Machine Learning Through Mathematics

Week 2 Exercise Sheet: Mathematical Foundations of Supervised Learning

Brief recap

- **Supervised learning** uses examples with known answers.
- A model tries to learn a function that maps inputs to outputs.
- Inputs are often written as vectors:

$$\mathbf{x} = (x_1, x_2, \ldots, x_n)$$

- **Regression** predicts continuous values, while **classification** predicts categories.
- A simple linear model has the form:

$$y = wx + b$$

- Training usually means minimising a loss function.

Exercises

Regression or classification?

For each situation, decide whether it is a **regression** or a **classification** problem.

a) Predicting the price of a used car.

b) Deciding whether an email is spam or not spam.

c) Predicting how many hours a student studies in a week.

d) Identifying whether a photo contains a cat, dog, or bird.

Linear models

A model is given by

$$\hat{y} = wx + b$$

with $w = 3$ and $b = -2$.

a) Compute $\hat{y}$ for $x = 4$.

b) Compute $\hat{y}$ for $x = 0$.

c) Describe in words what the weight w and bias b do.

d) Find an x such that $\hat{y} = 7$.

Inputs as vectors

A model takes a 3-dimensional input vector

$$\mathbf{x} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

and predicts

$$\hat{y} = x_1 + 2x_2 - x_3.$$

a) Compute $\hat{y}$ for

$$\mathbf{x} = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

b) Find one vector $\mathbf{x}$ such that $\hat{y} = 0$.

Loss functions

Suppose a model makes the following predictions (y is actual, $\hat{y}$ is the prediction:

$$(y, \hat{y}) = (5, 3),\ (2, 2),\ (7, 4)$$

Use the mean squared error

$$L = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$$

for the questions below.

a) Compute the squared error for each pair.

b) Compute the mean squared error.

c) Which prediction was the worst? Explain using the errors.

Modelling

Take the following data points and find the correct model linear model using the optimal choice we derived for linear regression in class.

$$(1, 2), (2, 4), (3, 6), (4, 8).$$