

Course Syllabus

Artificial Neural Networks and Deep Learning



Semester & Location: Spring - DIS Copenhagen

Type & Credits: Core Course - 3 credits

Major Disciplines: Computer Science, Mathematics

Core Course Study Tours: London

Faculty Members:

Program Director: Iben de Neergaard ids@dis.dk

Time & Place:

Description of Course

Artificial Neural Networks are programs that write themselves when given an objective, some training data, and abundant computing power. Recently, these programs have brought about a wide array of future-like innovations, such as self-driving cars, face recognition, and human-like speech generators.

This course offers you an introduction to Deep Artificial Neural Networks (i.e. “Deep Learning”). With focus on both theory and practice, we cover models for various applications, how they are trained and tested, and how they can be deployed in real-world applications.

Course elements

- Python programming
- Machine learning basics
- Neural networks
- Convolutional neural networks
- Recurrent neural networks

Course outline

- Intro to machine learning and neural networks: supervised learning, linear models for regression, basic neural network structure, simple examples and motivation for deep networks. Textbook: parts of Bishop chapters 1 and 3, or Goodfellow chapter 5. (2 sessions)
- Lab 0: intro to tensorflow, simple ML examples. (2 sessions)
Maybe include exercises for a small assignment to be handed in.
- Neural networks: forward propagation, cost functions, error backpropagation, training by gradient descent, bias/variance and under/overfitting, regularization. Textbook: Bishop sections 5.1, 5.2, 5.3, 5.5, or Goodfellow chapter 5. (2 sessions)
- 1 week travel (core course week)
- Lab 1: neural networks. Exercises on NNs, solving a problem with NNs on tensorflow. Students should have studied at home and started working on the assignment. (2 sessions)
- Convolutional Neural Networks. Textbook: Goodfellow chapter 9 (1 session)
- Lab 2: convolutional networks. Exercises on CNNs, solving a problem with CNNs on tensorflow. Students should have studied at home and started working on the assignment. (1 sessions)
- Recurrent Neural Networks. Textbook: Goodfellow chapter 10 (1 sessions)
- Lab 3: recurrent networks. Exercises on RNNs, solving a problem with RNNs on tensorflow. Students should have studied at home and started working on the assignment. (1 sessions)
- Final project
- Possibly: limitations of deep learning (e.g. data-hungry, results are at most as good as the data, ethical issues, advanced topics).

Learning Objectives

Upon successfully completing the course, the student will be able to:

- Understand generic machine learning terminology
- Understand motivation and functioning of the most common types of deep neural networks
- Understand the choices and limitations of a model for a given setting
- Apply deep learning techniques to practical problems
- Critically evaluate model performance and interpret results
- Write reports in which results are assessed and summarized in relation to aims, methods and available data

Prerequisites

Basics of probability theory, linear algebra and calculus at university level. Programming skills (Python will be used throughout the course).

Faculty

xxx

Readings

Most of the learning will be based on parts of the following book chapters:

- Bishop, Pattern Recognition and Machine Learning. Chapters 1, 3, 5.
- Goodfellow et al., Deep Learning. Chapters 5, 6, 7, 9, 10

Additional possible sources include blog posts, videos available online, and scientific papers.

Expectations of the Students

Students are expected to reach the preparation goal leading up to each session. Students who have little or no experience coding in Python should either follow a Python tutorial before the course starts, or prepare to invest some hours getting up to speed with the language once we start. Students should have a working laptop computer. It is advised that each machine has a least 4 GB of RAM and a reasonable processor (if it's bought after 2012 you should be fine).

Evaluation

During the course you will hand in three assignments containing selected exercises solved in class. Furthermore, you will complete a larger project that uses tools which have been taught in the class. You will be allowed to define your own project, but you can also get assistance from the teacher.

Both project and assignments are group efforts. The teacher will rate all the assignments, but you will also participate using the peer evaluation system Peergrade.io, where each hand-in is double-blind peer-reviewed by 3-4 students which, together with the teacher's evaluation composes indicators towards the final grade. This creates more and fairer feedback for each group as well as evaluation that is less sensitive to mistakes.

During the programming projects, you are allowed to consult freely with any of the other students and the instructor. Contributions from other students, however, must be acknowledged with citations in your final report, as required by academic standards. Contributions to your presentations must similarly be acknowledged. Needless to say, the right to consult does not include the right to copy — programs, papers, and presentations must be your own original work.

Grading

Assignment	Percent
Mandatory assignments	40%
Final project	40%
Active participation	10%
Peer reviews	10%

Academic Regulations

Please make sure to read the [Academic Regulations](#) on the DIS website. There you will find regulations on:

- [Course Enrollment and Grading](#)
- [Attendance](#)
- [Coursework, Exams, and Final Grade Reports](#)