Diffusion models for inverse problems

Valentin De Bortoli & Jiaming Song
ENS CNRS & NVIDIA

Generative modeling is the task of drawing new samples from an underlying distribution known only via an empirical measure. There exists a myriad of models to tackle this problem with applications in image and speech processing, medical imaging, forecasting and protein modeling to cite a few. Among these methods diffusion models are a new powerful class of generative models that exhibit remarkable empirical performance. They consist of a ``noising'' stage, whereby a diffusion is used to gradually add Gaussian noise to data, and a generative model, which entails a ``denoising'' process defined by approximating the time-reversal of the diffusion. In this tutorial we discuss three aspects of diffusion models. First, we will present some of their theoretical guarantees with an emphasis on their behavior under the so-called manifold hypothesis. Such theoretical guarantees are non-vacuous and provide insight on the empirical behavior of these models.

Diffusion models can also be trained for specific inverse problems, but such models are limited to their particular use cases and are expensive to train. This talk introduces several of my recent works on using the same, generic diffusion model for solving different inverse problems. First, we will present Denoising Diffusion Restoration Models (DDRM) for solving noisy, linear inverse problems. Next, we will mention PhysDiff, which incorporates physical constraints into a motion diffusion model. Finally, we introduce Pseudoinverse-guided Diffusion Models (ΠGDM). Despite using a generic diffusion model, ΠGDM achieves the performance of diffusion models specifically trained on the inverse problem. The flexibility of ΠGDM also allows it to solve a much wider set of inverse problems.