

Neural Networks from Scratch
Winter Term 2024

Exercise 1

- Return electronically until Tuesday, October 29, 2024, 09:00
- Include your names on the top sheet. Hand in only one PDF.
- A maximum of three students are allowed to work jointly in a group.
- When you include plots add a short explanation of what you expected and observed.
- Hand in source code if the exercise required programming. You can bundle the source code along with the PDF in a .zip file.

1.1 Reading

Read the following paper and provide a review as explained in the first lecture (see slides):

- Solveig Badillo, et al. 2020. An Introduction to Machine Learning.

1.2 Polynomial curve fitting

This exercise is meant as a gentle introduction to Python and the issue of overfitting.

- Install numpy on your machine, if you don't have Python (and/or work in Windows) it is easiest to install anaconda, see <https://docs.conda.io/projects/conda/en/latest/index.html>. If you have Python installed you can use `pip install numpy torch matplotlib`
- Start with the template `template.py` in the „template/” folder. The file should already produce two plots, with as of yet incorrect data.
- Extend the `ground_truth_function` to generate data of the shape $h(x) = \sin(2\pi x)$, as in the lecture.
- Extend the `error_function` to reflect the non-regularized error function from the lecture.
- For this exercise we will be using the `numpy` function `polynomial.Polynomial.fit` to fit the polynomial for us. If everything is setup correctly you should by now get a plot for a polynomial fit of third degree, similar to the one shown in the lecture.
- Create a new plot for an overfitted Polynomial of 11-th degree.
- Now vary the degree of the Polynomial from 0 to 11. With this data reproduce the plot of "Polynomial degree against the train and test error" from the lecture. Make use of the already created data and make sure to use the RMS-Error instead of the direct error.
- Now vary the size of the data, but keep the degree of the Polynomial constant. As a sensible starting point choose a 10-th degree polynomial. Then vary the sample size of the dataset from 10 until the RMS-Error of the train and test dataset are of similar size, i.e. the function doesn't over fit anymore. Note that the required sample size can easily be of the order of a few 100 for this to fully happen.