

## Linear Regression

Training set  $(x_i, y_i)$ ,  $x_i \in \mathbb{R}^P$ ,  $y_i \in \mathbb{R}$ ,  $i = 1, \dots, n$

Hypothesis:

$$y_i = \vartheta_0 + x_{i1}\vartheta_1 + \dots + x_{ip}\vartheta_p + \varepsilon_i, i = 1, \dots, n$$

$$y = X\vartheta + \varepsilon, x_{ij}, y_i \in \mathbb{R}$$

$$\rightarrow \hat{\vartheta} = (X^T X)^{-1} X^T y$$

New observation  $x \in \mathbb{R}^P$

$$\hat{y} = (1, x^T) \hat{\vartheta} \text{ best predictor}$$

## 7.1.2 Logistic Regression

Training set  $(x_i, y_i)$ ,  $x_i \in \mathbb{R}^P$ ,  $y_i \in \{0, 1\}$ ,  $i = 1, \dots, n$

Hypothesis

$$h_{\vartheta}(x) = \frac{1}{1+e^{-\vartheta^T x}} \in [0, 1]$$

$$\text{Set } x_{i0} = 1, \vartheta^T x_i = \vartheta_0 + \sum_{j=1}^P \vartheta_j x_{ij}$$

Find the best  $\vartheta$ !

Probabilistic interpretation:

Assume  $P(y=1|x, \vartheta) = h_{\vartheta}(x)$

$$P(y=0|x, \vartheta) = 1 - h_{\vartheta}(x)$$

Write:  $p(y|x, \vartheta) = (h_{\vartheta}(x))^y (1-h_{\vartheta}(x))^{1-y}, y \in \{0, 1\}$

Assume  $n$  independent training samples,

$$X = \begin{pmatrix} 1 & x_1^T \\ \vdots & \vdots \\ 1 & x_n^T \end{pmatrix} = (x_{ij})_{\substack{1 \leq i \leq n \\ 0 \leq j \leq p}}$$

Likelihood:

$$L(\vartheta) = p(y|X, \vartheta) = \prod_{i=1}^n p(y_i|x_i, \vartheta)$$

$$= \prod_{i=1}^n (h_{\vartheta}(x_i))^{y_i} (1-h_{\vartheta}(x_i))^{1-y_i}$$

Log-Likelihood:

$$\ell(\vartheta) = \log L(\vartheta) = \sum_{i=1}^n (y_i \log h_{\vartheta}(x_i) + (1-y_i) \log (1-h_{\vartheta}(x_i)))$$

Objective:  $\max_{\vartheta \in \mathbb{R}^{p+1}} \ell(\vartheta)$

An appr. algorithm: gradient descent

$$\vartheta^{(k+1)} = \vartheta^{(k)} + \alpha \nabla_{\vartheta} \ell(\vartheta^k), \alpha \text{ learning parameter}$$

Compute  $\nabla$  for each element  $i$  in  $(x)$ , set  $(x_i, y_i) = (x, y)$

$$\begin{aligned}
 & \frac{\partial}{\partial \vartheta_j} \left[ y \log h_{\vartheta}(x) + (1-y) \log (1-h_{\vartheta}(x)) \right] \\
 &= \left( y \frac{1}{g(\vartheta^T x)} - (1-y) \frac{1}{1-g(\vartheta^T x)} \right) \frac{\partial}{\partial \vartheta_j} g(\vartheta^T x) \\
 &= \left( \dots \right) g(\vartheta^T x) (1-g(\vartheta^T x)) \frac{\partial}{\partial \vartheta_j} \vartheta^T x \\
 &= \left( y (1-g(\vartheta^T x)) - (1-y) g(\vartheta^T x) \right) x_j \\
 &= (y - h_{\vartheta}(x)) x_j
 \end{aligned}$$

Hence:

$$\frac{\partial}{\partial \vartheta_j} \ell(\vartheta) = \sum_{i=1}^n (y_i - h_{\vartheta}(x_i)) x_{ij}, \quad j=0, 1, \dots, p$$

Update rule:

$$\vartheta_j^{(k+1)} = \vartheta_j^{(k)} + \alpha \sum_{i=1}^n (y_i - h_{\vartheta^{(k)}}(x_i)) x_{ij} \quad j=0, \dots, p.$$

Alternate: Newton's method

$$\vartheta^{(k+1)} := \vartheta^{(k)} - H^{-1} \nabla_{\vartheta} \ell(\vartheta^{(k)})$$

### 7.1.3 The perceptron learning algorithm

In logistic regression values in  $(0, 1)$ . (Soft decision)

Force 0 or 1! by taking

$$g(z) = \begin{cases} 1 & , z \geq 0 \\ 0 & , z < 0 \end{cases}$$

$$h_{\omega}(x) = g(\omega^T x)$$

Use the update rule

$$\omega_j^{(k+1)} = \omega_j^{(k)} + \alpha \sum_{i=1}^n (y_i - h_{\omega^{(k)}}(x_i)) x_{i,j}$$

the so called perceptron learning algorithm.

Lacking: meaningful prob. interpretation.

## 7.2. Reinforcement Learning

Basis are Markov decision processes. (MDP)

### 7.2.1 MDP

MDP is a tuple  $(S, A, \{P_{s,a}\}, \gamma, R)$

- $S$  set of states
- $A$  set of actions
- $P_{s,a}$  transition probabilities

$P_{s,a}$  prob. distribution over the state space  $S$   
 Distr. of states we will transition to  
 from state  $s$  taking action  $a$ .

- $\gamma \in [0, 1)$  discount factor
- $R : S \times A \rightarrow \mathbb{R}$  reward function,  
 often  $R : S \rightarrow \mathbb{R}$

Dynamics of the process

$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \rightarrow \dots$$

Total payoff:

$$R(s_0, a_0) + \gamma R(s_1, a_1) + \gamma^2 R(s_2, a_2) + \gamma^3 R(s_3, a_3) + \dots$$

$$\text{or } R(s_0) + \gamma R(s_1) + \gamma^2 R(s_2) + \gamma^3 R(s_3) + \dots$$

Goal:

$$\max_{\text{over actions } a \in \Pi} E[R(s_0) + \gamma R(s_1) + \gamma^2 R(s_2) + \dots]$$

Policy:  $\pi : S \rightarrow A : s \mapsto a = \pi(s)$

Value Function:

$$V^\pi(s) = E[R(s_0) + \gamma R(s_1) + \dots | s_0 = s, \pi]$$

$$V^*(s) = \max_{\pi} V^\pi(s)$$

It holds:  $V^*(s) = R(s) + \max_{a \in \Pi} \sum_{s' \in S} P_{sa}(s') V^*(s')$

Optimal policy:  $\pi^*(s) = \arg \max_{a \in \Pi} \sum_{s' \in S} P_{sa}(s') V^*(s')$

It holds: Any opt. policy  $\pi^*$  is independent of the initial state  $s \in S$  and achieves the same value  $V^{\pi^*}(s) = V^*(s)$ .