

Exploring With Denoising Diffusion Probabilistic Models

Tzer-jen Wei

College of Artificial Intelligence, National Yang Ming Chiao Tung University

E-mail: tjwei@nycu.edu.tw

Abstract

Recent advances in DDPM (Denoising Diffusion Probabilistic Models) allow us to generate high-fidelity and quality images using deep models. It is the current state-of-the-art of image generative models and beats GAN-based models in various tasks.

This talk will first briefly overview DDPM, including its relationship to score-based generative models, the diffusion process, stochastic differential equations, classifier guidance, and classifier-free guidance methods.

We will then introduce a few simple experiments to explore the conditional probability space corresponding to the diffusion models. We can see the denoising process by running the algorithm on low-dimensional spaces. We can use a technique akin to classifier guidance to guide the model to generate images with arbitrary geometric shapes or that satisfy other easily defined conditions by combining an unconditional DDPM with some known probability density. We will investigate related issues in a classifier-free guidance setting and contrast the outcomes.

Keywords Denoising Diffusion Probabilistic Models, score based generative model

References

- [1] HO, Jonathan; SALIMANS, Tim. Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598, 2022.
- [2] DHARIWAL, Prafulla; NICHOL, Alexander. Diffusion models beat gans on image synthesis. Advances in Neural Information Processing Systems, 2021, 34: 8780-8794.
- [3] Khruklov, Valentin and I. Oseledets. "Understanding DDPM Latent Codes Through Optimal Transport." ArXiv abs/2202.07477 2022.