AIS5103 tutorial 1 problems

Exercise 1 — Single neuron

Given

$$x = 2,\ w = 1.5,\ b = -1,$$

calculate $z = w^T x + b$ and the output for the ReLU(z) activation function.

Exercise 2 — Affine transformation

Given

$$W = \begin{pmatrix} 1 & 2 \\ -1 & 1 \end{pmatrix},\ \boldsymbol{b} = \begin{pmatrix} 1 \\ -2 \end{pmatrix},\ \boldsymbol{x} = \begin{pmatrix} 2 \\ 1 \end{pmatrix},$$

calculate

$$\boldsymbol{z} = W\boldsymbol{x} + \boldsymbol{b}.$$

Then to reproduce it using PyTorch `nn.Linear`.

Exercise 3 — Why activation functions?

A neural network contains three nn.Linear layers but no activation functions. Show mathematically that the entire network can be reduced to a single affine transformation.

Exercise 4 — Gradient descent

Given the loss function

$$L(w) = (w - 3)^2,$$

start with $w = 0$ and $\eta = 0.1$.

Calculate the first **three gradient-descent updates by hand**.

Then reproduce the same calculation in PyTorch.

Exercise 5 1 − 4 − 1network

Show the PyTorch code for a regression problem with 4 data points as x = (0.0, 1.0, 2.0, 3.0), y=(1.0,2.0, 1.5, 2.7), fitting to a curve with the neural network architecture: one input, 4 hidden units with tanh activation function, and 1 output, with 8+4+1=13 parameters; where and what are these parameters in your code? Write the training loop. Show how the loss decreases over epochs, and draw the final fitted curve along with the data points.

Define the model and inspect the actual PyTorch model:

```
model = nn.Sequential(
    nn.Linear(1, 4),
    nn.Tanh(),
    nn.Linear(4, 1)
)
```