



The Australian Association of Computational Neuroscientists and Neuromorphic Engineers

... presents a short-course, in association with NeuroEng 2016

DEEP LEARNING: AN INTRODUCTION WITH MATLAB

Deep learning is one of the most exciting technology breakthroughs of the last three years, and will change our society in the coming decade.

For example, in March 2016 Google's AlphaGo software defeated a world-champion human player of the game of Go, which had never previously happened. This success, which had been predicted to still be ten years away, relied on deep learning.

Meanwhile, deep learning methods now frequently achieve *better-than-human* performance on difficult pattern recognition tasks such as determining the breed of dog or cat in an image, automatic speech transcription, image-based medical diagnosis, question answering, automatic image captioning and artistic style transfer.

But what is deep learning? Is it just artificial neural networks? If so, why are they suddenly so successful when they were mediocre for so long? How do I program them myself? Do I need a GPU? Does a deep neural network take weeks to train? How do I build intuition?

This one-day hands-on course aims to answer these questions, with a hands-on introduction.

Course Style and Logistics

This course offers face-to-face instruction and hands-on programming tutorials. A small amount of lecture-style material will be combined with extensive exercises in Matlab.

- Instructor:** A/Prof. Mark McDonnell, University of South Australia
When: Wednesday 30th November 2016 (the day after NeuroEng finishes)
Where: University of Queensland, St Lucia campus (room TBA)
Time: 9:00 for 9:30 to 16:30 (lunch: 12:45-13:30; breaks: 11:00-11:15; 15:30-15:45)
Cost: \$50 registration for catering/administration
Prerequisites: Working knowledge of Matlab programming; undergrad. linear algebra
Provided: Access to a PC running Matlab, for use during course;
Course notes and access to Matlab source code introduced in the course;
Lunch, coffee, morning and afternoon tea.

There is a maximum of 15 attendees: pre-registration is essential.

You may bring your own laptop pre-installed with matlab and matconvnet, and wireless access to the internet for downloading code and data, but I won't be able to guarantee support during the workshop for trouble-shooting package/toolbox installation.

Is this Course for Me?

There are many deep-learning packages (e.g. Caffe, Theano, Tensorflow, DIGITS), and online tutorials/courses. But often a strong background in computer science and/or mathematics is assumed. This can be an obstacle to getting started with implementations or to building intuition.

This course aims to introduce deep networks and build your intuition.

So it might be for you if...

- ... you don't know anything about deep learning, but are curious about it, and how to implement its methods by yourself, using classic example datasets;
- or if ... you can implement algorithms in code, but don't know how to implement deep neural networks and could use some expert help to get started;
- or if ... you don't have a strong mathematics or programming background, and want to build some intuition, aided by visualisations, about how and what deep neural networks learn;
- or if ... you can implement basic deep neural networks, but will benefit from tuition on the latest breakthroughs, such as batch normalization and residual connections.

Course Content

1. What is "deep learning"?
 - Focus: The Imagenet competition; resources for deep-learning
 - Matlab tutorial: visualising images and running pre-trained models
2. Basics of supervised learning
 - Focus: Classifying MNIST handwritten digit images
 - Matlab tutorial: K-nearest neighbours; linear least-squares classifiers
3. Training neural network classifiers
 - Focus: The backpropagation and stochastic gradient descent algorithms
 - Matlab tutorial: Training classical single-hidden-layer neural networks
4. Deep convolutional neural networks
 - Focus: Classifying the CIFAR-10 tiny-images dataset
 - Matlab tutorial: Training a simple convolutional neural network classifier
5. The overfitting problem
 - Focus: cross-validation, regularization and data-augmentation
 - Matlab tutorial: Training an "Alex-net"
6. Tricks and tips for accelerating training and enhancing performance
 - Focus: visualisation of what a deep-neural network learns; going deeper
 - Matlab tutorial: Training a "VGG-net" and a "residual network"
7. Beyond image classification
 - Focus: a quick survey of other deep-learning methods
 - Matlab tutorial: Recurrent networks