



AUGUST 13 - 17, 2018 / UNIVERSITY OF COPENHAGEN

DEEP LEARNING

Deep learning is arguably the hottest trend in data analysis. It has pushed the boundaries in Big Data Analytics and Artificial Intelligence (AI) and has been outperforming the state-of-the-art in numerous applications across a wide range of domains. These include object classification in images, information retrieval along with web search, natural language processing tasks such as automatic translation, and bioinformatics. Moreover, the deep neural network powered AlphaGo became the first computer program to beat the human champion in the game of Go. This is widely regarded as a milestone in AI development.

Increasingly, it is not only leading players such as Google and Facebook, but also small and medium-sized companies that are successfully applying deep learning techniques to solve commercially relevant problems in a broad variety of areas as diverse as drug design, customer relation management, and mortgage risk estimation. This course will give you detailed insight into deep learning, introducing you to the basics as well as to the latest 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. You will learn how to use deep convolutional neural networks and recurrent neural networks for image, text, and time series analysis tasks. You will also become acquainted with advanced tools and become familiar with using appropriate computational resources to train and apply deep learning models.

The course will also teach you the theoretical foundations of deep neural networks, which will provide you with 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 enable data to be analyzed at different levels of abstraction.

Core elements:

- Thorough introduction to the basics of neural networks including how to train them (e.g. back propagation)
- Introduction to convolutional neural networks
- Introduction to recurrent neural networks, for example long short-term memory networks (LSTMs) and gated recurrent units (GRUs), for time series modelling and predicting
- Training and applying convolutional and recurrent neural networks for text- and image analysis
- Utilizing data augmentation and other preprocessing steps to further improve the generalization performance of such models
- Introduction to generative adversarial networks (GANs)
- Using modern software tools for deep learning, in particular TensorFlow (used by DeepMind, Google Brain, Ebay, Twitter, Qualcomm, SAP, and many more) as well as Keras
- 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



Advanced topics (depending on participants' interests)

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

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

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

“A lot of information in a short amount of time – great!”

Former participant on Deep Learning 2017

PARTICIPANTS

The course is for professionals who need state-of-the-art skills in deep learning analytics. Participants should:

- Be acquainted with data analysis i.e. hold a relevant bachelor degree or equivalent and/or have several years of data analysis experience
- Have elementary programming knowledge
- Have an interest in programming (the programming on the course will mostly be done in Python, which you can easily pick up using the course material)
- Have a background in statistics and/or conventional data analysis (it is assumed that participants have elementary knowledge of linear algebra and calculus and can recall what a derivative and a scalar product is)

You will have the opportunity of refreshing your linear algebra and calculus as necessary using the course material, which will be sent out to participants before the start of the course.

COURSE DATES

5 days, August 13-17, 2018, 9:00 – 16:30 at the University of Copenhagen, South Campus.

COURSE DIRECTOR

Christian Igel, Professor, Department of Computer Science, Constituent Head of the SCIENCE AI Center, 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

Wouter Boomsma, Assistant Professor, Department of Computer Science, University of Copenhagen

Isabelle Augenstein, Assistant Professor, Department of Computer Science, University of Copenhagen

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

Akshay Sadananda Uppinakudru Pai, Postdoctoral Researcher, Department of Computer Science, University of Copenhagen

COURSE FEE

EUR 2,680/DKK 19,900 excl. Danish VAT. Fee includes teaching, course materials and all meals during the course.



FOR MORE INFORMATION AND REGISTRATION:
copenhagensummeruniversity.ku.dk