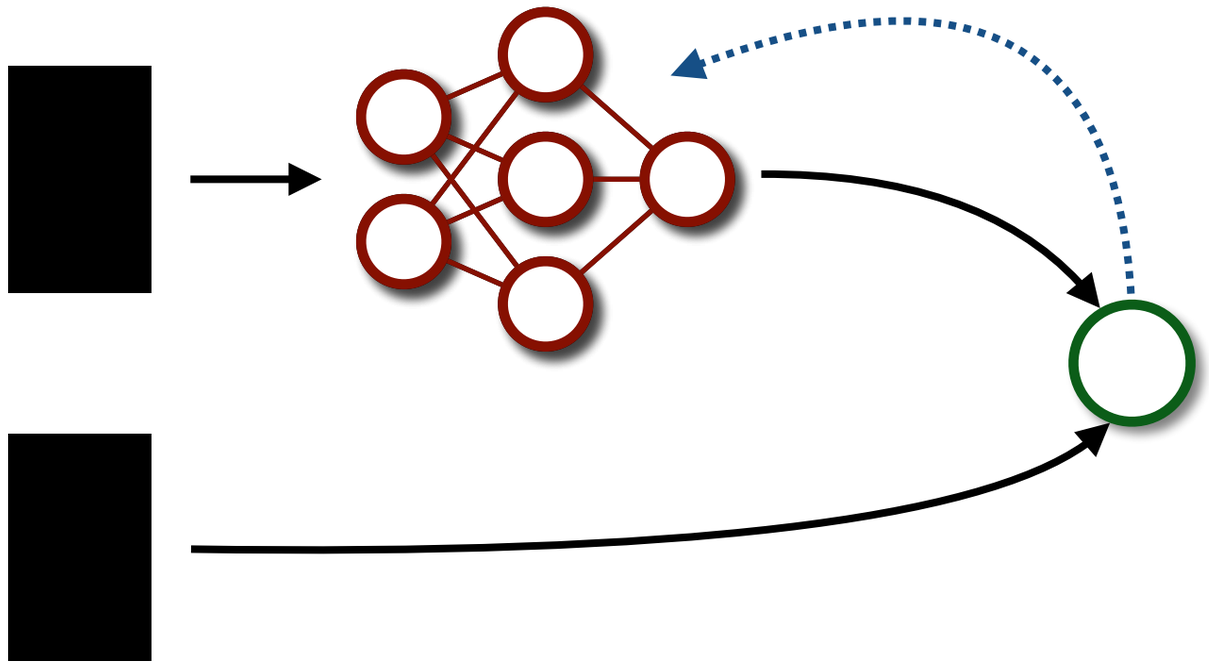

Neural Network Building Blocks (3/3)

Docent: Dennis Noll (RWTH)
Tutor: Boyang Yu (LMU)

Deep Learning School "Basic Concepts"

09.08.22

RWTH AACHEN UNIVERSITY

Exercise 3.1 - Regression

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- You are given data which follows a very complex function. Your task is to approximate this function using a deep neural network.
  - Open this notebook
  - Visualize the data
  - Create the model:
    - What is a suitable size?
    - How many inputs and outputs?
    - What are suitable activations?
  - Compile the model:
    - Which loss function should be used? (Documentation)
    - Which optimizer should be used?
  - Train the model and visualize its prediction. What do you observe?
  - Try to improve your model. You may try the following configurations. Describe your observations.
    - Different activation functions
      - ReLU, Sigmoid, Thanh
    - Different learning rates
      - 0.0001, 0.001, 0.01 ,0.1, 1, 10
  - What happens if you increase the uncertainty of the data points?

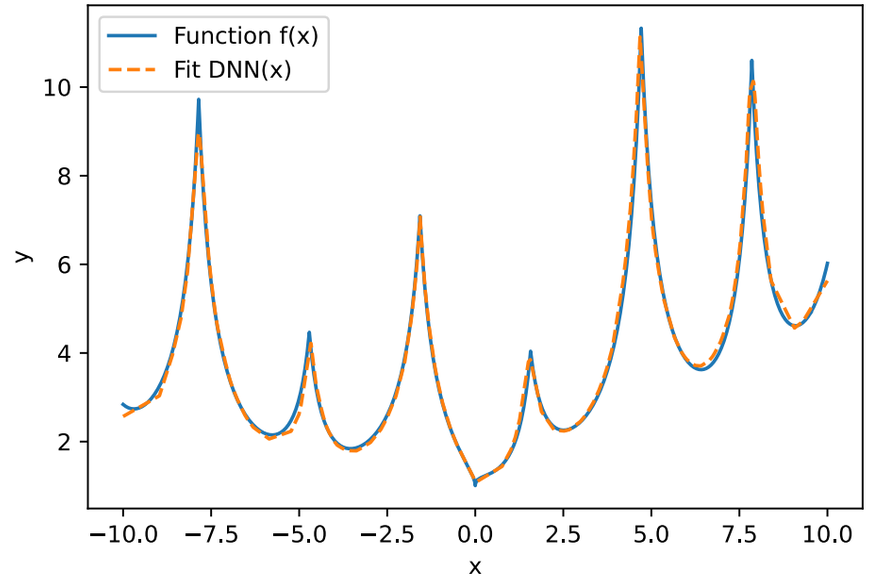

Exercise 3.1 - Regression

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- **Notes/Solutions**:

Exercise 3.2 - Classification (Gaussians)

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- You are given data from two classes. In each class the data follows a distribution out of one or many gaussian distributions with class dependent parameters. Your task is to build a model which can classify between the two classes.
  - Open this notebook
  - Visualize the data
  - Create the model:
    - What is a suitable size?
    - How many inputs and outputs?
    - What are suitable activations?
      - **Hint:** Think about the activation of the last layer.
  - Compile the model:
    - Which loss function should be used? (Documentation)
    - Which optimizer should be used?
  - Train the model
  - Visualize the output of the network along with the training data. Describe your observation.
  - Now we will make our exercise more difficult:
    - Make the functions more complex (N_PEAK) and train the classifier again. Describe your observations.
    - Raise the number of dimensions (N_DIM) to 2 (and 10) and train the classifier again. Describe your observations.

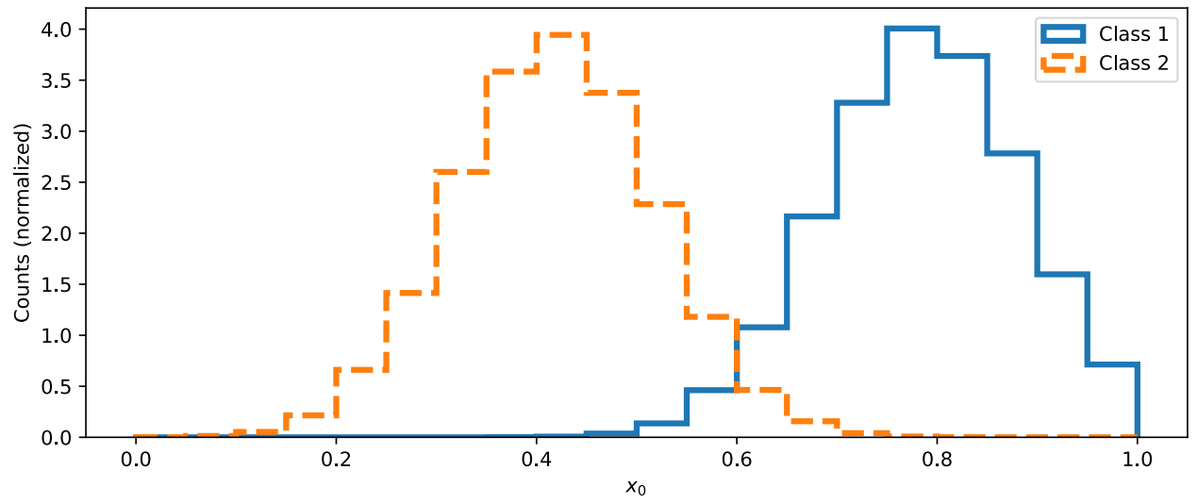

Exercise 3.2 - Classification (Gaussians)

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- **Notes/Solutions**:

Exercise 3.3 [advanced] - Classification

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- You are given training data from two one-dimensional gaussian probability distributions. The number of samples from each of the distributions is assumed to be very large. The parameter of the probability distributions are as follows:
  - Gauss: $f(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}$
  - Distribution A: $\mu = 0,\ \sigma = 1.5$
  - Distribution B: $\mu = 1,\ \sigma = 1.2$
- You train a sufficiently large network (f_{DNN}) with a single output utilizing a sigmoid activation with the cross-entropy to perform a classification between the two distributions.
- Tasks:
  - Using pen & paper: Predict the output of the network $f_{DNN}(x)$ in dependence of x.
    - **Hint:** Start with the cross entropy loss. What do you know about its derivative?
  - Use this notebook to perform the task. Compare the output of your training to your calculation.

Exercise 3.3 [advanced] - Classification

RWTH AACHEN UNIVERSITY
Dennis Noll - 09.08.22

- **Notes/Solutions**:

All Notebooks (Repo)

All Notebooks (Colab)