Andrea Vedaldi

University of Oxford, United Kingdom

Lecture 1: Advanced Convolutional Neural Networks

Since their introduction in 2012 with the AlexNet model, modern Deep Convolutional Neural Network (DCNNs) architectures have stormed computer vision. In this lecture, I will review some of the most significant advances in DCNNs obtained in the past three years. I will start by reviewing several design guidelines and architecture enhancements that have almost tripled the performance of the original AlexNet. Then, I will discuss applications of CNNs beyond image classification. I will review the idea of fully-convolutional network for image segmentation as well as region-based networks for object instance detection. Finally, I will briefly summarise a number of other applications of CNNs to illustrates their breadth.

Lecture 2: Understanding CNNs using visualisation and transformation analysis

Learnable representations, and Deep Convolutional Neural Network (DCNNs) in particular, have become the preferred way of extracting visual features for image understanding tasks, from object recognition to semantic segmentation. Part of the reason for the power of these models is that they contain millions of parameters learned from millions of images. However, this also means that deep networks are black boxes of difficult interpretation. In this lecture, I will focus on the problem of understanding deep networks. First, I will discuss a number of techniques to visualise the feature learned in a CNN: inversion, exploring which visual information is retained in a deep representation, activation maximisation, which finds patterns that maximally excite representation components, and caricaturisation, which emphasise which visual properties are captured by different representation layers. Then, I will discuss methods to study how image transformations are encoded by representations, and whether different representations are equivalent.