Linear Classification

A simple example

VS

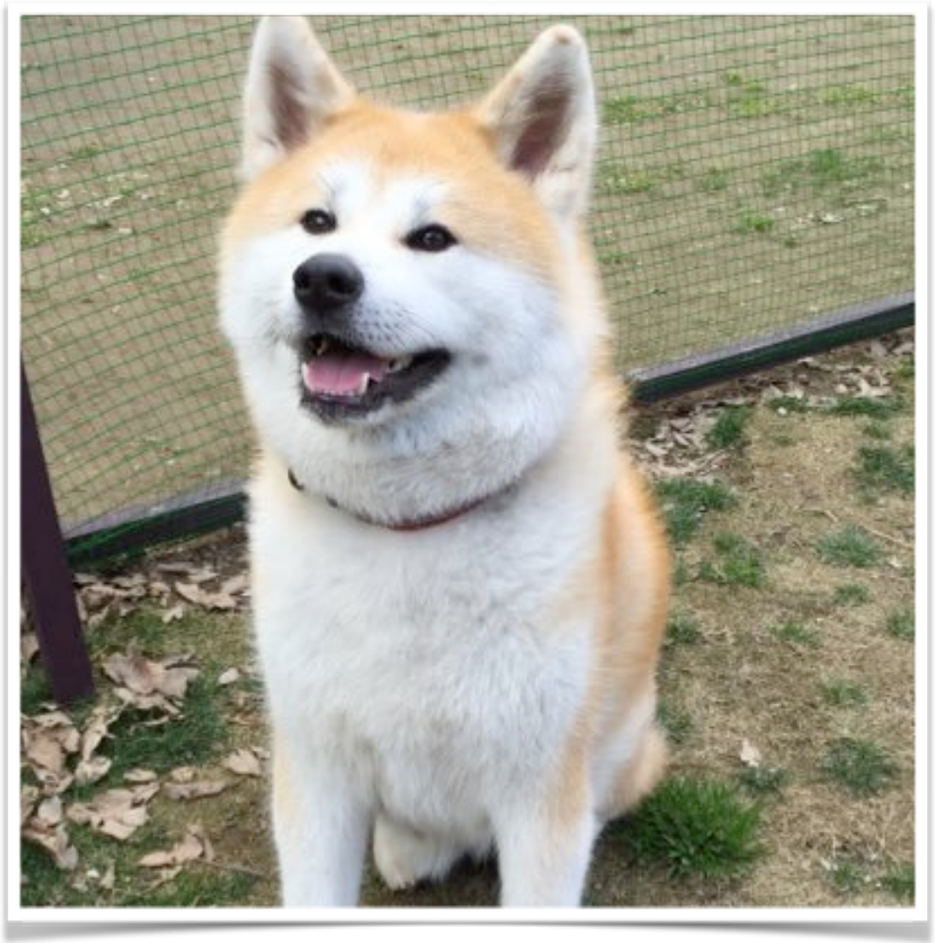

Linear binary classifier

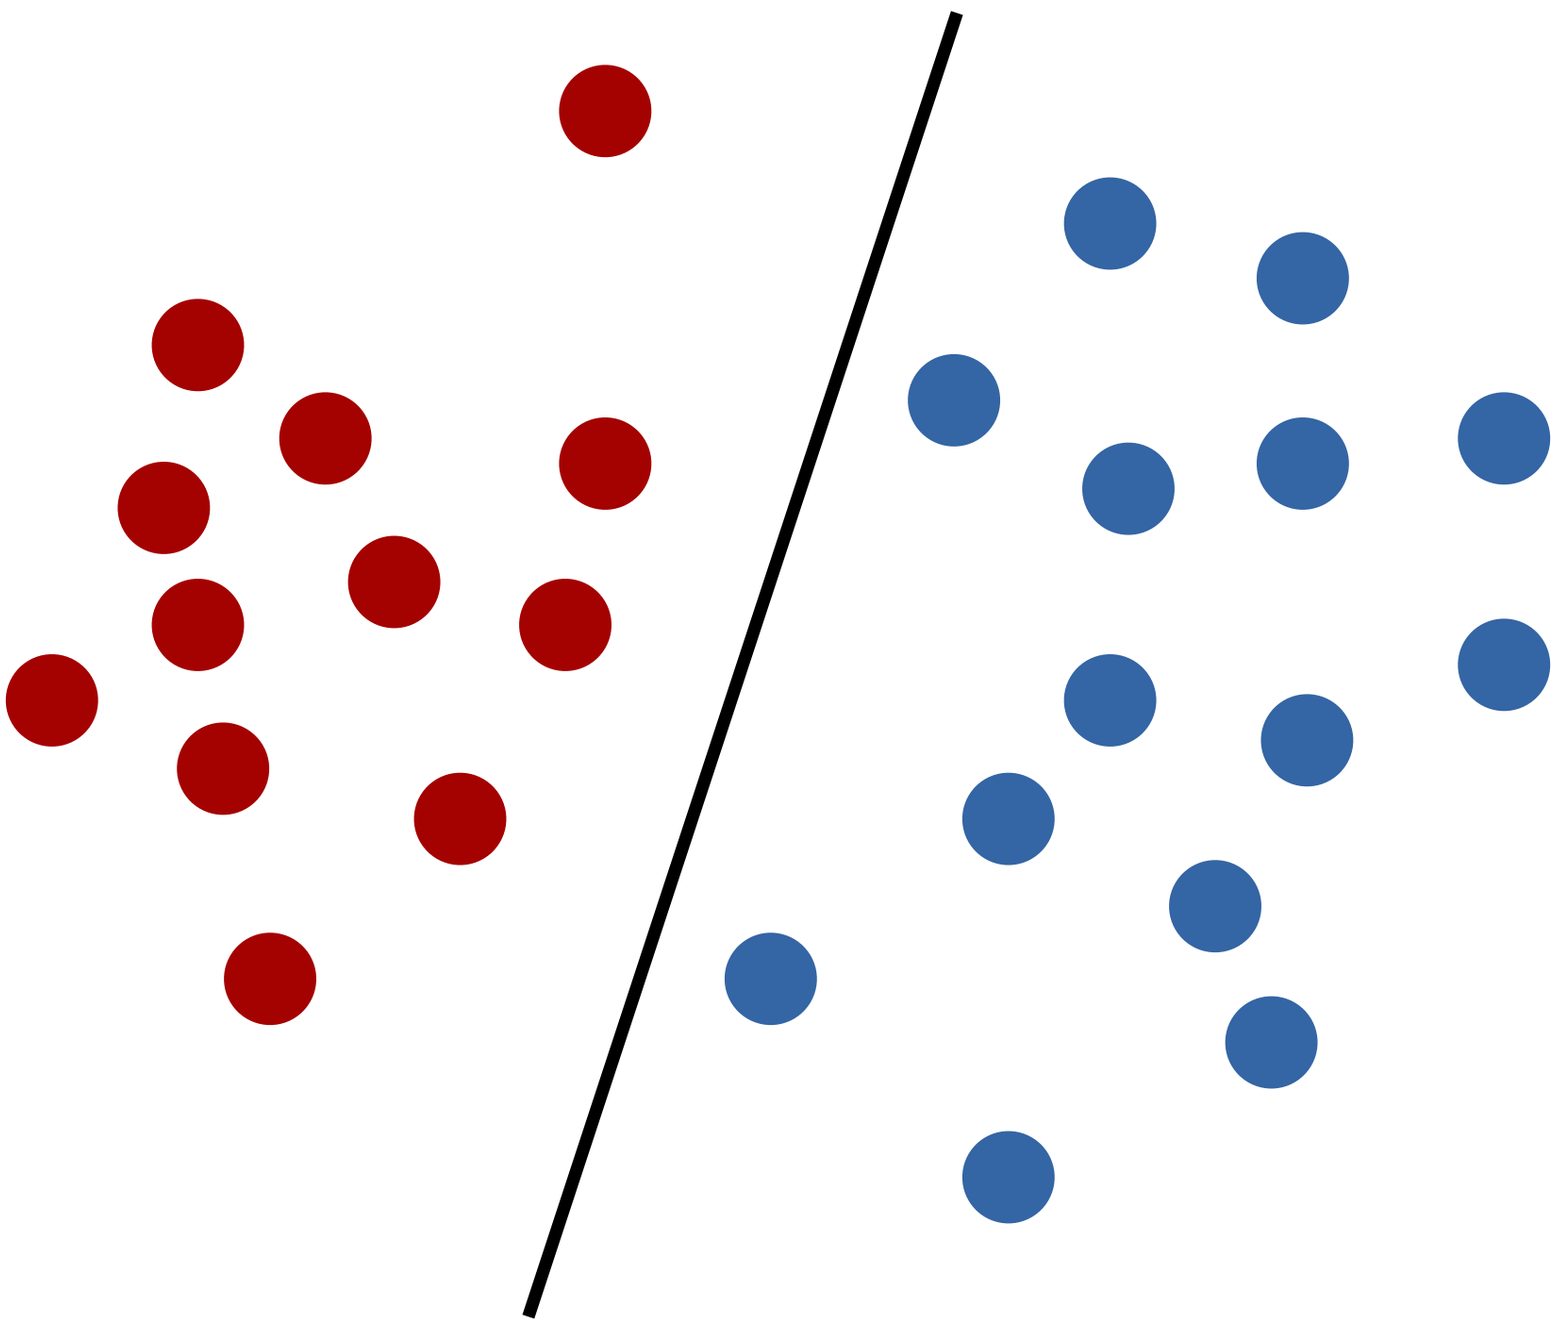

Linear binary classifier

- Input: $\mathbf{x}$ (tensor)
- Label: $y \in \{0,1\}$
- Parameters: $\mathbf{w}, b$
- Prediction $\hat{y} = \begin{cases} 1 & \text{if } \mathbf{w}^\top \mathbf{x} + b > 0 \\ 0 & \text{otherwise} \end{cases}$

Sigmoid function

- Maps $\mathbb{R}$ to $[0,1]$

- $o = \mathbf{w}^\top \mathbf{x} + b$

$$p(y = 1) = \sigma(o) = \frac{1}{1 + e^{-o}}$$

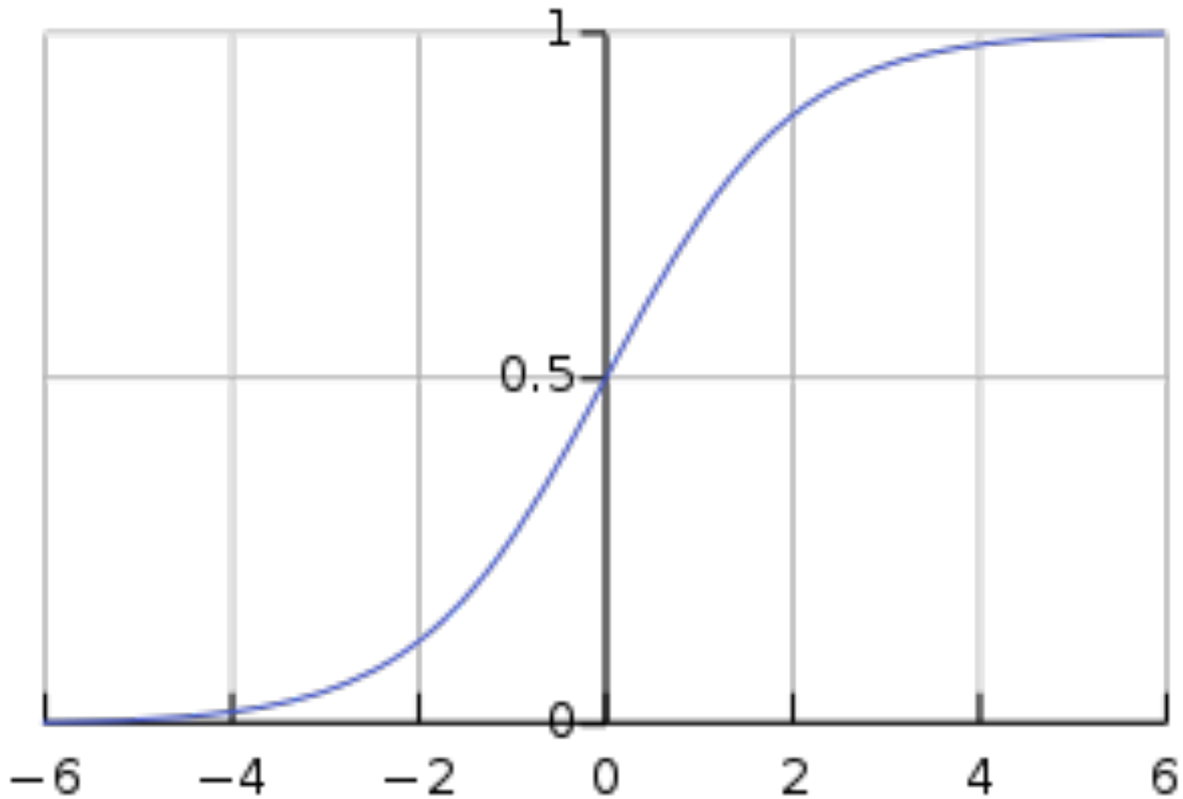

Logistic regression

- Input: $\mathbf{x}$ (tensor)
- Label: $y \in \{0,1\}$
- Parameters: $\mathbf{w}, b$
- $o = \mathbf{w}^\top \mathbf{x} + b$
  - $$\begin{aligned} p(y=1) &= \sigma(o) \\ p(y=0) &= 1 - \sigma(o) \end{aligned} \quad \leftrightarrow \quad p(y) = \sigma(o)^y (1-\sigma(o))^{1-y}$$
- Loss (negative log likelihood):
  - $$-\log p(y) = -y \log(\sigma(o)) - (1-y)\log(1-\sigma(o))$$