Theory of Machine Learning

Exercise sheet — Session 1

Exercise I (Risk relation between predictor and inverted predictor). Find the relation between the expected risk of f and $-f$ in the setting of binary classification with $0-1$ loss ($\mathcal{Y} = \{-1, 1\}$).

Exercise II (MSE loss). Consider the function $\ell : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+$ defined by $\ell(y, y') := (y - y')^2$.

1. What is the Bayes predictor for this loss? (*Hint:* convex optimization problem)
2. Given a random variable Y and constant $c \in \mathbb{R}$, prove that:

$$\mathbb{E}\left[(Y - c)^2\right] = \mathrm{Var}\,(Y) + (\mathbb{E}\,[Y] - c)^2\,.$$

3. Prove Question 1. using Question 2.

Exercise III (Absolute value loss). Consider the function $\ell : \mathbb{R} \times \mathbb{R} \to \mathbb{R}_+$ defined by $\ell(y, y') := |y - y'|$.

1. In the context of regression, does it seem like a reasonable loss function to you?
2. What is the Bayes predictor for this loss?

Exercise III (Bayes predictor for binary classification). In this exercise, we want to find the expression of the Bayes predictor for binary classification with $0-1$ loss ($\mathcal{Y} = \{0, 1\}$). As in the lecture, we set $\eta(x) := \mathbb{P}\,(Y = 1 \,|\, X = x)$ for all $x \in \mathcal{X}$ and we let

$$f^\star(x) = \begin{cases} 1 & \text{if } \eta(x) \geq 1/2 \\ 0 & \text{otherwise.} \end{cases}$$

1. Let $f : \mathcal{X} \to \mathcal{Y}$ be a predictor. Show that

$$\mathbb{P}\,(f(X) \neq Y \,|\, X = x) = \eta(x) \cdot \mathbb{P}\,(f(X) = 0 \,|\, X = x) + (1 - \eta(x)) \cdot \mathbb{P}\,(f(X) = 1 \,|\, X = x)\;.$$

2. Deduce that

$$\mathbb{P}\,(f^\star(X) \neq Y \,|\, X = x) = \min(\eta(x), 1 - \eta(x))\,.$$

3. Show that, for any predictor $f : \mathcal{X} \to \mathcal{Y}$,

$$\mathbb{P}\,(f(X) \neq Y \,|\, X = x) \geq \mathbb{P}\,(f^\star(X) \neq Y \,|\, X = x)\;.$$

4. Deduce that $f^\star$ is risk optimal, that is, for any predictor f,

$$\mathcal{R}(f^\star) \leq \mathcal{R}(f)\,.$$