MISS 2016

31 July - 6 Aug 2016 Favignana, Sicily

Medical Imaging meets Machine Learning

Ben Glocker

Imperial College London, United Kingdom

Lecture 1: Random Forests and their applications in medical imaging

In this lecture, we will discuss Random Forests, a powerful machine learning technique that has been successfully applied to many different image analysis tasks. We will cover the basics of Random Forests, explore how they can be applied to classification and regression problems in medical imaging, and how they compare to other supervised learning techniques. We will look at example applications including anatomy localisation, multi-organ segmentation, and image synthesis. We will also discuss some subtle, non-obvious implementation details that are important when using Random Forests but which are rarely mentioned in research papers.

Lecture 2: Solving continuous problems with discrete optimization: Tracking and Registration with Markov Random Fields

The estimation of continuous parameters such as the transformation in image registration normally requires the employment of continuous optimization methods, e.g. gradient descent. Continuous optimization, however, is generally very sensitive to initialization and prone to converge to local minima of the cost function. In this presentation, we will discuss how we can employ powerful discrete optimization techniques such as iterative graph-cuts to solve continuous problems. Discrete optimization is computationally efficient, and often yields very good solutions close to the global optimum. It also allows the incorporation of any non-differentiable objective function, which enables novel similarity measures for multi-modal image registration. We will discuss two example applications of continuous problems that are solved with discrete optimization, interventional tool tracking and 3D non-rigid registration.

Lecture 3: Deep Learning for Brain Lesion Segmentation

The presentation will cover recent work on using deep learning for brain lesion segmentation. We discuss a very efficient multi-scale, 3D convolutional neural network approach which achieves state-of-the-art performance on various, challenging lesion segmentation tasks including segmentation of brain tumors, ischemic stroke, and traumatic brain injuries. The network architecture is devised to learn both local and contextual features related to lesions from multi-modal 3D brain MRI scans. We further explore some insights that can be obtained when looking at the feature maps that are learned in the hidden layers of the network, and we discuss current limitations and opportunities for future research in deep learning for medical image analysis.