CS 378 Lecture 4

Classification 3: Multiclass

"Now! ... *That* should clear up a few things around here!"

Announcements

- A1 due in one week
- Social impact response: in-class Tues / open for a week

Recap: Logistic regression

$$P(y=+1|\bar{x}) = \frac{e^{\bar{w}^T f(\bar{x})}}{1+e^{\bar{w}^T f(\bar{x})}}$$

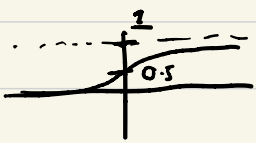

Train to maximize likelihood of data:

$$\max_{\bar{w}} \prod_{i=1}^{D} P(y=y^{(i)}|\bar{x}^{(i)}) \Rightarrow \min_{\bar{w}} \sum_{i=1}^{D} -\log P(y=y^{(i)}|\bar{x}^{(i)})$$

Update from SGD

$$\bar{w} \leftarrow \bar{w} + \alpha f(\bar{x}^{(i)}) \left(1 - P(y=+1 \mid \bar{x}^{(i)})\right)$$

if $y^{(i)} = +1$

$$\bar{w} \leftarrow \bar{w} - \alpha f(\bar{x}^{(i)}) \left(1 - P(y=-1 \mid \bar{x}^{(i)})\right)$$

if $y^{(i)} = -1$

Today

- Multiclass classification
- MC Perceptron / LR

Multiclass Classification

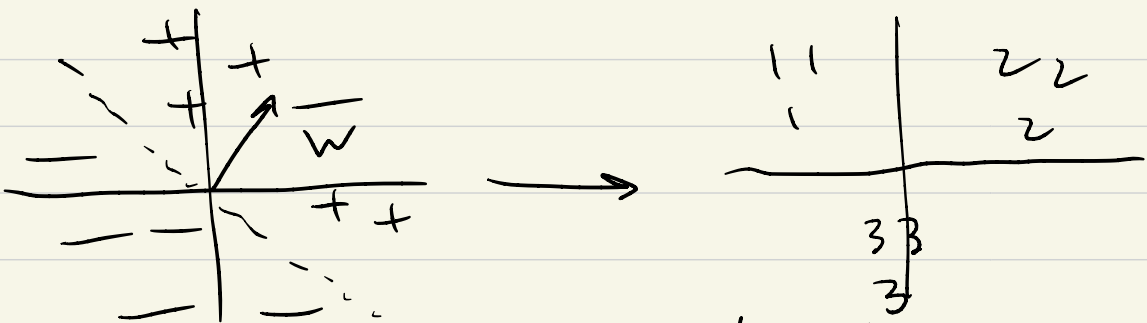

Output space $Y = \{1, 2, 3\}$

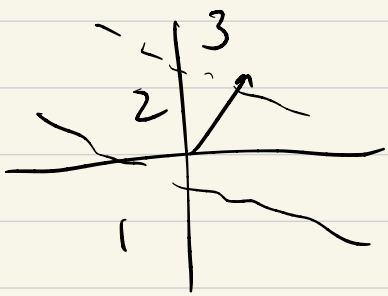

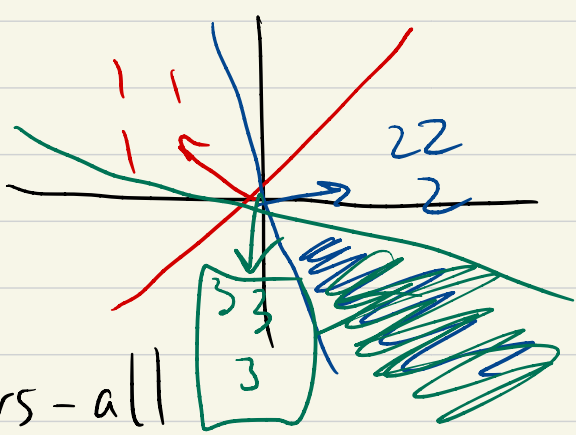

one-vs-all

m classifiers

+2 +3

??

Reconciling multiple positives:

- Ordering (1 > 2 > 3)
- Further partition
- Prefer closest vector (highest dot product)

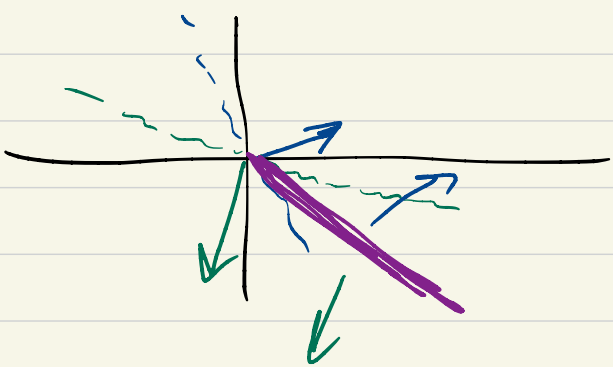

MC classification: like one-vs-all, but trained differently

Two ways of thinking:

① Different weights per class

② Different features per class (one weight vector)

Different weights

$\bar{w}_1$ $\bar{w}_2$ $\bar{w}_3$ weight vectors per class

$$\text{pred label} = \underset{y \in Y}{\text{argmax}} \; \bar{w}_y^T f(\bar{x})$$

loop over y

dot prod. w/ features

drug patients baseball

Ex $f(\bar{x}) = [1, 1, 0]$

$y = \{\text{health}, \text{sports}, \text{science}\}$

$\bar{w}_1 = [2, 5.6, -3]$

$\bar{w}_2 = [1.2, -3.1, 5.7]$

$\bar{w}_3 = [1, 1.2, -0.5]$

score for class 1 = 7.6

2 = -1.9

3 = 2.2

return 1

Multiclass Perceptron

for t in epochs

for i in data

$f(\bar{x}^{(i)})$ feats, $y^{(i)} \in \mathcal{Y}$ label

$$y_{pred} = \underset{y \in \mathcal{Y}}{\arg\max}\ \bar{w}_y^T f(\bar{x})$$

if $y_{pred} \neq y^{(i)}$:

$$\bar{w}_{y^{(i)}} \leftarrow \bar{w}_{y^{(i)}} + \alpha f(\bar{x})$$

$$\bar{w}_{y_{pred}} \leftarrow \bar{w}_{y_{pred}} - \alpha f(\bar{x})$$

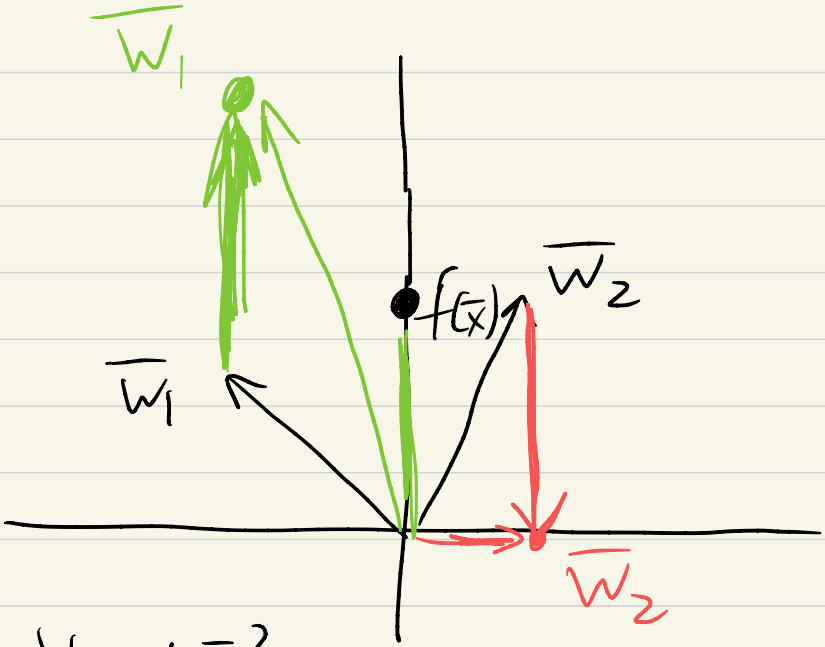

$y_{pred} = 2$

$y^{(i)} = 1$

Multiclass Logistic Regression

$$P(y = \hat{y} \mid \bar{x}) = \frac{e^{\bar{w}_{\hat{y}}^T f(\bar{x})}}{\sum_{y' \in y} e^{\bar{w}_{y'}^T f(\bar{x})}}$$

distribution over classes

$$P(y = \text{class } 1 \mid \bar{x}) = \frac{e^{\bar{w}_1^T f(\bar{x})}}{e^{\bar{w}_1^T f(\bar{x})} + e^{\bar{w}_2^T f(\bar{x})} + e^{\bar{w}_3^T f(\bar{x})}}$$

Update:

For $y^{(i)}$

$$\bar{w}_{y^{(i)}} \leftarrow \bar{w}_{y^{(i)}} + \alpha f(\bar{x}^{(i)})\left(1 - P(y = y^{(i)} \mid \bar{x})\right)$$

For all other y':

$$\bar{w}_{y'} \leftarrow \bar{w}_{y'} - \alpha f(\bar{x}^{(i)})\left(P(y = y' \mid \bar{x})\right)$$

1, 2, 3: $[0.1, 0.8, 0.1]$ $y^{(i)} = 1$

$\bar{w}_2$ "subtract" $0.8 f(\bar{x})$

$\bar{w}_3$ "subtract" $0.1 f(\bar{x})$ $\bar{w}_1$ add $0.9 f(\bar{x})$

Instapoll:

$(\bar{x}, y) = ([1\ 1\ 0], 1)$

$\bar{w}_1 = \bar{w}_2 = \bar{w}_3 = [0\ 0\ 0]$

① Probs $= [1/3 \quad 1/3 \quad 1/3]$

② After one update:

$\bar{w}_1$ $[2/3 \quad 2/3 \quad 0]$

$\bar{w}_2$ $[-1/3 \quad -1/3 \quad 0]$

$\bar{w}_3$ $[-1/3 \quad -1/3 \quad 0]$