Neural Stochastic Differential Equations for Time Series Modelling

James Foster (Oxford University, UK)

Stochastic differential equations (SDEs) are a popular model for describing continuous-time phenomena and have seen particular success in the pricing and hedging of financial derivatives. However, given the current data science revolution and following the seminar paper “Neural Ordinary Differential Equations”, it is natural to investigate how SDE methodologies could be improved using tools from machine learning. This has led to several recent works on so-called “Neural SDEs”, which seek to combine the modelling capabilities of SDEs with the flexibility and efficient training of neural networks. In this talk, I will give an overview of these developments and show how SDEs can be viewed as time series models that fit nicely with well-known ideas from data science, such as generative adversarial networks (GANs).

The course focuses on differential deep learning (DL), arguably the strongest application. We will show how standard DL trains neural networks (NN) on punctual examples, whereas differential DL teaches them the shape of the target function, resulting in vastly improved performance, illustrating it with a number of numerical examples, both idealized and real world. We will also discuss how to apply differential learning to other ML models, like classic regression or principal component analysis (PCA).