



jumping rivers

Python and Tensorflow

Deep learning is a cutting-edge machine learning technique for classification and regression. In the past few years, it has produced state-of-the-art results in fields such as image classification, natural language processing, bioinformatics and robotics. This course will cover the main ideas of deep learning, and how to implement it in practice with tensorflow: a software framework for efficient and scalable deep learning. We'll introduce the high-level *keras* library, which allows deployment of standard neural networks with just a few lines of code, as well as the details of raw *tensorflow*, allowing a deeper level of customisation.

Course Outline

- Supervised learning: how to frame classification, regression and prediction tasks
- Multilayer perceptrons: a simple neural network architecture
- Training neural networks: stochastic gradient descent and back propagation
- Deep learning: learning complex features from big data
- Parameter tuning: regularisation methods to avoid overfitting, choosing activation functions
- Convolutional neural networks: learning from image data
- Scaling: using GPUs with big data

Prior Knowledge

Basic knowledge of Python and a familiarity with mathematical concepts such as exponentiation, logarithms and summation notation. Some knowledge of calculus, matrix algebra and probability would be helpful but is not essential. The ideas needed from these subjects will be explained during the course.