Georgia Tech

Naïve Bayes and Logistic Regression

Mahdi Roozbahani
Georgia Tech

These slides are inspired based on slides from Le Song, Eric Eaton, and Chao Zhang.


THE BEST WAY TO
EXPLAIN OVERFITTING

Outline

- Generative and Discriminative Classification
- The Logistic Regression Model
- Understanding the Objective Function
- Gradient Descent for Parameter Learning
- Multiclass Logistic Regression

Classification

- Represent the data

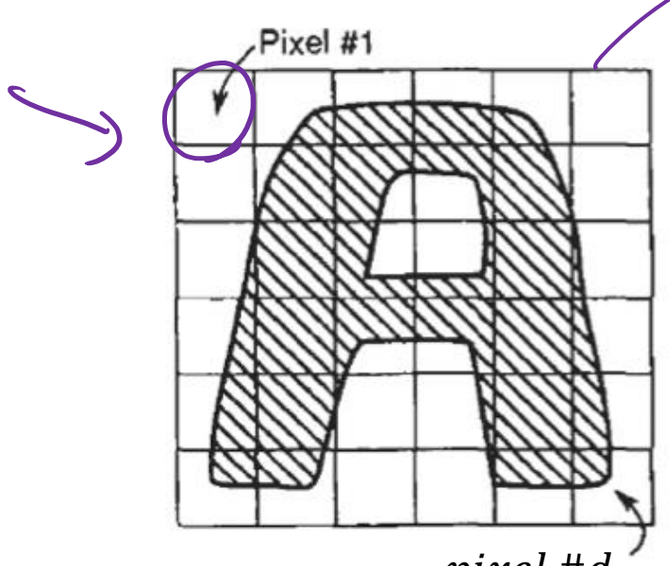

$$\Rightarrow \quad X = \begin{bmatrix} x_1 \\ x_2 \\ \cdot \\ \cdot \\ \cdot \\ x_d \end{bmatrix}$$

- A label is provided for each data point, eg., $y \in \{-1, +1\}$

- Classifier

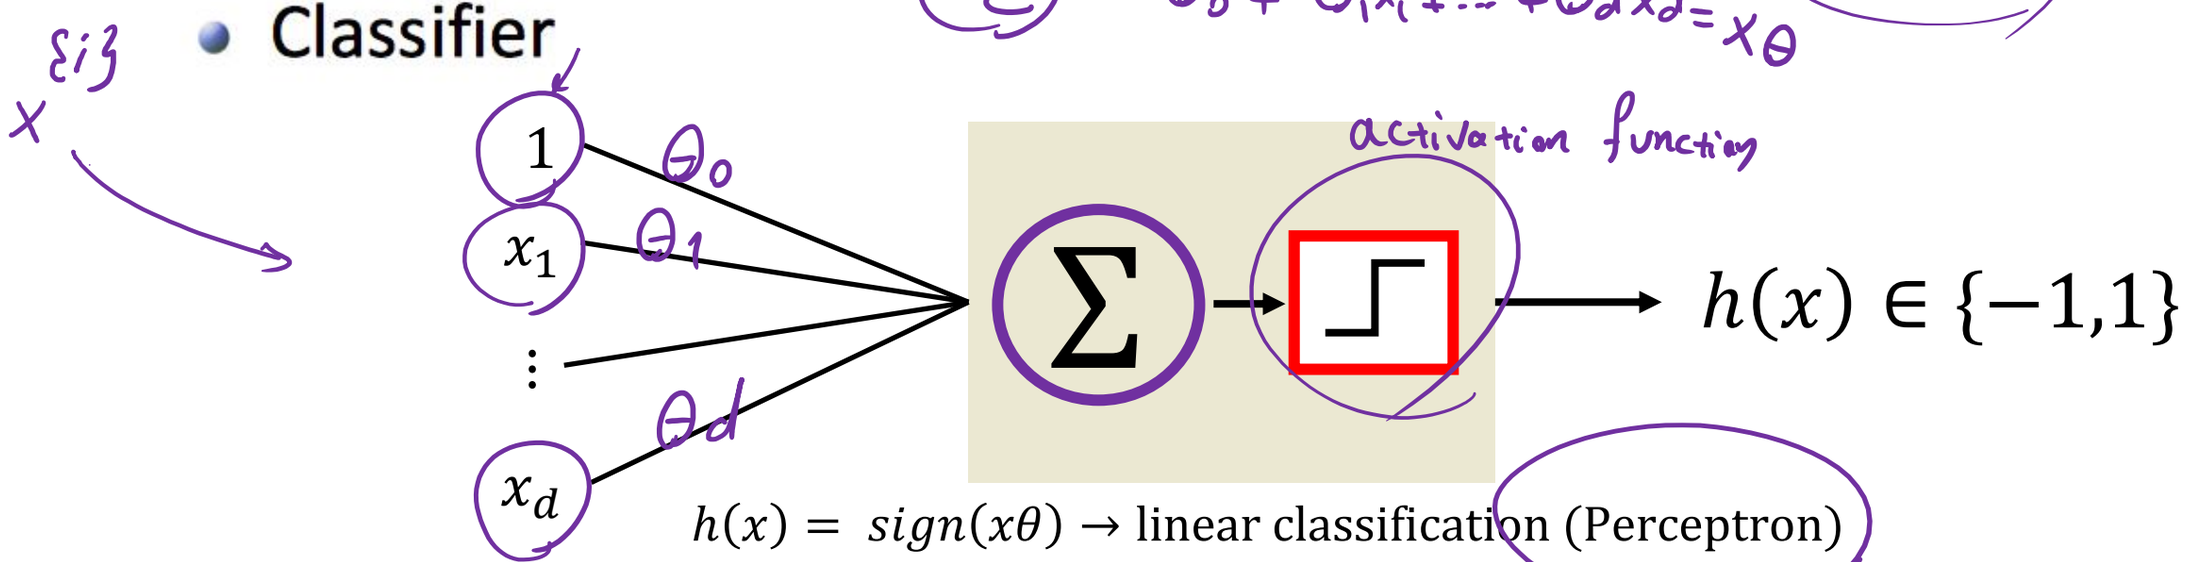

$h(x) = sign(x\theta) \rightarrow$ linear classification (Perceptron)

Decision Making: Dividing the Feature Space

- Distributions of sample from normal (positive class) and abnormal (negative class) tissues

$X = \begin{bmatrix} h & w \\ & \\ & \end{bmatrix}$ $Y = \begin{bmatrix} cat \\ \vdots \\ dog \\ \vdots \end{bmatrix}$

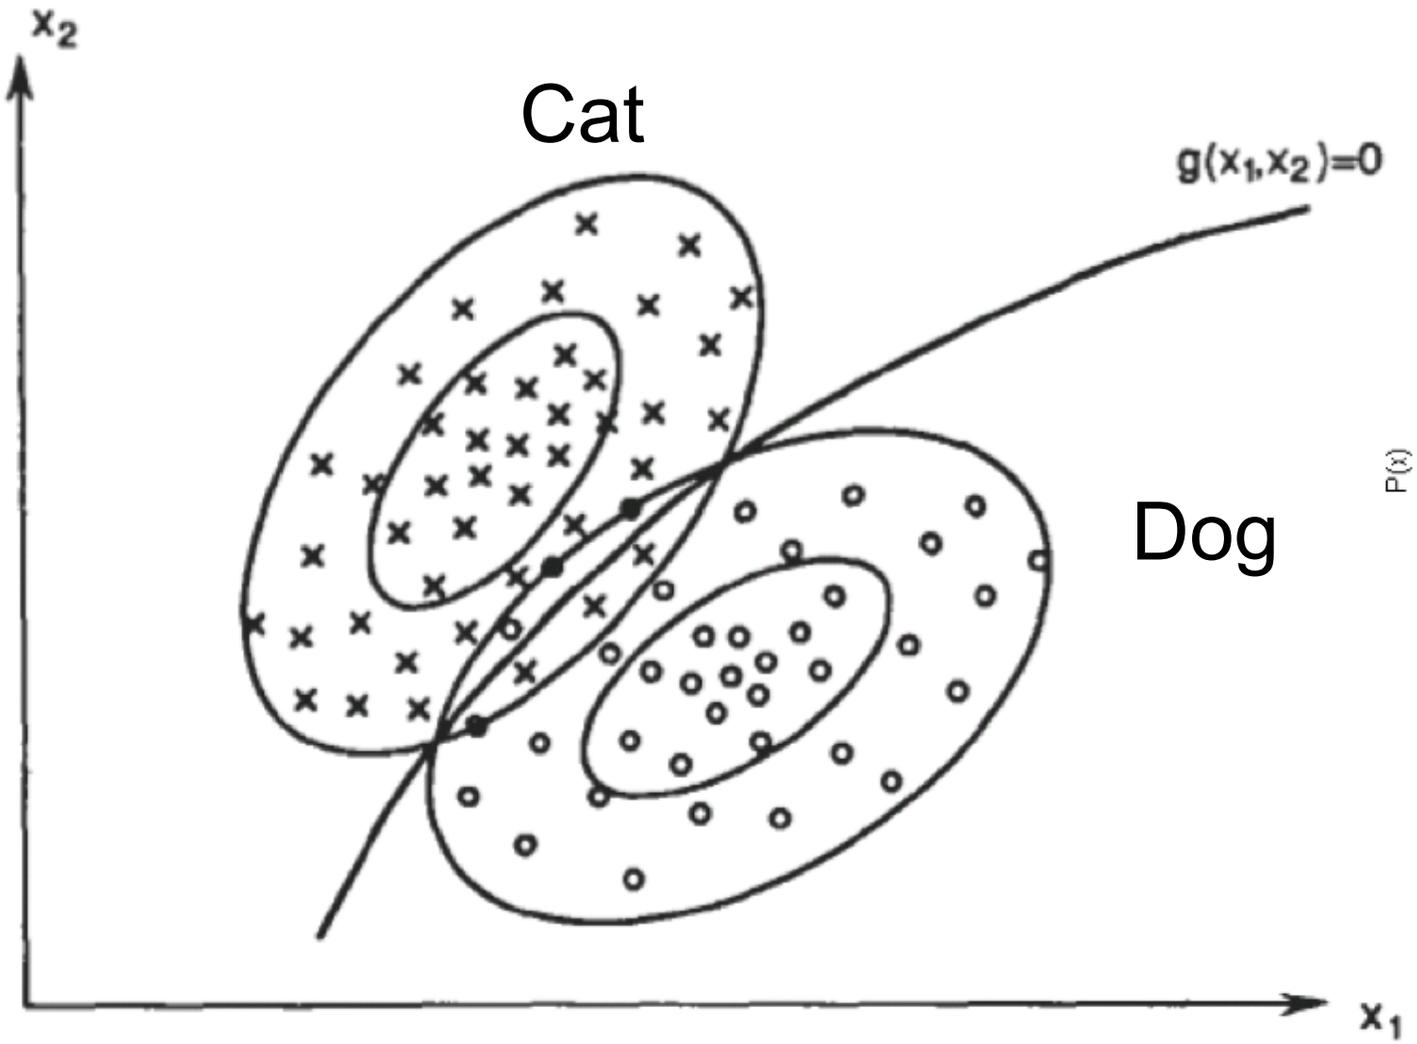

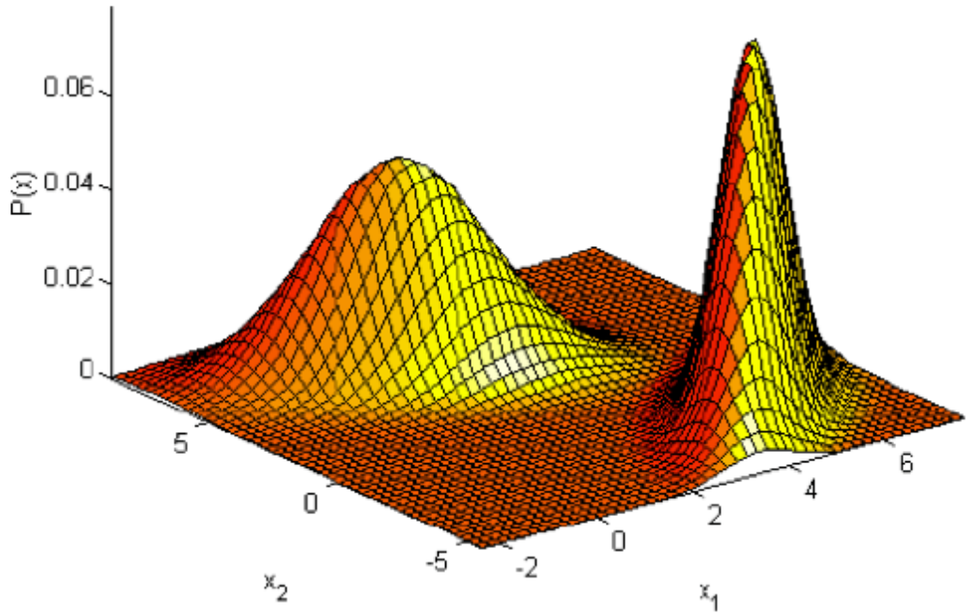

How to Determine the Decision Boundary?

- Given class conditional distribution: $P(x|y=1), P(x|y=-1)$, and class prior: $P(y=1), P(y=-1)$

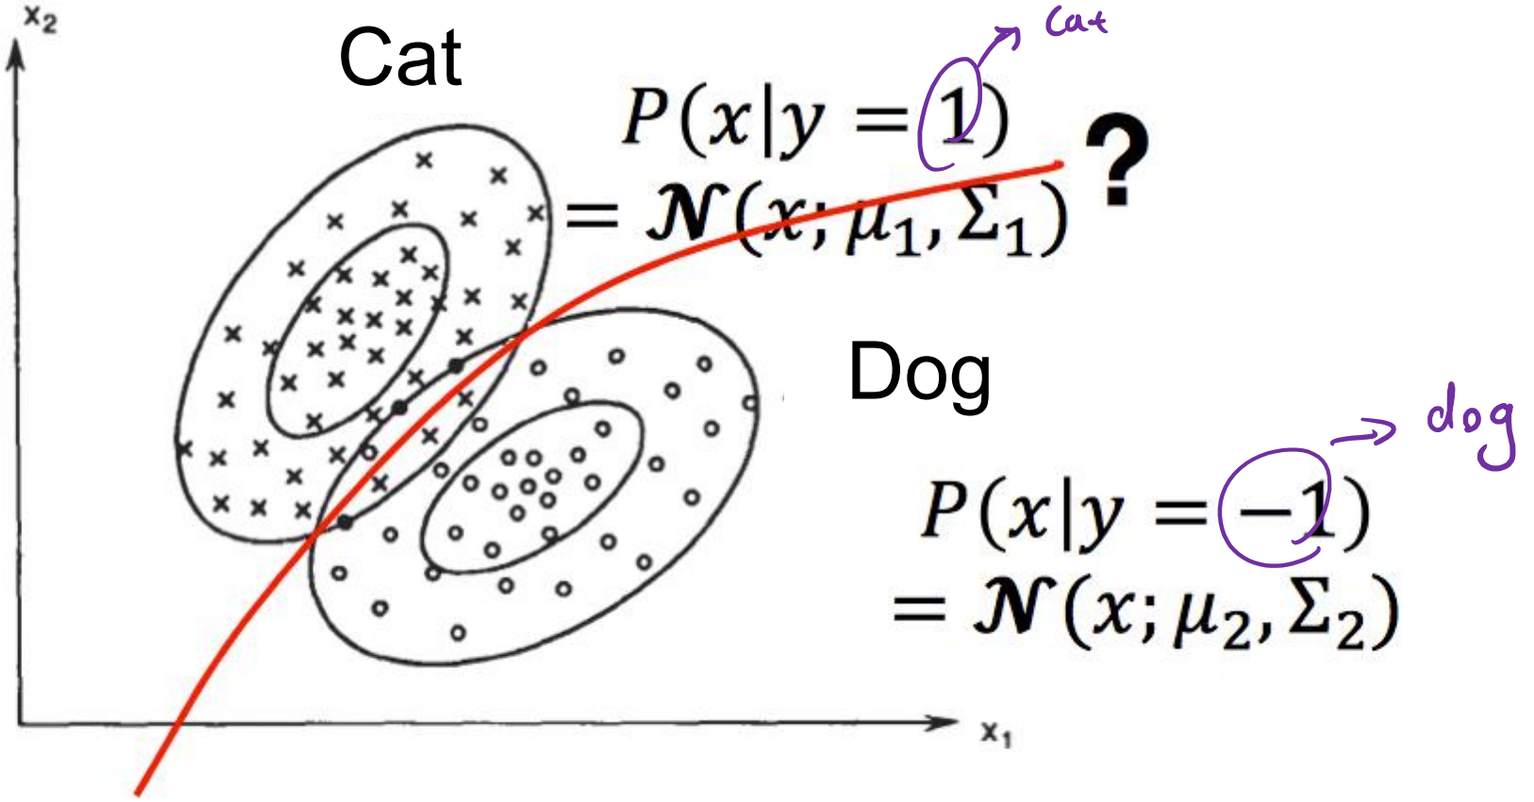

Bayes Decision Rule

likelihood Prior

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

posterior normalization constant

Prior: $P(y)$

Likelihood (class conditional distribution : $p(x|y) = \mathcal{N}(x|\mu_y, \Sigma_y)$

Posterior: $P(y|x) = \frac{P(y)\mathcal{N}(x|\mu_y, \Sigma_y)}{\sum_y P(y)\mathcal{N}(x|\mu_y, \Sigma_y)}$

Bayes Decision Rule

- Learning: prior: $p(y)$,class conditional distribution : $p(x|y)$

- The poster probability of a test point

$$q_i(x) \coloneqq P(y = i|x) = \frac{P(x|y)P(y)}{P(x)}$$

- Bayes decision rule:
  - If $q_i(x) > q_j(x)$, then $y = i$, otherwise $y = j$

- Alternatively:
  - If ratio $l(x) = \frac{P(x|y=i)}{P(x|y=j)} > \frac{P(y=j)}{P(y=i)}$, then $y = i$, otherwise $y = j$
  - Or look at the log-likelihood ratio $\mathrm{h(x)} = -\ln \frac{q_i(x)}{q_j(x)}$

$$X = \begin{matrix} h & w \\ \end{matrix} \begin{bmatrix} 1 & 2 \\ 3 & 4 \\ 5 & 6 \\ 7 & 8 \\ 9 & 10 \\ 11 & 12 \end{bmatrix} \qquad Y = \begin{bmatrix} \text{cat} \\ \text{cat} \\ \text{cat} \\ \text{dog} \\ \text{dog} \\ \text{dog} \end{bmatrix}$$

$$x^{[+]} \rightarrow \text{test} \qquad x^{[+]} = \begin{matrix} h & w \\ [13, & 14] \end{matrix} \qquad y = ?$$

$y = +1 \rightarrow \text{cat} \qquad y = -1 \rightarrow \text{dog}$

$$M = [M_h \quad M_w]$$

$$p(y=+1 \mid x^{[+]}) = \frac{p(y=+1, x^{[+]})}{p(x^{[+]})} = \frac{\overbrace{\underbrace{p(x^{[+]} \mid y=+1)}_{\text{likelihood}} \; p(y=+1)}^{a} \rightarrow \text{prior}}{p(x^{[+]}) = \sum_y p(x^{[+]}, y) =}$$

$$\underbrace{p(x^{[+]}, y=+1)}_{a} + \underbrace{p(x^{[+]}, y=-1)}_{b}$$

$$p(y=-1 \mid x^{[+]}) = \frac{b}{a+b}$$

$$P(y=+1 \mid x^{\{+\}}) = \frac{N(x^{\{+\}} \mid \mu, \Sigma)\, P(y=+1)}{P(x^{\{+\}})}$$

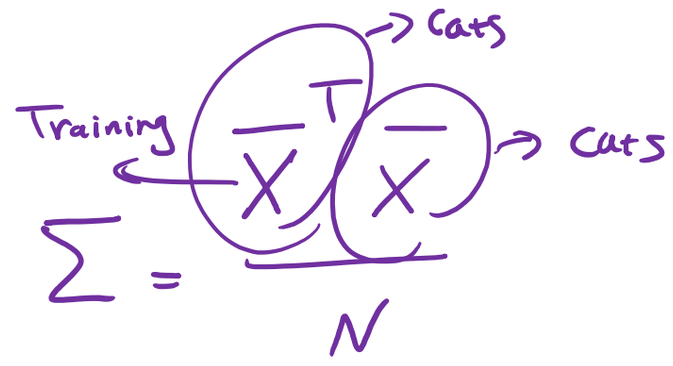

cats

$$N(x^{\{+\}} \mid \mu, \Sigma) = \frac{1}{\sqrt{2\pi\Sigma}^{d/2}} \exp\left[(x^{\{+\}} - \mu)\, \Sigma^{-1}\, (x^{\{+\}} - \mu)^T \right]$$

$$P(y=-1 \mid x^{\{+\}}) = ?$$

Naïve Bayes → "all features are conditionally independent"

$$Cov = \begin{bmatrix} \sigma_h^2 & 0 \\ 0 & \sigma_\omega^2 \end{bmatrix} \quad \Sigma^{-1}$$

$$X^{\{+\}} = [x_1, \dots, x_d]$$

$$P(y=+1 \mid x^{\{+\}}) = \frac{P(x_1^{\{+\}}, x_2^{\{+\}}, \dots, x_d^{\{+\}} \mid y=+1)\, P(y=+1)}{P(x^{\{+\}})}$$

$$= \frac{P(x_1^{\{+\}} \mid y=+1)\, P(x_2^{\{+\}} \mid y=+1) \cdots P(x_d^{\{+\}} \mid y=+1)\, P(y=+1)}{P(x^{\{+\}})}$$

What do People do in Practice?

- Generative models
  - Model prior and likelihood explicitly
  - “Generative” means able to generate synthetic data points
  - Examples: Naive Bayes, Hidden Markov Models
- Discriminative models
  - Directly estimate the posterior probabilities
  - No need to model underlying prior and likelihood distributions
  - Examples: Logistic Regression, SVM, Neural Networks

$P(y=+1 \mid x^{\{+3}) \Rightarrow$ directly

Generative Model: Naive Bayes

- Use Bayes decision rule for classification

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

Conditionally

- But assume $p(x|y = 1)$ is fully factorized : Dimensions are independent.

$$p(x|y = 1) = \prod_{i=1}^{d} p(x_i|y = 1)$$

- Or the variables corresponding to each dimension of the data are independent given the label

“Naïve” conditional independence assumption

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

Joint probability model:

$$P(x, y_{label=1}) = \mathrm{P}(\mathrm{x}_1, \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}}) = \mathrm{P}(\mathrm{x}_1|\mathrm{x}_2, , \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})P(\mathrm{x}_2, \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})$$

$$= \mathrm{P}(\mathrm{x}_1|\mathrm{x}_2, , \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})P(\mathrm{x}_2|\mathrm{x}_3 \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})P(\mathrm{x}_3, \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})$$

$$= \cdots$$

$$= \mathrm{P}(\mathrm{x}_1|\mathrm{x}_2, , \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}})P(\mathrm{x}_2|\mathrm{x}_3 \ldots, \mathrm{x}_\mathrm{d}, \mathrm{y}_{\mathrm{label=1}}) \ldots P(\mathrm{x}_{\mathrm{d}-1}|x_d, y_{label=1})P(x_d|y_{label=1})\mathrm{P}(y_{label=1})$$

Naïve Bayes assumption: let’s rewrite it as:

$$P(x, y_{label=1}) = P(x_1|y_{label=1})P(x_2|y_{label=1}) \ldots P(x_d|y_{label=1})P(y_{label=1}) =$$

$$P(y_{label=1}) \prod_{i=1}^{d} P(x_i|y_{label=1})$$

Gaussian naïve Bayes
A typical assumption

Example

Discriminative Models

- Directly estimate decision boundary $\mathrm{h}(\mathrm{x}) = -\ln \frac{q_i(x)}{q_j(x)}$ or posterior distribution $p(y|x)$
  - Logistic regression, Neural networks
  - Do not estimate $p(x|y)$ and $p(y)$

- Why discriminative classifier?
  - Avoid difficult density estimation problem → Generative model
  - Empirically achieve better classification results

Outline

- Generative and Discriminative Classification
- The Logistic Regression Model
- Understanding the Objective Function
- Gradient Descent for Parameter Learning
- Multiclass Logistic Regression

Gaussian Naïve Bayes GNB

$$P(y=1|x) = \frac{P(x|y=1)P(y=1)}{P(x)} = \frac{P(y=1)\prod_{i=1}^{d} P(x_i|y=1)}{P(x)}$$

$$\prod_{i=1}^{d} p(x_i|y=1,\mu_{1i},\sigma_{1i})$$

$$= \prod_{i=1}^{d} \frac{1}{\sqrt{2\pi}\sigma_{1i}} \exp\left(-\frac{1}{2\sigma_{1i}^2}(x_{1i}-\mu_{1i})^2\right)$$

Prior: $p(y=1) = \pi_1$

Posterior: $p(y = 1 \mid x, \mu, \sigma, \pi)$

$$= \frac{\pi_1 \prod_{i=1}^{d} \frac{1}{\sqrt{2\pi}\sigma_{1i}} \exp\left(-\frac{1}{2\sigma_{1i}^2}(x_i - \mu_{1i})^2\right)}{\sum_{k=1}^{2} \pi_k \prod_{i=1}^{d} \frac{1}{\sqrt{2\pi}\sigma_{ki}} \exp\left(-\frac{1}{2\sigma_{ki}^2}(x_i - \mu_{ki})^2\right)}$$

labels

get $\exp(\ln(u))$ of numerator and denominator

$$= \frac{\exp\left(-\sum_{i=1}^{d}\left(\frac{1}{2\sigma_{1i}^2}(x_i - \mu_{1i})^2 + \log\sigma_{1i} + C\right) + \log\pi_1\right)}{\sum_{k=1}^{2} \exp\left(-\sum_{i=1}^{d}\left(\frac{1}{2\sigma_{ki}^2}(x_i - \mu_{ki})^2 + \log\sigma_{ki} + C\right) + \log\pi_k\right)}$$

$$\frac{a}{a+b} \quad \frac{\frac{a}{a}}{\frac{a}{a}+\frac{b}{a}} = \frac{1}{1+\frac{b}{a}}$$

$$= \frac{\exp\left(-\sum_{i=1}^{d}\left(\frac{1}{2\sigma_i^2}(x_i-\mu_{1i})^2+\log\sigma_i+C\right)+\log\pi_1\right)}{\sum_{k=1}^{2}\exp\left(-\sum_{i=1}^{d}\left(\frac{1}{2\sigma_i^2}(x_i-\mu_{ki})^2+\log\sigma_i+C\right)+\log\pi_k\right)}$$

$$P(y=+1 \mid x) \overset{\{+\}}{=} \frac{1}{1+\exp\left(-\sum_{i=1}^{d}\left(x_i\underbrace{\frac{1}{\sigma_i}(\mu_{1i}-\mu_{2i})}_{\theta_i}+\underbrace{\frac{1}{\sigma_i^2}(\mu_{1i}^2-\mu_{2i}^2)}_{\theta_0}\right)+\log\frac{\pi_2}{\pi_1}\right)}$$

$$\frac{\sum \theta_i x_i}{1+\exp\left(-\left[\sum_{i=1}^{d} x_i\theta_i+\theta_0\right]\right)} = \frac{1}{1+\exp\underbrace{\left[-(\theta_0+\theta_1x_1+\cdots+\theta_dx_d)\right]}_{LCF}}$$

$$P(y=1|x) = \frac{1}{1+\exp\left(-\sum_{i=1}^{d}\left(x_i\frac{1}{\sigma_i}(\mu_{1i}-\mu_{2i}) + \frac{1}{\sigma_i^2}(\mu_{1i}^2-\mu_{2i}^2)\right) + \log\frac{\pi_2}{\pi_1}\right)}$$

Number of parameters:

$2d+1 \rightarrow d$ *mean*, d *variance*, *and* 1 *for prior*

Summation

$$P(y=1|x) = \frac{1}{1+\exp[-(\sum_{i=1}^{d}(\theta_i x_i)+\theta_0)]} = \frac{1}{1+\exp(-s)}$$

Sigmoid

$S = LCF = \theta_0 + \theta_1 x_1 + \cdots + \theta_d x_d = x\theta$

Number of parameters = $d+1 \rightarrow \theta_0, \theta_1, \theta_2, \ldots, \theta_d$

Why not directly learning $P(y=1|x)$ or θ parameters?

Gaussian Naïve Bayes is a subset of logistic regression

Why $\frac{1}{1+\exp(-x\theta)}$ is a probability?

$X\theta = LCF = S \in \mathbb{R}$

Summation

$+\infty$

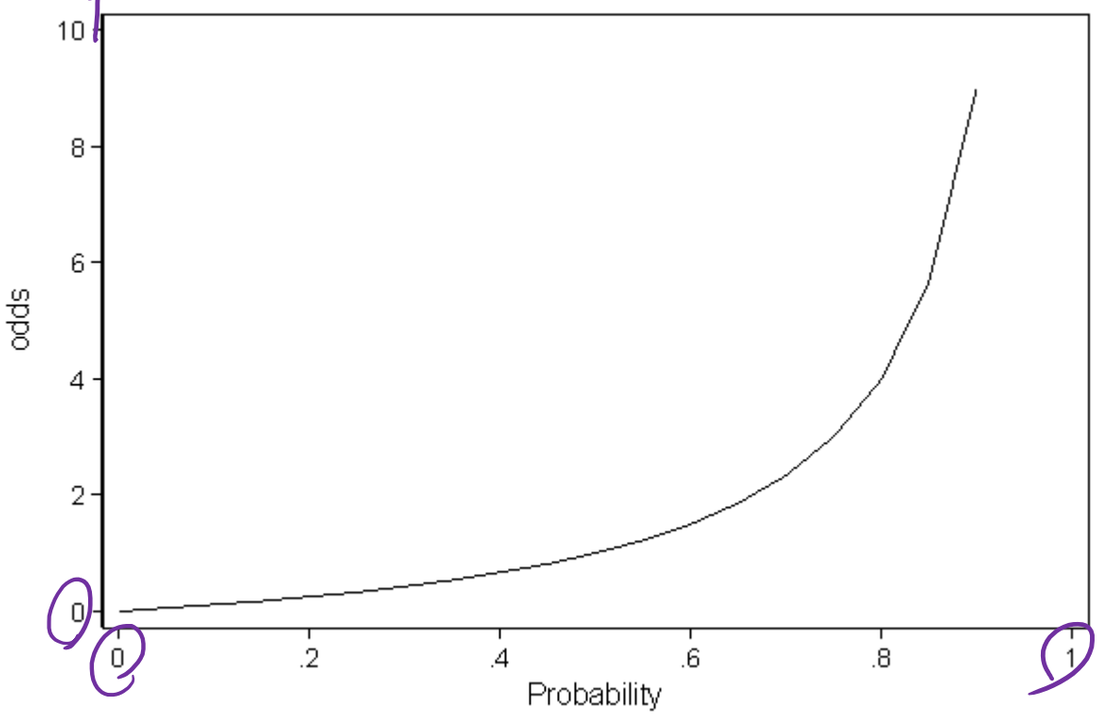

$\frac{P(y=1|x)}{1-P(y=1|x)}$ is called Odds

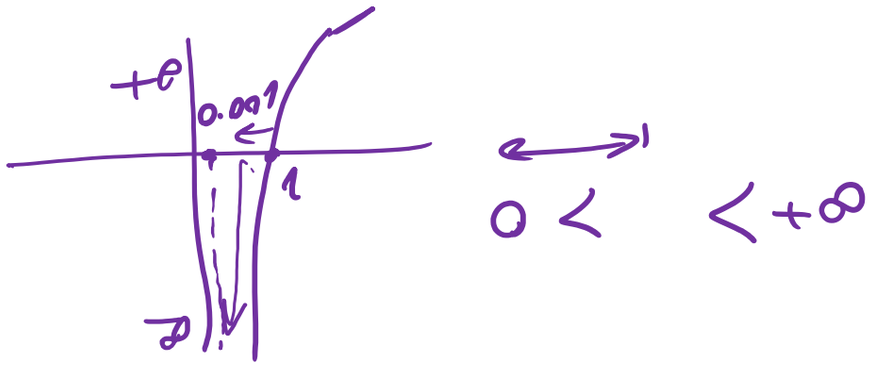

$\in \mathbb{R}$

$log(odds)\ vs\ odds$

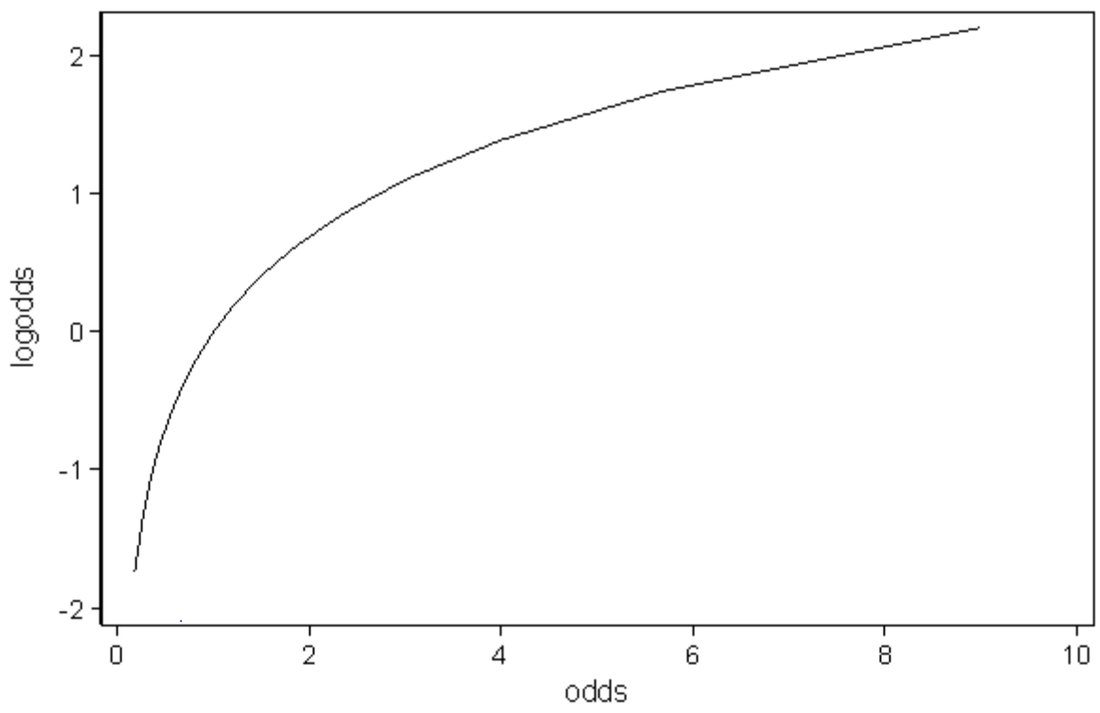

What could be $x\theta$ domain?

What is logit function?

$$logit(p) = log(odds) = \log\left(\frac{p}{1-p}\right)$$

$$\log\left(\frac{p}{1-p}\right) = \theta_0 + \theta_1 x_1 + \cdots + \theta_d x_d = \sum_{i=0}^{d} x_i \theta_i = x\theta$$

$$exp\left(log\left(\frac{p}{1-p}\right)\right) = \exp(x\theta)$$

$$P(y=+1 \mid x^{(i)}) = p = \frac{e^{x\theta}}{1+e^{x\theta}} = \frac{1}{1+e^{-x\theta}}$$

Sigmoid

Logistic function for posterior probability

$-\infty < S = LCF = X\theta = \theta_0 + \theta_1 x_1 + \cdots + \theta_d x_d < +\infty$

$g(s) = \frac{1}{1 + e^{-x^{[test3]}\theta}} > 0.5$ Cat

$\leq$ dog

Let's use the following function:

$$s = x\theta$$

$$g(s) = P(y = 1|x) = \frac{e^s / e^s}{1 + e^s} = \frac{1}{1 + e^{-s}}$$

This formula is called sigmoid function

Many equations can give us this shape

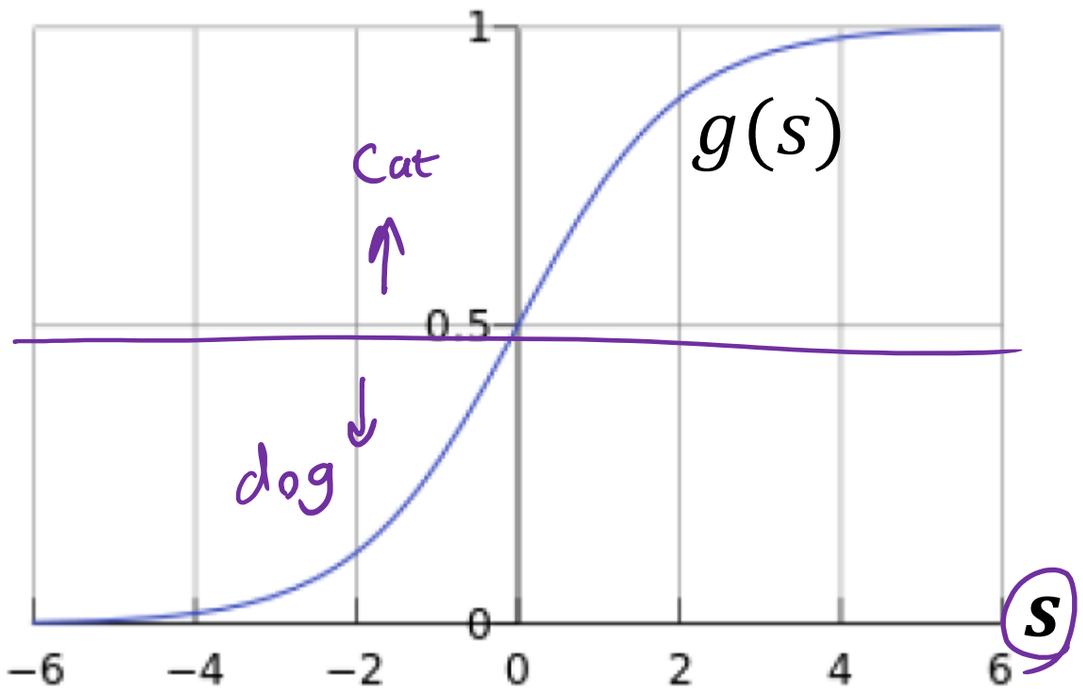

It is easier to use this function for optimization

Is 0.5 threshold cut-off a good choice?

Learn about ROC and AUC (False positive rate and True positive rate) (Interactive)

Sigmoid Function

$$g(s) = \frac{e^s}{1+e^s} = \frac{1}{1+e^{-s}}$$

$$s = \sum_{i=0}^{d} x_i\theta_i = \theta_0 + \theta_1 x_1 + \cdots + \theta_d x_d$$

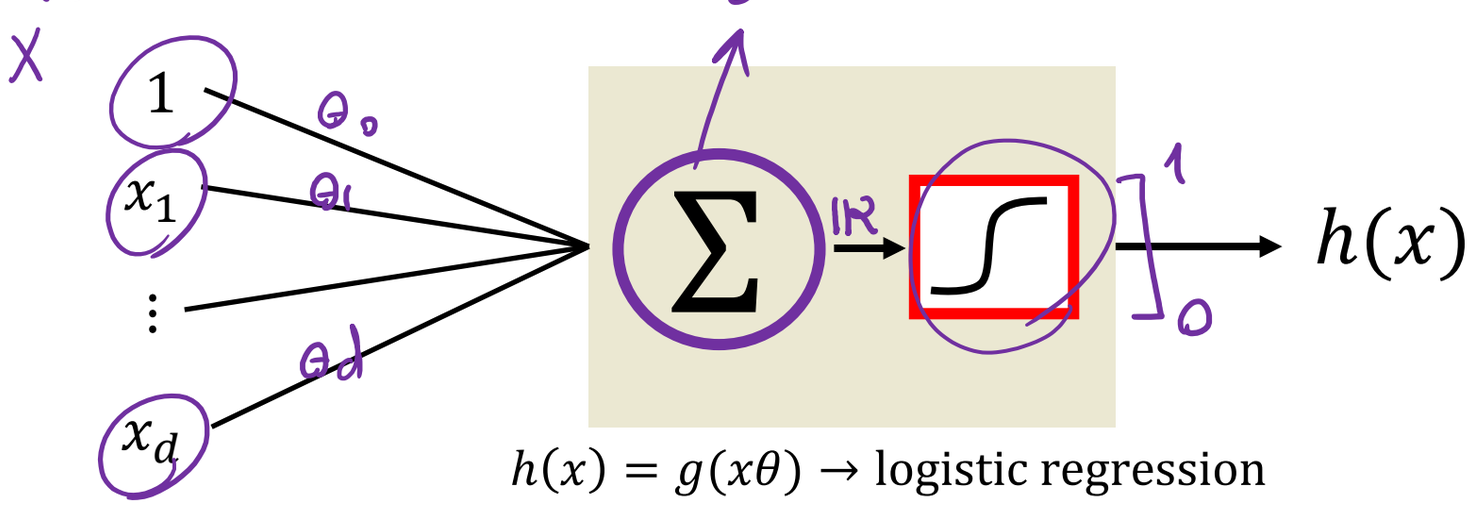

$h(x) = g(x\theta) \rightarrow$ logistic regression

Soft classification
Posterior probability

$$s = \sum_{i=0}^{d} x_i \theta_i = \theta_0 + \theta_1 x_1 + \cdots + \theta_d x_d$$

Three linear models

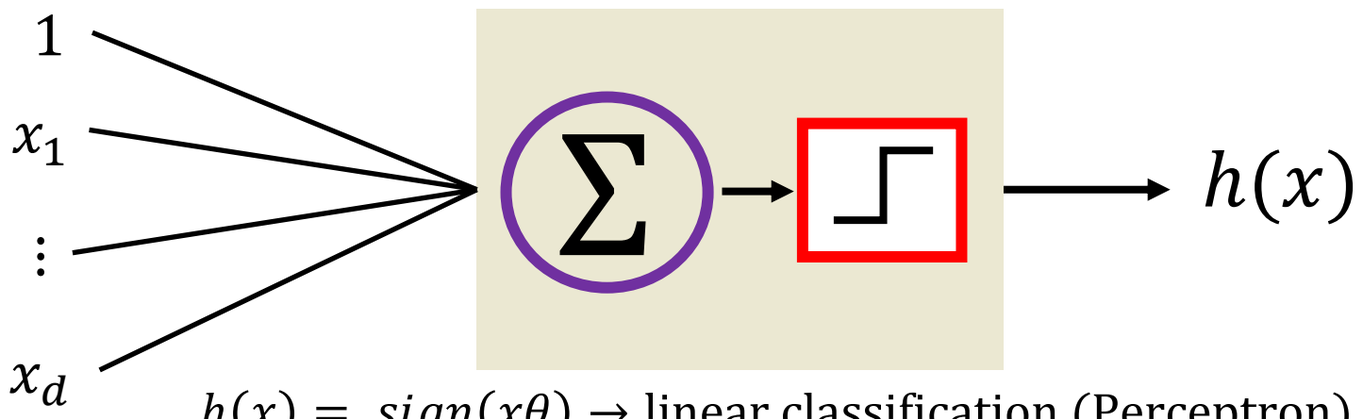

Hard classification

$h(x) = sign(x\theta) \rightarrow$ linear classification (Perceptron)

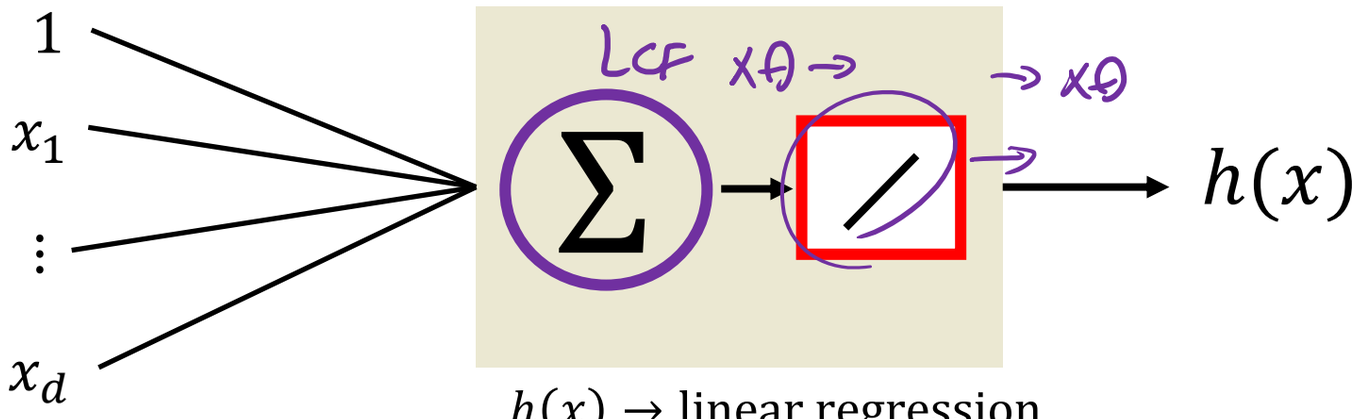

$h(x) \rightarrow$ linear regression

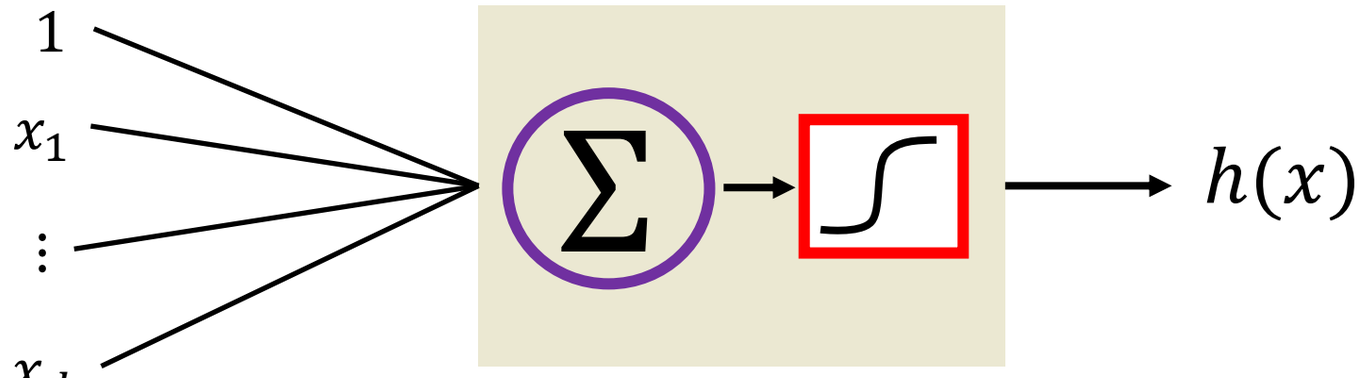

Soft classification
Posterior probability

$h(x) = g(x\theta) \rightarrow$ logistic regression

$g(s)$ is interpreted as probability

Example: Prediction of heart attacks

Input $\boldsymbol{x}$: cholesterol level, age, weight, finger size, etc.

$g(s)$: probability of heart attack within a certain time

We can't have a hard prediction here

$s = x\theta$ Let's call this risk score

$$\prod_{i=1}^{N} p(y^{(i)}|x^{(i)})$$

$$p(y|x) = \frac{1}{1+\exp(-x\theta)}$$

$d+1$

$$h_\theta(x) = p(y|x) = \begin{cases} g(s), & y = 1 \\ 1 - g(s), & y = 0 \end{cases}$$

sigmoid

→ cat

dog

Using posterior probability directly

Logistic regression model

$$p(y|x) = \begin{cases} \dfrac{1}{1+\exp(-x\theta)} & y = 1 \\ 1 - \dfrac{1}{1+\exp(-x\theta)} = \dfrac{\exp(-x\theta)}{1+\exp(-x\theta)} & y = 0 \end{cases}$$

We need to find θ parameters, let's set up log-likelihood for **n** datapoints

Maximize log likelihood

$$l(\theta) := log \prod_{i=1}^{n} p(y^{\{i\}}|x^{\{i\}}, \theta) = \log \prod_{i=1}^{N} P(y^{[i]}|x^{[i]},\theta)^{y^{[i]}} (1 - P(y^{[i]}|x^{[i]},\theta))^{1-y^{[i]}}$$

$$= \sum_{i} \theta^T (x^{\{i\}})^T \left(y^{\{i\}} - 1\right) - \log\left(1 + \exp\left(-x^{\{i\}}\theta\right)\right)$$

This form is concave, negative of this form is convex

The gradient of $l(\theta)$

$$\min \quad l(\theta) = -log \prod_{i=1}^{n} p(y^{\{i\}}|x^{\{i\}}, \theta) = CE = \sum y_a^{\{i\}} \ln y_p^{\{i\}}$$

$$= \sum_i \theta^T (x^{\{i\}})^T \left(y^{\{i\}} - 1\right) - \log(1 + \exp(-x^{\{i\}}\theta))$$

- Gradient

$$\min \quad \frac{\partial l(\theta)}{\partial \theta} = -\sum_i (x^{\{i\}})^T \left(y^{\{i\}} - 1\right) + (x^{\{i\}})^T \frac{\exp(-x^{\{i\}}\theta)}{1 + \exp(-x^{\{i\}}\theta)}$$

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} + \alpha \frac{\partial L}{\partial \theta}$$

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} - \alpha \frac{\partial L}{\partial \theta}$$

- Setting it to 0 does not lead to closed form solution

The Objective Function

- Find θ, such that the conditional likelihood of the labels is maximized

$$\max_{\theta} l(\theta) := log \prod_{i=1}^{n} p(y^{\{i\}}|x^{\{i\}}, \theta)$$

- Good news: $l(\theta)$ is concave function of θ, and there is a single global optimum.

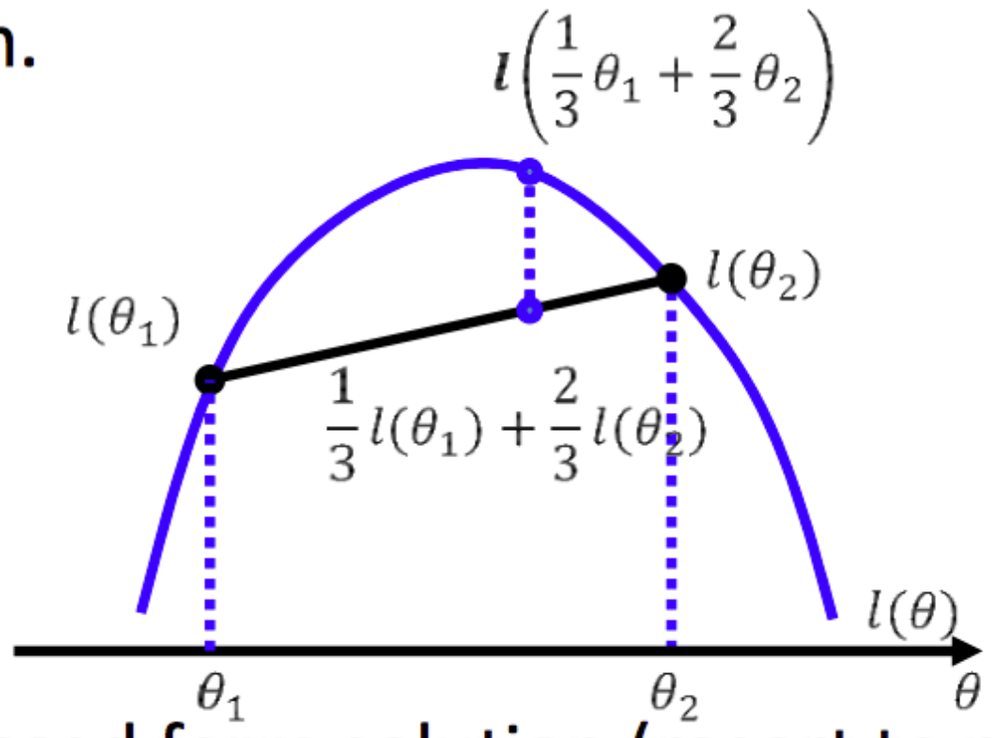

- Bad new: no closed form solution (resort to numerical method)

Gradient Descent

- One way to solve an *unconstrained* optimization problem is gradient descent

- Given an initial guess, we *iteratively* refine the guess by taking the direction of the negative gradient

- Think about going down a hill by taking the steepest direction at each step

- Update rule

$$x_{k+1} = x_k - \gamma_k \nabla f(x_k)$$

γ_k is called the step size or learning rate

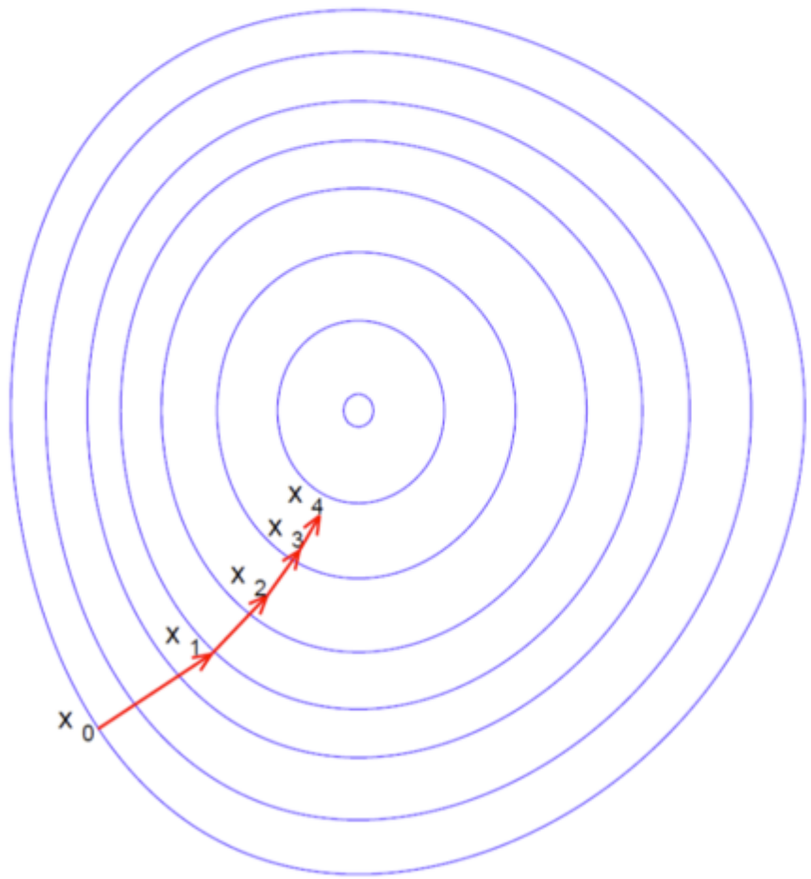

Logistic Regression

$$g(s) = \frac{e^s}{1+e^s} = \frac{1}{1+e^{-s}}$$

$$s = x\theta$$

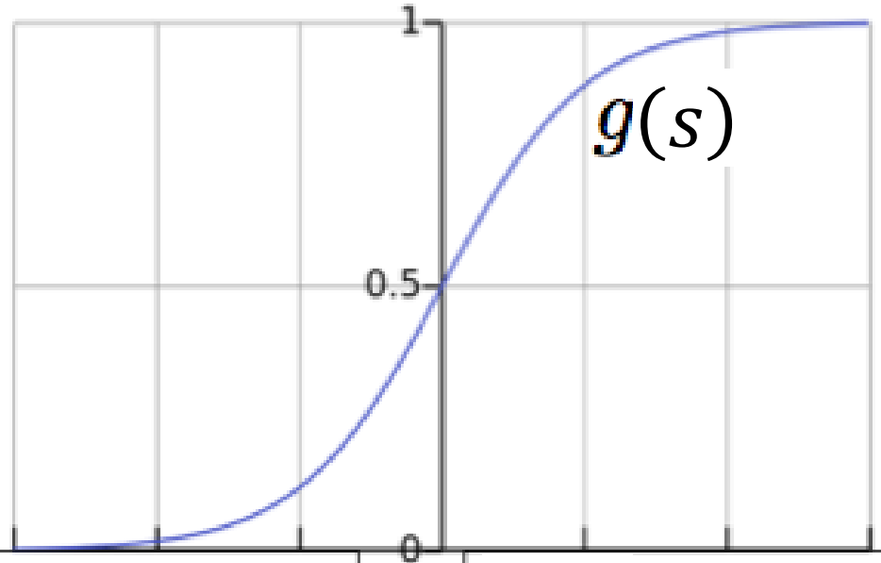

$x\theta$ should be large negative values for negative instances	$x\theta$ should be large positive values for positive instances

- Assume a threshold and...
  - Predict $y = 1$ if $g(s) \geq 0.5$
  - Predict $y = 0$ if $g(s) < 0.5$

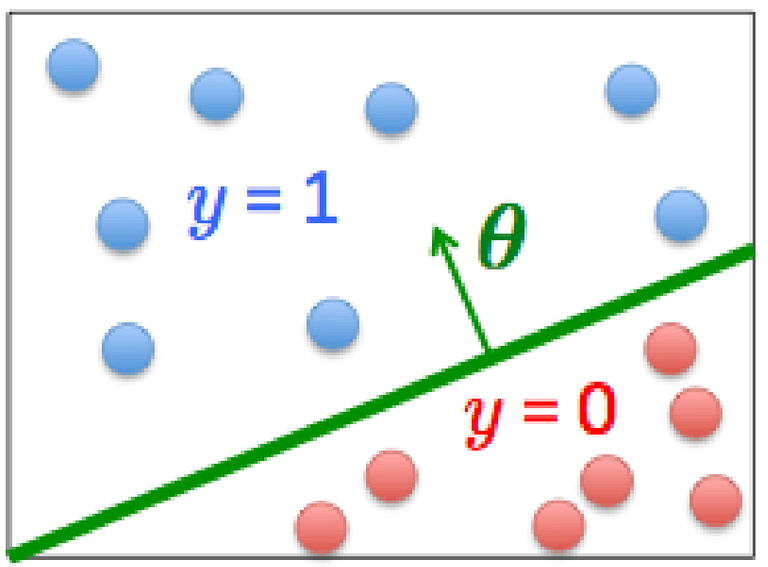

Outline

- Generative and Discriminative Classification
- The Logistic Regression Model
- Understanding the Objective Function
- Gradient Descent for Parameter Learning
- Multiclass Logistic Regression 

Multiclass Logistic Regression

Binary classification:

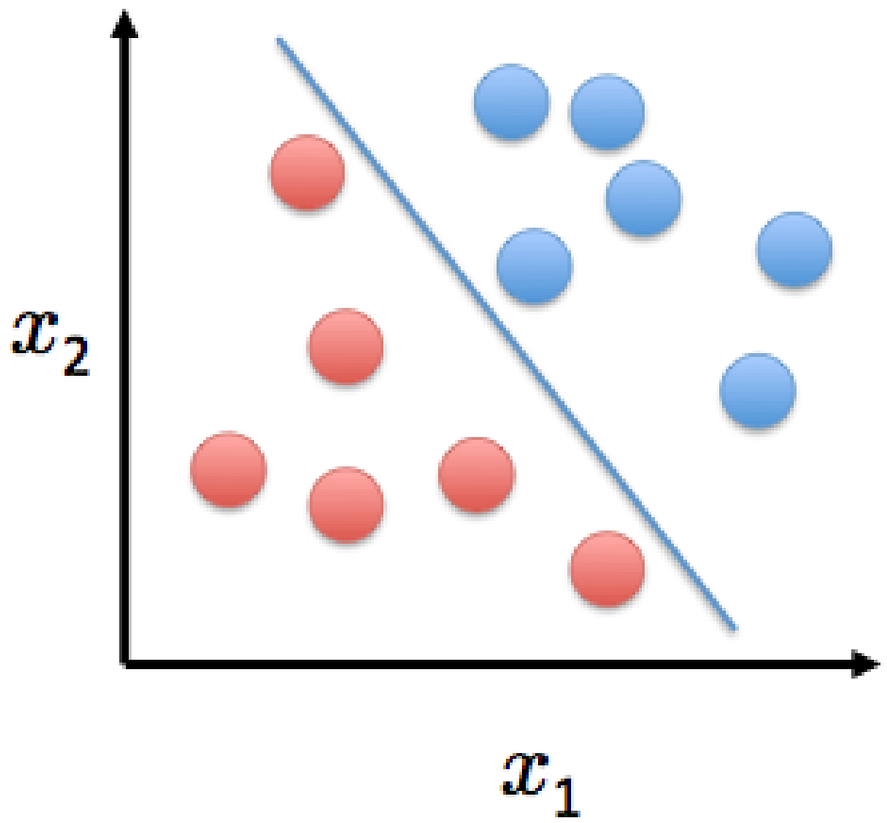

Multi-class classification:

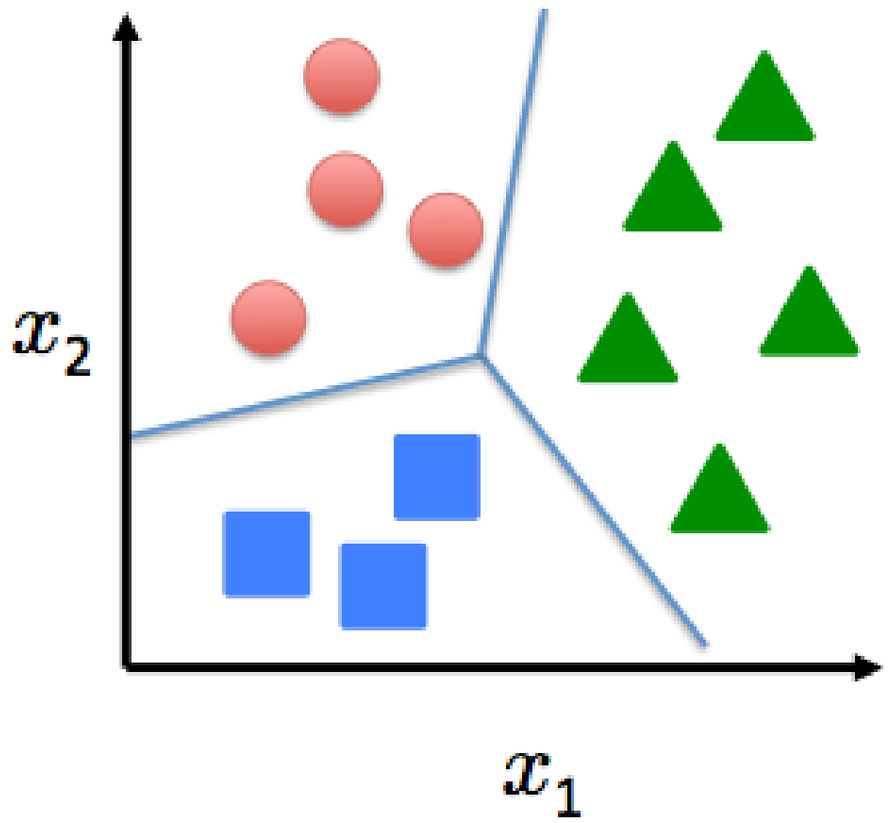

Disease diagnosis: healthy / cold / flu / pneumonia

Object classification: desk / chair / monitor / bookcase

Gradient Ascent(concave)/~~Descent(convex)~~ algorithm

Negative log likelihood = Cross Entropy = $-\sum y_a^{\{i\}} \log y_p^{\{i\}}$

$\frac{\partial CE}{\partial \theta} = \bar{X}^T (y_p - y_a)$

$\frac{\partial L}{\partial \theta}$

$y_p = \frac{1}{1+\exp(-s)} = \frac{1}{1+\exp(-x\theta)}$

y_p

- Initialize parameter θ^0
- Do

$$\theta^{t+1} \leftarrow \theta^t + \eta \sum_i (x^{\{i\}})^T \left(y^{\{i\}} - 1\right) + (x^{\{i\}})^T \frac{\exp\left(-x^{\{i\}}\theta\right)}{1+\exp(-x^{\{i\}}\theta)}$$

- While the $||\theta^{t+1} - \theta^t|| > \epsilon$

$\theta^t - \sum x^{\{i\}T} \left(y_p^{\{i\}} - y_a^{\{i\}} \right)$

$$P(y|x) = \text{Sigmoid} = \frac{1}{1+\exp(-s)} \qquad s = x\theta$$

$$\log \prod_{i=1}^{N} P(y^{[i]} | x^{[i]})$$

$$P(y=1|x) = \frac{\exp(s_1)}{\exp(s_1)+\exp(s_2)} \quad \text{cat}$$

$$P(y=2|x) = \frac{\exp(s_2)}{\exp(s_1)+\exp(s_2)} \quad \text{dog}$$

$s_1 = x\theta_{cat}$ for cats

$s_2 = x\theta_{dog}$ for dogs

$$\theta = \begin{bmatrix} \theta_0 & \theta_0 \\ \vdots & \vdots \\ \theta_d & \theta_d \end{bmatrix} \begin{matrix} \text{cat} & \text{dog} \end{matrix}$$

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} - \alpha \underbrace{X^T}_{d\times n} \underbrace{(Y_p - Y_a)}_{n\times 2}$$

$d \times 2$

$$Y_p = \begin{bmatrix} & \\ & \\ & \end{bmatrix}_{n\times 2} \begin{matrix} \text{cat} & \text{dog} \end{matrix}$$

$\frac{1}{1+\exp(-x\theta)}$

$$P(y=1|x) = \frac{1}{1+\exp(\underbrace{s_2 - s_1}_{s})} = \frac{1}{1+\exp(-s)}$$

$$P(y=2|x) = \frac{\exp(s_2-s_1)}{1+\exp(\underbrace{s_2-s_1}_{s})}$$

$$s_2 - s_1 = x(\theta_{dog} - \theta_{cat}) = x\theta$$

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} - \alpha\, x^T (Y_p - Y_a)$$

$S_1 = X\theta_{cat}$ $S_2 = X\theta_{dog}$ $S_3 = X\theta_{fish}$

$$\theta = \begin{bmatrix} \theta_{cat} & \theta_{dog} & \theta_{fish} \end{bmatrix}$$

$Y_a^{\{i\}} = [1 \;\; 0 \;\; 0]$ $\hat{Y}_p = [0.7 \;\; 0.2 \;\; 0.1]$ (0.2: dog, 0.1: fish)

cat — Reference

$$P(y=1|x) = \frac{\exp(s_1)}{\exp(s_1)+\exp(s_2)+\exp(s_3)} = \frac{1}{1+\exp(\underbrace{s_2-s_1}_{\theta_{(1)}})+\exp(\underbrace{s_3-s_1}_{\theta_{(2)}})}$$

dog

$$P(y=2|x) = \frac{\exp(s_2)}{C} = \frac{\exp(s_2-s_1)}{1+\exp(s_2-s_1)+\exp(s_3-s_1)} \quad (1)$$

fish

$$P(y=3|x) = \frac{\exp(s_3)}{C} = \frac{\exp(s_3-s_1)}{1+\exp(s_2-s_1)+\exp(s_3-s_1)} \quad (2)$$

$$\theta^{\{t+1\}} \leftarrow \theta^{\{t\}} - \alpha \bar{X}^T (\hat{Y}_p - Y_a)$$

θ: $d \times 2$ — $\bar{X}$: $d \times N$ — $\hat{Y}_p - Y_a$: $N \times 2$ (columns ① ②: dog, fish)

One-vs-all (one-vs-rest)

Multi-class classification:

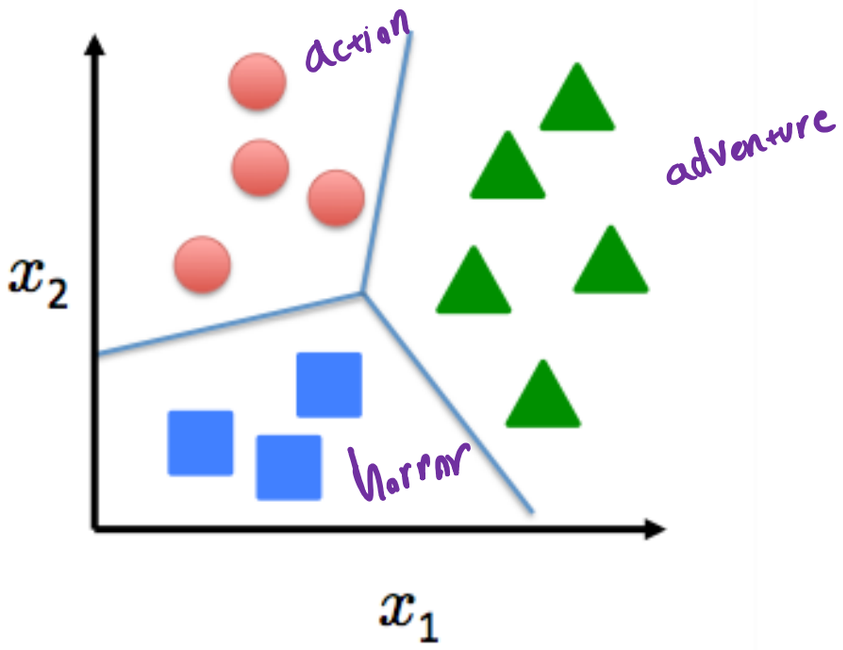

$$h_\theta^{(m)}(x) = p(y = 1|x, \theta)\ (m = 1,2,3)$$

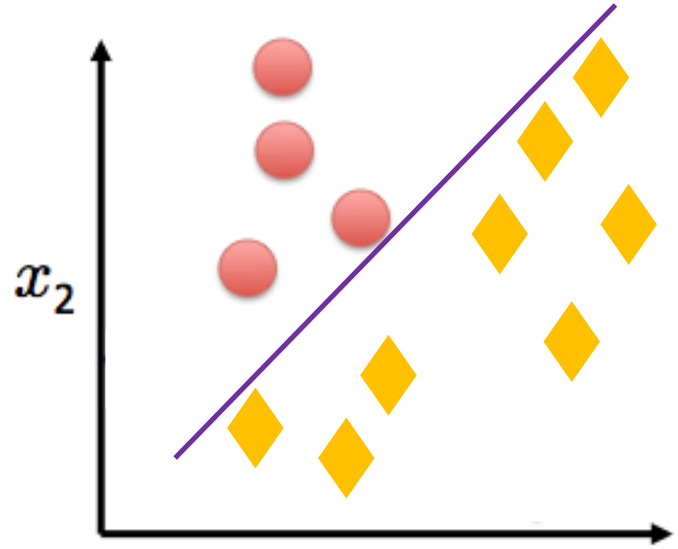

$h_\theta^1(x)$

0.6

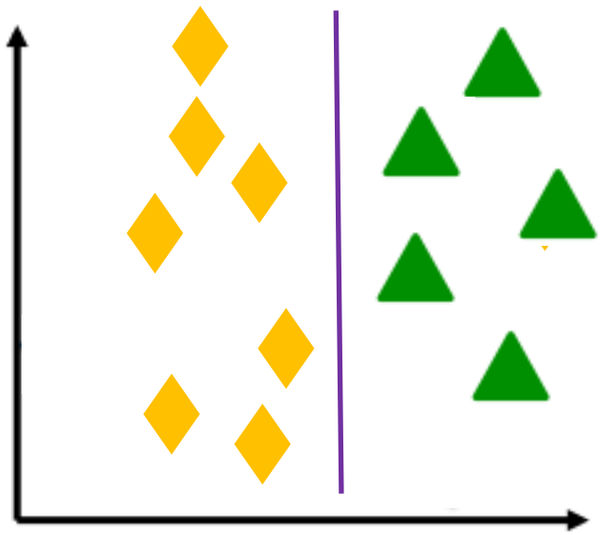

$h_\theta^2(x)$

0.75

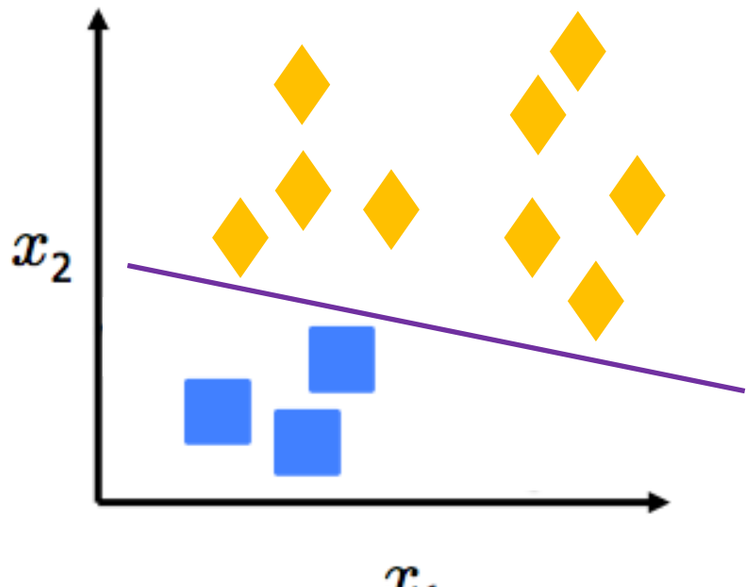

$h_\theta^3(x)$

0.4

One-vs-all (one-vs-rest)

Train a logistic regression $h_\theta^{(m)}(x)$ for each class m

To predict the label of a new input x, pick class m that maximizes:

$$\max_i h_\theta^{(m)}(x)$$

Using Softmax

$$L(\theta) = -\sum_{i=1}^{N} y_a^{\{i\}} * \log\left(y_p^{\{i\}}\right)$$

$y_a = [cat, dog, fish] = [1,0,0]$
$\Rightarrow there\ are\ M\ classes\ (M = 3\ in\ this\ example)$

$$y_{p\ for\ class\ \boldsymbol{m}} = softmax(x\theta) = \frac{\exp(x\theta)_{\boldsymbol{m}}}{\sum_{j=0}^{M} \exp(x\theta)_j}$$

$$y_p = [0.6,0.3,0.1]$$

$$SGD \Rightarrow \theta^{t+1} \leftarrow \theta^t - \alpha\, \nabla L(\theta)$$

$$\theta^{t+1} \leftarrow \theta^t - \alpha\, x^T (y_p - y_a)$$

Take-Home Messages

- Generative and Discriminative Classification
- The Logistic Regression Model
- Understanding the Objective Function
- Gradient Descent for Parameter Learning
- Multiclass Logistic Regression