

Denoising diffusion models without diffusions

Denoising diffusion models have enabled remarkable advances in generative modeling across various domains. These methods rely on a two-step process: first, sampling a noisy version of the data—an easier computational task—and then denoising it, either in a single step or through a sequential procedure. Both stages hinge on the same key component: the score function, which is closely tied to the optimal denoiser mapping noisy inputs back to clean data. In this talk, I will introduce an alternative perspective on denoising-based sampling that bypasses the need for continuous-time diffusion processes. This framework not only offers a fresh conceptual angle but also naturally extends to discrete settings, such as binary data. Joint work with Saeed Saremi and Ji-Won Park (<https://arxiv.org/abs/2305.19473>, <https://arxiv.org/abs/2502.00557>).

Orateur: BACH, Francis