle x_i, \beta \rangle)^2 / 2\sigma^2})}$$

$$= \arg \min_{\beta} \sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2$$



$$\|X\beta - y\|_2^2$$



$$\min \|y - X\beta\|_2^2$$

Conclusion: Choose  $\beta$  to minimize:

$$\beta = (X^T X)^{-1} X^T z$$


$$\sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2 = \underbrace{\|y - X\beta\|_2^2}.$$

**This is a completely different justification for squared loss!**

Minimizing the  $\ell_2$  loss is optimal in a certain sense when you assume your data follows a linear model with i.i.d. Gaussian noise.

## BAYESIAN REGRESSION

If we had modeled our noise  $\eta$  as Laplace noise, we would have found that minimizing  $\|\mathbf{y} - \mathbf{X}\beta\|_1$  was optimal.



$$Pr(\eta = z) \sim e^{-|z|} \quad e^{-z^2} \rightarrow \text{gaussian}$$

Laplace noise has “heavier tails”, meaning that it results in more outliers.

This is a completely different justification for  $l_1$  loss.



We can add another layer of probabilistic modeling by also assuming  $\beta$  is random and comes from some distribution, which encodes our prior belief on what the parameters are.

Return to Maximum a posteriori (MAP estimation):

$$\Pr(\underline{\beta} | \underline{X}, \underline{y}) = \frac{\Pr(\underline{X}, \underline{y} | \underline{\beta}) \Pr(\underline{\beta})}{\Pr(\underline{X}, \underline{y})}$$

Assume values in  $\underline{\beta} = [\underline{\beta}_1, \dots, \underline{\beta}_d]$  come from some distribution.

- **Common model:** Each  $\underline{\beta}_i$  drawn from  $\underline{N}(0, \gamma^2)$ , i.e. normally distributed, independent.
- Encodes a belief that we are unlikely to see models with very large coefficients.

Goal: choose  $\beta$  to maximize:

$$\beta^{\text{opt}} = \underset{\beta}{\text{argmax}} \Pr(\beta | X, y) = \frac{\Pr(X, y | \beta) \Pr(\beta)}{\cancel{\Pr(X, y)}}$$

- We can still ignore the “evidence” term  $\Pr(X, y)$  since it is a constant that does not depend on  $\beta$ .
- $\Pr(\beta)$  =  $\Pr(\beta_1)$  ·  $\Pr(\beta_2)$  · ... ·  $\Pr(\beta_d)$
- $\Pr(\beta) \sim \prod_{i=1}^d e^{-\frac{1}{2} \beta_i^2 / \sigma^2}$

# BAYESIAN REGULARIZATION


Easier to work with the **log likelihood**:

$$\begin{aligned} & \arg \max_{\beta} \Pr(X, y | \beta) \cdot \Pr(\beta) \\ &= \arg \max_{\beta} \prod_{i=1}^n e^{-(y_i - \langle x_i, \beta \rangle)^2 / 2\sigma^2} \cdot \left( \prod_{i=1}^d e^{-(\beta_i)^2 / 2\gamma^2} \right) \\ &= \arg \max_{\beta} \sum_{i=1}^n -\frac{(y_i - \langle x_i, \beta \rangle)^2}{2\sigma^2} + \left( \sum_{i=1}^d -\frac{(\beta_i)^2}{2\gamma^2} \right) \\ &= \arg \min_{\beta} \left( \sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2 + \frac{\sigma^2}{\gamma^2} \sum_{i=1}^d (\beta_i)^2 \right) \end{aligned}$$

Choose  $\beta$  to minimize  $\|y - X\beta\|_2^2 + \frac{\sigma^2}{\gamma^2} \|\beta\|_2^2$ .

Completely different justification for ridge regularization!

$\lambda$



**Test your intuition:** What modeling assumption justifies LASSO



regularization:  $\min_{\beta} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1$ ?

# LINEAR CLASSIFICATION

## MOTIVATING PROBLEM

**Breast Cancer Biopsy:** Determine if a breast lump in a patient is malignant (cancerous) or benign (safe).

- Collect cells from lump using fine needle biopsy.
- Stain and examine cells under microscope.
- Based on certain characteristics (shape, size, cohesion) determine if likely malignant or not).




Demo: `demo_breast_cancer.ipynb`

Data: UCI machine learning repository

## Breast Cancer Wisconsin (Original) Data Set

Download: [Data Folder](#), [Data Set Description](#)

**Abstract:** Original Wisconsin Breast Cancer Database



<b>Data Set Characteristics:</b>	Multivariate	<b>Number of Instances:</b>	699	<b>Area:</b>	Life
<b>Attribute Characteristics:</b>	Integer	<b>Number of Attributes:</b>	10	<b>Date Donated</b>	1992-07-15
<b>Associated Tasks:</b>	Classification	<b>Missing Values?</b>	Yes	<b>Number of Web Hits:</b>	564320

`https://archive.ics.uci.edu/ml/datasets/breast+cancer+wisconsin+(original)`

**Features:** 10 numerical scores about cell characteristics (Clump Thickness, Uniformity, Marginal Adhesion, etc.)

## MOTIVATING PROBLEM

**Data:**  $(x_1, y_1), \dots, (x_n, y_n)$ .

$x_i \in [1, 5, 4, \dots, 2]$  contains score values.

Label  $y_i \in \{0, 1\}$  is 0 if benign cells, 1 if malignant cells.

**Goal:** Based on scores (which would be collected manually, or even learned on their own using an ML algorithm) predict if a sample of cells is malignant or benign.

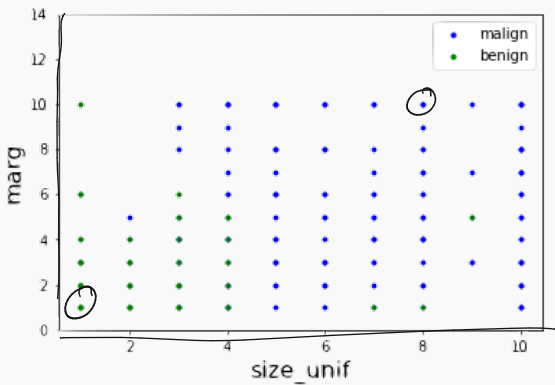
**Approach:**

- Naive Bayes Classifier can be extended to  $\mathbf{x}$  with numerical values (instead of binary values as seen before). Will see on homework.
- **Today:** Learn a different type of classifier.



## BEGIN BY PLOTTING DATA

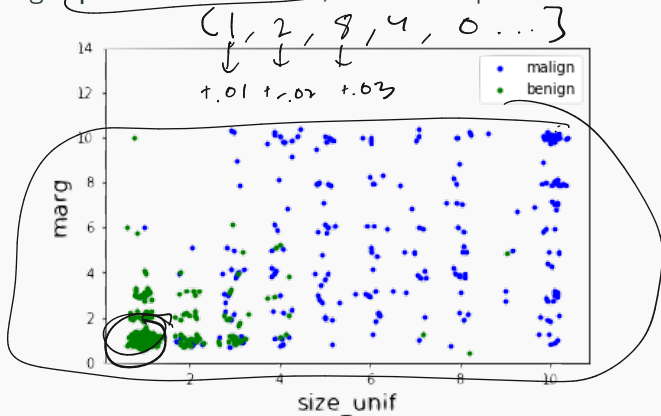
We pick two variables, Margin Adhesion and Size Uniformity and plot a scatter plot. Points with label 1 (malignant) are plotted in blue, those with label 2 (benign) are plotted in green.



Lots of overlapping points! Hard to get a sense of the data.


## PLOTTING WITH JITTER

Simple + Useful Trick: data jittering. Add tiny random noise (using e.g. `np.random.randn()`) to data to prevent overlap.



Noise is only for plotting. It is not added to the data for training, testing, etc.

Any ideas for possible classification rules for this data?

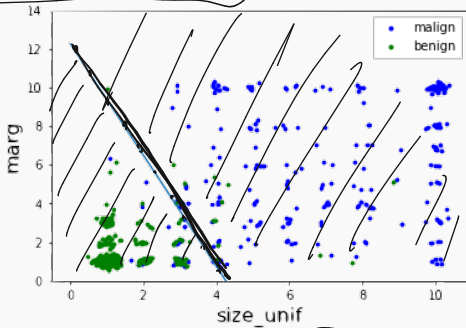


# LINEAR CLASSIFIER

Given vector of predictors  $\mathbf{x}_i \in \mathbb{R}^d$  (here  $d = 2$ ) find a parameter vector  $\boldsymbol{\beta} \in \mathbb{R}^d$  and threshold  $\lambda$ .

- Predict  $y_i = 0$  if  $\langle \mathbf{x}_i, \boldsymbol{\beta} \rangle \leq \lambda$ .
- Predict  $y_i = 1$  if  $\langle \mathbf{x}_i, \boldsymbol{\beta} \rangle > \lambda$

[ margin, size\_unif ]  $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$



Line has equation  $\langle \mathbf{x}, \boldsymbol{\beta} \rangle = \lambda$ .


## LINEAR CLASSIFIER

As long as we append a 1 onto each data vector  $\mathbf{x}_i$  (i.e. a column of ones onto the data matrix  $\mathbf{X}$ ) like we did for linear regression, an equivalent function is:

- Predict  $y_i = 0$  if  $\langle \mathbf{x}_i, \boldsymbol{\beta} \rangle \leq 0$ .
- Predict  $y_i = 1$  if  $\langle \mathbf{x}_i, \boldsymbol{\beta} \rangle \geq 0$

[w<sub>0</sub>, size, 1]

$\beta_0$   
 $\beta_1$   
 $\beta_2$



Line has equation  $\langle \mathbf{x}, \boldsymbol{\beta} \rangle = 0$ .

## 0 - 1 LOSS

**Question:** How do we find a good linear classifier automatically?

Loss minimization approach (first attempt):

- Model<sup>1</sup>:

$$f_{\beta}(x) = \mathbb{1}[\langle x, \beta \rangle > 0]$$

- Loss function: "0 - 1 Loss"

$$L(\beta) = \sum_{i=1}^n |f_{\beta}(x_i) - y_i|$$


0, 1

# of training examples

---

<sup>1</sup> $\mathbb{1}[\text{event}]$  is the indicator function: it evaluates to 1 if the argument inside is true, 0 if false.

## Problem with 0 – 1 loss:



- The loss function  $L(\beta)$  is not differentiable because  $f_{\beta}(\mathbf{x})$  is discontinuous.
- Impossible to take the gradient, very hard to minimize loss to find optimal  $\beta$ .
- Non-convex function (will make more sense next lecture).


Loss minimization approach (second attempt):

- Model:

$$\underline{f_{\beta}(x)} = \mathbb{1}[\underline{\langle x, \beta \rangle} > \underline{1/2}]$$

- Loss function: "Square Loss"

$$\text{min}_{\beta} \quad L(\beta) = \sum_{i=1}^n (\langle x, \beta \rangle - y_i)^2$$




Intuitively tries to make  $\langle x, \beta \rangle$  close to 0 for examples in class 0, close to 1 for examples in class 1.



## LINEAR CLASSIFIER VIA SQUARE LOSS


We can solve for  $\beta$  by just solving a least squares multiple linear regression problem.



Do you see any issues here?

## LINEAR CLASSIFIER VIA SQUARE LOSS

Problem with square loss:




- Loss increases if  $\langle \mathbf{x}, \boldsymbol{\beta} \rangle > 1$  even if correct label is 1. Or if  $\langle \mathbf{x}, \boldsymbol{\beta} \rangle < 0$  even if correct label is 0.
- Intuitively we don't want to "punish" these cases.

# LOGISTIC REGRESSION

Let  $h_{\beta}(x)$  be the **logistic function**:

$$h_{\beta}(x) = \frac{1}{1 + e^{-\langle \beta, x \rangle}}$$

$$\frac{1}{1 + e^0} = \frac{1}{2}$$



Loss minimization approach (this works!):

- Model: Let  $h_{\beta}(x) = \frac{1}{1+e^{-\langle \beta, x \rangle}}$

$\langle x, \beta \rangle > 0$

$$f_{\beta}(x) = \mathbb{1}[\underline{h_{\beta}(x)} > 1/2]$$

- Loss function: “Logistic loss” aka “binary cross-entropy loss”

$$L(\beta) = \sum_{i=1}^n \left( y_i \log(h_{\beta}(x_i)) + (1 - y_i) \log(1 - h_{\beta}(x_i)) \right)$$


$y_i = 0$

$y_i = 1 - \log(h_{\beta}(x))$

# LOGISTIC LOSS

Logistic Loss:


$$L(\beta) = - \sum_{i=1}^n y_i \log(h_{\beta}(x)) + (1 - y_i) \log(1 - h_{\beta}(x))$$



# LOGISTIC LOSS


Logistic Loss:

$$L(\beta) = - \sum_{i=1}^n y_i \log(h_{\beta}(x)) + (1 - y_i) \log(1 - h_{\beta}(x))$$



## LOGISTIC LOSS

- Convex function in  $\beta$ , can be minimized using gradient descent (next lecture).
- Works well in practice.
- Good Bayesian motivation: see posted lecture notes if you are interested.



Fit using logistic regression / log loss.

Once we have a classification algorithm, how do we judge its performance?


- **Simplest answer:** Error rate = fraction of data examples misclassified in test set.
- What are some issues with this approach?



# ERROR IN CLASSIFICATION

- **Precision:** Fraction of positively labeled examples (label 1) which are correct.
- **Recall:** Fraction of true positives that we labeled correctly with label 1.

**Question:** Which should we optimize for medical diagnosis?



How many selected items are relevant?

$$\text{Precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}$$

How many relevant items are selected?

$$\text{Recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}$$

### Logistic regression workflow:

- Select  $\beta$  via training and compute  $h_{\beta}(\mathbf{x}_i) = \frac{1}{1+e^{-\langle \mathbf{x}_i, \beta \rangle}}$  for all  $\mathbf{x}_i$ .
- Predict  $y_i = 0$  if  $h_{\beta}(\mathbf{x}_i) \leq \lambda$ ,  $y_i = 1$  if  $h_{\beta}(\mathbf{x}_i) > \lambda$ .
- Default value of  $\lambda$  is  $1/2$ . Increasing  $\lambda$  improves precision. Decreasing  $\lambda$  improves recall.


Once we have a classification algorithm, how do we judge its performance?

- **Simplest answer:** Error rate = fraction of data examples misclassified in test set.
- What are some issues with this approach?

# ERROR IN CLASSIFICATION

- **Precision:** Fraction of positively labeled examples (label 1) which are correct.
- **Recall:** Fraction of true positives that we labeled correctly with label 1.

**Question:** Which should we optimize for medical diagnosis?



How many selected items are relevant?

$$\text{Precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}}$$

How many relevant items are selected?

$$\text{Recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}}$$

### Possible logistic regression workflow:

- Learn  $\vec{\beta}$  and compute  $h_{\vec{\beta}}(\vec{x}_i) = \frac{1}{1+e^{-\langle \vec{x}_i, \vec{\beta} \rangle}}$  for all  $\vec{x}_i$ .
- Predict  $y_i = 0$  if  $h_{\vec{\beta}}(\vec{x}_i) \leq \lambda$ ,  $y_i = 1$  if  $h_{\vec{\beta}}(\vec{x}_i) > \lambda$ .
- Default value of  $\lambda$  is  $1/2$ . Increasing  $\lambda$  improves precision.  
Decreasing  $\lambda$  improves recall.

This is very heuristic. There are other methods for handling “class imbalance” which can often lead to good overall error, but poor precision or recall. Techniques include weighting the loss function to care more about false negatives, or subsampling the larger class.

What about when  $y \in \{1, \dots, q\}$  instead of  $y \in \{0, 1\}$

**Two options for multiclass data:**

- One-vs.-all (most common, also called one-vs.-rest)
- One-vs.-one (slower, but can be more effective)


**In both cases, we convert to multiple binary classification problem.**

# ONE VS. REST



- For  $q$  classes train  $q$  classifiers. Obtain parameters  $\vec{\beta}_1, \dots, \vec{\beta}_q$ .
- Assign  $y$  to class  $i$  with maximum  $\langle \vec{\beta}_i, \vec{x} \rangle$ .


# ONE VS. REST



- For  $q$  classes train  $\frac{q(q-1)}{2}$  classifiers.
- Assign  $y$  to class which  $i$  which wins in the most number of head-to-head comparisons.



Hard case for one-vs.-all.



- One-vs.-one would be a better choice here.
- Also tends to work better when there is class in balance.

We will see one more very common approach related to one-vs-rest either next lecture or when we do neural nets.

## ERROR IN (MULTICLASS) CLASSIFICATION

Confusion matrix for  $k$  classes:

Pred-->	1	2	...	K
Real↓				
1				
2				
...				
K				

- Entry  $i, j$  is the fraction of class  $i$  items classified as class  $j$ .
- Useful to see whole matrix to visualize where errors occur.