
How to train your Diffusion model faster?

Shivam Pal
(231110608)
pshivam@cse.iitk.ac.in

Abstract

Diffusion models are very effective for generating images, but they usually require a lot of compute power (training time). To make these models more efficient, the state-of-the-art approach is patch diffusion, which uses random small size patches instead of full resolution image. This reduces the training time significantly. However, the quality of generated images is reduced due to patch-only training. We propose a retrieval guidance based approach to patch diffusion. This uses nearest neighbor embeddings together with patches. Our approach achieves better generation quality with training time similar to patch diffusion.