

Bias-variance Tradeoff in Diffusion Models

Joakim Andersson Svendsen*

February 2026

Abstract

Denoising Diffusion Probabilistic Models (DDPMs) are state-of-the-art generative models whose training objective takes the form of a mean squared error. Despite this apparent similarity to a regression problem, the relevance of the classical bias–variance trade-off in DDPMs is not well understood. Unlike supervised learning, where prediction error and generalization can be directly assessed, generative models must be evaluated through indirect measures such as sample quality and diversity, complicating any direct transfer of bias-variance intuition.

This thesis examines how the bias-variance framework can be meaningfully interpreted in the context of DDPM training. We analyze the DDPM objective as a regression problem with a stochastic target and perform an empirical study on the MNIST dataset, focusing on training dynamics, noise-prediction residuals across diffusion timesteps, and the evolution of generated samples. Model behavior is assessed using loss-based diagnostics alongside established generative evaluation metrics, including the Inception Score and the Fréchet Inception Distance.

The results indicate that persistent error in the learned denoising function is the dominant factor limiting sample fidelity, while variance-related effects such as training instability (e.g. mode collapse) are not observed in these experiments. Most improvements in generative quality occur early in training and then level off, with sample diversity preserved throughout. These findings suggest that the classical bias–variance trade-off does not carry over directly to diffusion models: bias primarily governs fidelity through persistent denoising error, whereas variance plays a secondary role by influencing training stability and, indirectly, diversity.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden.
E-mail: Joakim.a.svendsen@gmail.com. Supervisor: Chun-Biu Li.