



21 – 25 AUGUST 2017 / UNIVERSITY OF COPENHAGEN

DEEP LEARNING FOR DATA ANALYTICS

Deep learning arguably is the hottest trend in data analysis. It has pushed the boundaries in Big Data Analytics and Artificial Intelligence and has been outperforming the state-of-the-art in many applications from various domains, for instance, various domains e.g. object classification in images, information retrieval including web search, natural language processing tasks such as automatic translation, and bioinformatics.

Not only large players such as Google and Facebook, but also more and more small and medium-sized companies are successfully applying deep learning techniques to solve commercially relevant problems in a broad variety of application areas as diverse as drug design, customer relation management and mortgage risk estimation.

This course will bring you to the forefront of deep learning, introducing the foundations as well as the newest tools and methods in this emerging field.

WHAT YOU WILL LEARN

By completing the course you will be able to set up and use basic deep learning techniques. In particular, you will learn how to use deep convolutional neural networks and recurrent neural networks for image and text analysis tasks. You will also be acquainted with a number of advanced tools and will become familiar with using appropriate computational resources to train and apply deep learning models.

The course will also teach the theoretical foundations of deep neural networks, which will provide the understanding necessary for adapting and successfully applying deep learning in your own applications.

COURSE CONTENT

Deep learning refers to machine learning algorithms that process data in multiple stages, each stage working on a different representation of the data. These representations are learned and allow for analyzing data on different levels of abstraction.

The course will cover the following methods and tools:

Core elements:

- Thorough introduction to the foundations of neural networks including their training e.i. backpropagation
- Introduction to convolutional neural networks and recurrent neural networks. In particular long short-term memory networks, also called LSTMs
- Training and applying convolutional and recurrent neural networks for text- and image analysis
- Making use of data augmentation and other preprocessing steps to further improve the generalization performance of such models
- Using modern software tools for deep learning like Theano or TensorFlow
- Application examples presented by experts with first-hand experience in applying deep learning in scientific and commercial applications
- Exploiting appropriate hardware systems to speed up the compute-intensive process of generating complex deep learning models, e.g. via graphics processing units

”I found the course to be of very high quality theoretically. The theoretical concept covered were very detailed and relevant.”

Anup Singh, General Manager, Mærsk Group on the course Big Data Analysis (2016)



Advanced topics (depending on participants' interests):

- Deep reinforcement learning i.e., how can deep neural networks learn to interact with an environment
- Introduction to Restricted Boltzmann Machines and Deep Belief Networks, which are deep generative models

All covered techniques can be easily implemented with Python, which will be the programming language used throughout the course.

All participants will receive a copy of the new MIT Press textbook *Deep Learning* by Ian Goodfellow, Yoshua Bengio and Aaron Courville (2016).



PARTICIPANTS

The course is strictly focused on deep learning, thus a background in statistics and/or conventional data analysis is assumed. This includes elementary linear algebra and calculus i.e., it is assumed that the participants can recall what a derivative and a scalar product is. Linear algebra and calculus can be refreshed using the course material sent to the participants before the course.

The course also assumes elementary programming knowledge. The little programming in the course will mostly be done in Python, which can be easily picked up using the course material.

COURSE DATES

5 days, 21-25 August 2017, 9:00 – 16:30 at the University of Copenhagen, Frederiksberg Campus.

COURSE DIRECTORS

Christina Lioma, Associate Professor, Department of Computer Science, University of Copenhagen

Christian Igel, Professor, Department of Computer Science, University of Copenhagen

OTHER COURSE TEACHERS

Mads Nielsen, Professor, Head of Department, Department of Computer Science, University of Copenhagen

Anders Søgaard, Professor, Department of Computer Science, University of Copenhagen

Fabian Gieseke, Assistant Professor, Department of Computer Science, University of Copenhagen

COURSE FEE

EUR 2,600/DKK 19,000 excl. Danish VAT. Fee includes teaching, course materials and all meals during the course.

FOR MORE INFORMATION AND REGISTRATION:
copenhagensummeruniversity.ku.dk