

SEQUENTIAL MONTE CARLO GUIDANCE OF (DISCRETE) DIFFUSION MODELS

FREDRIK LINDSTEN

Diffusion models have become a prominent approach to generative modelling in machine learning. These models are trained by simulating some “destructive” stochastic process, initialized at training data samples and with a limiting distribution that is easy to sample from. The generative model is obtained by learning how to approximately reverse the stochastic process. Most applications of diffusion models have been for continuous data and use Gaussian diffusion as the stochastic process. However, the same idea can be applied also to discrete data by suitable choices of destructive processes, e.g. based on discrete Markov chains and the introduction of absorbing states. The performance and applicability of diffusion generative models have further been improved by guidance, a technique for steering the generative process based on some auxiliary information or external model. Guidance can be used both for conditional generation (e.g. with classifier guidance) and for improved sample quality (discriminator guidance). In this talk I will discuss how Sequential Monte Carlo can be used for guidance of diffusion models. I will focus on the discrete setting where conventional score-based guidance techniques are not readily applicable.

Based on joint work with Filip Ekström Kelvinius (Discriminator Guidance for Autoregressive Diffusion Models, AISTATS 2024, <https://arxiv.org/abs/2310.15817>)