Deep neural networks as structure preserving optimal control problems by **Brynjulf Owren**

Abstract
A deep neural network model consists a large number of layers, each with a number of parameters associated to them. In supervised learning, these parameters are optimised to match training data in the best possible way. The data are propagated through the layers by nonlinear transformations, and in an important subclass of models (ResNet) the transformation can be seen as the numerical flow of a certain type of continuous vector field. Ruthotto and Haber (2017) as well as Cheng et al. have experimented in using different type of vector fields to improve the deep learning model. In particular it is of interest that the trained model has good long time behaviour and are stable in the deep limit, when the number of layers tends to infinity. The models presented in the literature have certain builtin structural properties, they can for instance be gradient flows or Hamiltonian vector fields.