Paper Summary: Denoising Diffusion Probabilistic Models (Ho et al., 2020)

Alvin

September 18, 2026

Introduction

This paper introduces denoising diffusion probabilistic models (DDPMs) and demonstrates that a certain parameterization allows these DDPMs to produce high quality samples that are competitive against generative adversarial networks (GANs), autoregressive models, and variational autoencoders (VAEs). A diffusion model is a Markov chain that learns to reverse a destructive diffusion process: the *forward process* is the fixed Markov chain with noising transitions to gradually transform an image into pure noise. The *reverse process* is the Markov chain with learnable denoising transitions which generates an image from pure noise.

Methods

Since the exact likelihood is intractable, the model is trained by maximizing the evidence lower bound (ELBO) on the data log-likelihood. As noise increments are sufficiently small in the forward diffusion process, learned reverse transitions can also be conveniently treated as conditional Gaussians to make the reverse process computationally tractable. In training loss, the decomposition into independent terms per-timestep allows optimized training over random t. Training loss is computed using KL divergences between the learned reverse transitions and the true posteriors.

The model is reparameterized for predicting the added noise instead of the posterior mean in favor of higher sample quality as validated by the authors' empirical results. The paper then demonstrates that this training strategy is mathematically identical to denoising score-matching while the sampling method resembles annealed Langevin dynamics. Considering this connection, the authors experiment with simplifying their objective by using an unweighted variational bound:

$$L_{\text{simple}}(\theta) = \mathbb{E}_{t,\,\mathbf{x}_0,\,\boldsymbol{\epsilon}}\left[\left\|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta\big(\sqrt{\bar{\alpha}_t}\,\mathbf{x}_0 + \sqrt{1-\bar{\alpha}_t}\,\boldsymbol{\epsilon},\, t\big)\right\|^2\right] \tag{1}$$

By discarding these weights, the new variational bound down-weights loss terms at low timesteps and instead emphasizes denoising at higher timesteps. The authors find that this modified reweighting improves sample quality while log-likelihoods become less competitive.

Key Contributions

The paper's primary contribution is their approach of ϵ-prediction which is a parameterization strategy where the model learns to predict the noise added at different timesteps. In addition, the paper demonstrates that a modified unweighted variational bound with a simplified training objective produces substantially improved sample perceptual-quality while sacrificing more competitive log-likelihoods. The paper also finds that their ϵ-prediction is mathematically equivalent to denoising score matching, which holds valuable practical insights relevant in further DDPM-related development.

Results

On unconditional CIFAR-10, the paper's model achieves an FID of 3.17 and Inception score of 9.46, with uncompetitive log-likelihoods (≤ 3.75 bits/dim). A hasty small-scale reproduction achieves an FID of 54.97.