

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.



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

Learning Outcomes

On successful completion of the course, delegates will...

- be able to leverage the popular **tensorflow** python package to build feed forward neural networks for both regression and classification tasks.
- gain an understanding of deep learning terminology and the mechanisms behind how such networks are trained.
- gain some experience in integrating with data pre-processing and pipeline creation from the **sklearn** python package.
- experience building convolutional neural networks (CNNs) and monitoring their training via **tensorboard**.

Attendee Feedback

- “The instructor was great and really helpful and willing to answer loads of wider questions as well as troubleshoot issues with the assignments.”
- “The delivery of the training was excellent.”
- “Things that were good included the math background and understanding the concept of neural nets, also how willing the instructor was to answer wider questions.”