

Exercise Sheet #4

Problem 1 (*Exceptions in Computing Cross-Entropy Losses*)

Cross-entropy loss is a commonly used loss function in machine learning when dealing with classification problems. When working with neural networks to solve the classification task with C distinct classes, the output of the network might be a vector

$$\vec{p}_{\text{pred}} = (p_1, \dots, p_C)^T,$$

where the entry p_c resembles the probability that the input to the network was classified to class $c \in [1, C]$. The cross-entropy loss

$$\text{CELoss}[\vec{p}_{\text{pred}}, \vec{p}_{\text{targ}}] = - \sum_{c=1}^C (\vec{p}_{\text{targ}})_c \log[(\vec{p}_{\text{pred}})_c]$$

quantifies the difference between the outputted predicted class probabilities and the correct classification.¹

- (a) Implement cross-entropy loss in Python.
- (b) Create two custom exceptions:
 - an exception for when \vec{p}_{pred} or \vec{p}_{targ} contain values outside $[0, 1]$.
 - an exception for when the shapes of \vec{p}_{pred} and \vec{p}_{targ} are different.
- (c) For each of the two exceptions, give an example where it is raised. If the vectors contain values outside $[0, 1]$, handle the exception by normalizing the vectors, thereby avoiding a crash of the program.

¹Often the correct classification is unambiguous. Then

$$\vec{p}_{\text{targ}} = (0, \dots, 1, \dots, 0)^T.$$

Problem 2 (*Generating XOR*)

Generators in Python are a type of iterable that allow you to iterate through a sequence of values without generating the entire sequence at once. Instead, they produce values one at a time as needed, which can be more memory-efficient, especially for large datasets or infinite sequences. This ‘yield on demand’ is called *lazy evaluation*.

The XOR time series is defined by

$$\begin{aligned}x(t) &= \text{XOR}[x(t-1), x(t-2)] \\ &= [x(t-1) + x(t-2)] \pmod{2}\end{aligned}$$

(a) Write a function in Python that generates the XOR time series.

Hint: Use the `yield` keyword.

(b) Use that function to infinity print the XOR time series.