Generative modeling through diffusions

Eduardo García-Portugués • Department of Statistics, UC3M

Diffusion models have gained significant attention in the field of generative modeling, emerging as a powerful alternative to Generative Adversarial Networks (GANs) and Variational Autoencoders (VAEs). These models rely on a gradual noise removal process to generate high-quality samples from data distributions. Despite their growing popularity, understanding and implementing diffusion models from scratch can be challenging due to their mathematical complexity, requiring from stochastic differential equations and Langevin dynamics, and their computational requirements. The literature on generative diffusion models is vast. A great resource is https://yang-song.net/blog/2021/score/ and references and notebooks therein.

This thesis will focus on two key objectives: (1) reviewing and digesting the recent advances in diffusion models to provide a comprehensive overview of their development, and (2) training a basic diffusion model to generate images from scratch, analyzing its performance and limitations. This work aims to provide an accessible introduction to diffusion models while also contributing practical insights through hands-on model training.