

Graham Taylor: Learning Multi-scale Temporal Dynamics with Recurrent Neural Networks

The last three years have seen an explosion of activity studying recurrent neural networks (RNNs), a generalization of feedforward neural networks which can map sequences to sequences. Training RNNs using backpropagation through time can be difficult, and was thought up until recently to be hopeless due to vanishing and exploding gradients used in training. Recent advances in optimization methods and architectures have led to impressive results in modeling speech, handwriting and language. Applications to other areas are emerging. In this talk, I will review some recent progress on RNNs and discuss our work on extending and improving the Clockwork RNN (Koutnick et al.), a simple yet powerful model that partitions its hidden units to model specific temporal scales. Our “Dense clockworks” are a shift-invariant form of the architecture which we show to be more efficient and effective than their predecessor. I will also describe a recent collaboration with Google in which we apply Dense clockworks to authenticating mobile phone users based on the movement of the device as captured by the accelerometer and gyroscope.