

Logistic Regression

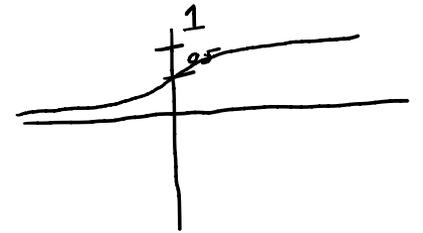
Discriminative probabilistic model / $P(y | \bar{x})$

(generative = $P(\bar{x}, y)$)

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

logistic

$$\frac{e^x}{1 + e^x}$$



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

Decision boundary?

return +1 if $P(y=+1 | \bar{x}) > 0.5$

$$\bar{w}^T f(\bar{x}) > 0$$

Training

For a dataset $(\bar{x}^{(i)}, y^{(i)})_{i=1}^D$, want to maximize

$$\prod_{i=1}^D P(y^{(i)} | \bar{x}^{(i)}) \quad \text{maximum likelihood}$$

$$\max_{\bar{w}} \sum_{i=1}^D \log P(y^{(i)} | \bar{x}^{(i)}) \quad \text{log likelihood}$$



$$\Rightarrow \min_{\bar{w}} \sum_{i=1}^D \underbrace{-\log P(y^{(i)} | \bar{x}^{(i)})}_{\text{loss}(\bar{x}^{(i)}, y^{(i)}, \bar{w})} \quad \left. \vphantom{\sum_{i=1}^D} \right] \begin{array}{l} \text{training objective} \\ \text{negative log likelihood} \\ \text{(NLL)} \end{array}$$

Need to compute $\frac{\partial}{\partial \bar{w}} \text{loss}(\bar{x}^{(i)}, y^{(i)}, \bar{w})$

LR Gradients

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

$$\frac{\partial}{\partial \bar{w}} \text{loss} = \frac{\partial}{\partial \bar{w}} \left[-\bar{w}^T f(\bar{x}) + \log(1 + e^{\bar{w}^T f(\bar{x})}) \right]$$

$$= -f(\bar{x}) + \frac{1}{1 + e^{\bar{w}^T f(\bar{x})}} - e^{\bar{w}^T f(\bar{x})} \cdot f(\bar{x})$$

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

$$\text{Update: } \bar{w} \leftarrow \bar{w} + \alpha f(\bar{x}) (1 - P(y=+1 | \bar{x}))$$

$P(y=+1 | \bar{x}) \approx 1$: no update

$P(y=+1 | \bar{x}) \approx 0$: perceptron update

Let $z = \bar{w}^T f(\bar{x})$

NLL: $\log(1 + e^z) - z$

