

"Deep Learning and its Applications to Generative Models"

Dr David Pfau



Deep learning has swiftly risen to prominence within machine learning in the last several years, resurrecting research into artificial neural networks and finding applications across a wide range of unexpected domains. In this talk I'll focus on several ways of unifying deep learning with unsupervised learning, otherwise known as generative models. By drawing on ideas from variational inference, a branch of machine learning closely connected to statistical physics, it is possible to use deep neural networks to perform approximate Bayesian inference in a class of models known as variational autoencoders. These models can both infer global variables that explain variations in the data, as well as generate novel data outside the training set. Another approach to unsupervised learning is to model low-level conditional dependencies between pixels. A number of closely related models of this type (PixelRNN, PixelCNN, WaveNet and ByteNet) have recently achieved state-of-the-art results in image compression, text-to-speech synthesis and language modeling. Lastly, generative adversarial networks (GANs) approach unsupervised learning from a dramatically different perspective of a minimax game between two opponents, one of which generates data and the other classifies data as real or generated. GANs have managed to produce extremely compelling samples in many cases, and have deep connections with many other branches of machine learning. Together, these methods illustrate the variety of ways that deep learning can be used to realize the potential of unsupervised learning.