

Deep Learning Demystified: a practical, hands-on course

Date:	5 & 6 June 2019
Time:	10.00 – 16.00 h
Location	TU Delft, room t.b.a.
Course leader(s):	Dr. Jan van Gemert, Osman Kayhan, Paul van Gent
Days:	1.5
ECTS:	1.5 (attendance, including preparatory assignments and lab sessions)
Course fee:	Free for TRAIL/Beta/OML members, others please contact the TRAIL office
Registration:	www.rstrail.nl

Objectives

Participants will:

- Become familiar with the theoretical underpinning of neural networks and convolutional neural nets;
- Learn to approach deep learning from a practical perspective;
- Take part in a lab where they will learn to apply deep learning and have ample time for hands-on experience;
- Learn to work with the Deep Learning frameworks 'Tensorflow' and 'Keras';
- Learn to structure their Deep Learning projects effectively;
- Learn to spot common problems in model fitting, and how to effectively fix them.

Active participation during the theoretical lectures is stimulated and encouraged.

Course description

Deep Learning as a methodology has seen more and more adoption recently. From systems that detect cancer in x-rays better than doctors, to systems that automatically extract vehicle information from the registration plates, or systems that detect and label other road users. The success and safety of self-driving cars will be deeply rooted in this methodology. But what exactly is it? How does it work?

In this course you will learn the basics underpinning different types of deep networks, and how they can be used to handle different types of data. More importantly, you will get extensive hands-on experience with implementing deep networks using a popular toolkit: Tensorflow combined with Keras. This will prepare you for using Deep Learning in your own research.

The theoretical part of the course will be taught in the first morning by Assistant Professor Jan van Gemert. The first hands-on lab will be in the afternoon and taught by Osman Kayhan and Paul van Gent. Here you will learn to apply the models to several popular datasets. The second day will have a hands-on lab that focuses on structuring your projects effectively and tackling common training problems.

Assignment

There will be an assignment beforehand to help students get familiar with Python. This ensures you will have enough prerequisite knowledge to enter the course. You will need between 2 and 5 hours to go through it.

In the afternoon of the first day, and the second day there will be a lab session to get hands-on experience with the concepts discussed in the first morning.

Program

Day 1: Neural Networks and Convolutional Networks

Morning program (theory)

- What is a neural network, how does it function? Building intuition
- Gradient descent: teaching the neural net.
- Convolutions, how do they function? Building intuitions
- Training a convolutional neural network.

Afternoon program (hands-on lab)

Three exercises designed to teach you how to

- structure your data.
- use Keras.
- train the networks.

Day 2: Structuring Your Deep Learning Projects

Hands-on lab (approx. 1/2 to 2/3 day)

- Effectively structuring your deep learning project
- Common pitfalls and mistakes: why is my model not fitting?
- Transfer Learning: how to handle limited data.

Course material

TBD

Methodology

Python

Deep Learning

Tensorflow

Keras

Prerequisite

We assume a basis in the Python programming language. Materials and exercises are provided to familiarize yourself with Python.