



Introduction to Machine Learning

Advanced Risk Minimization

0-1-Loss



Learning goals

- Derive the risk minimizer of the 0-1-loss
- Derive the optimal constant model for the 0-1-loss

Learning goals

- Derive the risk minimizer of the 0-1-loss
- Derive the optimal constant model for the 0-1-loss



0-1-LOSS


- Let us first consider a discrete classifier $h: \mathcal{X} \rightarrow \mathcal{Y}$.
- The most natural choice for $L(y, h)$ is the 0-1 loss

$$L(y, h(x)) = \mathbb{I}_{\{y \neq h(x)\}} = \begin{cases} 1 & \text{if } y \neq h(x) \\ 0 & \text{if } y = h(x) \end{cases}$$

- For the binary case ($g = 2$) we can express the 0-1-loss for a scoring classifier $f(x)$ based on the margin $\nu := yf(x)$

$$L(y, f(x)) = \mathbb{I}_{\{\nu < 0\}} = \mathbb{I}_{\{yf(x) < 0\}}.$$

- Analytic properties: Not continuous, even for linear f the optimization problem is NP-hard and close to intractable.



0-1-LOSS: RISK MINIMIZER

By the law of total expectation we can in general rewrite the risk as
(this all works for the multiclass case with 0-1)

$$\begin{aligned}\mathcal{R}(f) &= \mathbb{E}_{xy} [L(y, f(\mathbf{x}))] = \mathbb{E}_{\mathbf{x}} [\mathbb{E}_{y|\mathbf{x}} [L(y, f(\mathbf{x}))]] \\ &= \mathbb{E}_{\mathbf{x}} \left[\sum_{k \in \mathcal{Y}} L(k, f(\mathbf{x})) \mathbb{P}(y = k | \mathbf{x}) \right],\end{aligned}$$

with $\mathbb{P}(y = k | \mathbf{x})$ the posterior probability for class k . For the binary case we denote $\eta(\mathbf{x}) := \mathbb{P}(y = 1 | \mathbf{x})$ and the expression becomes

$$\mathcal{R}(f) = \mathbb{E}_{\mathbf{x}} [L(1, \pi(\mathbf{x})) \cdot \eta(\mathbf{x}) + L(0, \pi(\mathbf{x})) \cdot (1 - \eta(\mathbf{x}))].$$

