Distribution Learning Via Neural ODEs

Flow-based methods for sampling and generative modeling use continuous-time dynamical systems to represent transport maps that push a source measure forward to a target measure. In this talk, we examine the approximation and statistical properties of Neural Ordinary Differential Equations (Neural ODEs) in this context. Given a fixed set of independent and identically distributed samples from a target distribution, the goal is to either estimate the underlying density or generate new samples. While many existing approaches constrain trajectories to be straight lines (i.e., with zero acceleration), we show that a specific class of curved trajectories can lead to improved approximation and learning.