1 week deep learning hands-on course for companies

	Monday	Tuesday	Wednesday	Thursday	Friday
9-9.30	Breakfast	Breakfast	Breakfast	Breakfast	Breakfast
9.30-10.30	Lecture 1: Introduction to deep learning and statistical artificial intelligence	Lecture 3: Tricks of the trade - faster convergence and improved performance	Lecture 4: Recurrent neural networks (RNN)	Lecture 5: Un- and semi-supervised learning	Lecture 6: Reinforcement learning
10.30-11.00	Coffee break	Coffee break	Coffee break	Coffee break	Coffee break
11.00-12.00	Lecture 2: Feed-forward neural networks (FFNN)	Lecture 4: Convolutional neural networks (CNNs)	Lab 4: RNNs	Lab 6: un- and semi-supervised learning	Lab 7: Reinforcement learning
12.00-13.00	Lunch	Lunch	Lunch	Lunch	Lunch
13.00-14.00	Lab 1: Setting up software and Lab 2: FFNN	Lab 3: CNNs	Lab 4: continued	Lab 6: continued	Lab 7: continued
14.00-15.00	Lab 2: continued	Guest speaker 1: TBA	Guest speaker 2: TBA	Guest speaker 3: TBA	Guest speaker 4: TBA
15.00-15.30	Coffee break	Coffee break	Coffee break	Coffee break	Coffee break
15.30-16.30	Lab 2: continued	Lab 3: continued	Lab 5: take part in Kaggle competition	Lab 5: continued or finishing other labs	Lab 5: continued or finishing other labs

Dates: November 14th to 18th, 2016.

Location: TBA

Price including breakfast, lunch and access to computer resources: 20.000 DKK

Prerequisites: programming preferably in Python, basic probability theory and basic linear algebra.

Computer frameworks: TensorFlow and Theano/Lasagne

Personnel: Ole Winther, professor in Data science and complexity, DTU Compute, lectures and supervision; guest lecturers and TAs from DTU Compute.

Motivation: Statistical artificial intelligence and deep learning have received a lot of attention in the recent years due to the impressive advances within the field in areas such as speech recognition, image classification and artificial intelligence for games such as Google Deepmind's AlphaGo system. The leading technology companies see machine learning as a key to developing their services and a lot of resources are invested in order to apply the same type of methodology in many different application areas. These methods can be expected to play a major role in development of intelligent user-interfaces and automation.

The purpose of this course is to teach the participants about the latest developments in the field, about opportunities and pitfalls and give the participating companies access to computational frameworks that will allow them to go directly home and apply into in their own context. It covers both more well-established methods like feed-forward, convolutional and recurrent neural networks and frontiers like un-, semi- and reinforcement learning that can be expected to play a larger role in the coming years. The course is taught by leading experts in the field.

Detailed contents:

1. Introduction to statistical artificial intelligence and deep learning. The "deep learning revolution": Speech recognition, image classification and games case stories.
2. Feed-forward networks. The basic building block in deep neural network models. How to set training and how to train neural networks with the stochastic gradient descent methods. The error back propagation algorithm.
3. Tricks of the trade. Deep learning is a challenging to get to work in new contexts. The course will cover the state-of-the-art methods to ensure fast convergence of training and regularization methodology to ensure that flexible deep learning model finds solutions with good predictive performance. Generalization to data outside the training set is the key objective of predictive machine learning methods.
4. Convolutional neural networks (CNNs) are neural network models that use learned filters and other translationally invariant operations. CNN architectures and applications will be covered. CNNs have achieved unprecedented performance in image classification tasks and is the workhorse behind modern image and video processing for example in self-driving vehicles.

5. Recurrent neural networks (RNNs) are neural network models for sequence data. RNNs are applied in many types of tasks involving sequences: speech recognition, machine translation, dialogue systems and biological sequence analysis. Architectures such a gated recurrent units (GRUs) and long-short term memory (LSTM), applications and challenges in training RNNs will be covered.
6. Un- and semi-supervised. In unsupervised learning we don't have access to labeled data. Since most natural learning is unsupervised and most data we have access to is unlabeled. Unsupervised learning is still not as developed as supervised learning but is arguably the most important part of machine learning. State-of-the-art unsupervised learning with autoencoders and their extension to semi-supervised learning (combining labeled and unlabeled data) will be covered.
7. Reinforcement learning is about learning in a context where the consequence of an action now is only revealed later. This is for example relevant in games, driving and advertisement. The course will cover the basic concepts, policy gradient methods and go through Google DeepMind's AlphaGo as a case.

There will be hand-on labs to cover all topics in the lectures and guest speakers from companies that apply deep learning. The participants will learn how to use both CPU, fast GPU (graphics processing unit) based programming frameworks and cloud computing solutions with many examples. Two frameworks will be used: Theano/Lasagne which is a more well-established framework that has been around from some time and TensorFlow from Google that has been introduced recently and gained widespread use. After the course the participants will be able to run all code on their own applications.

For more information please contact Ole Winther, olwi@dtu.dk or Christina Horn Nexø, chne@dtu.dk