

## A BACKGROUND

### A.1 DENOISING DIFFUSION PROBABILISTIC MODEL (DDPM)

DDPM is a latent-variable generative model that gradually transforms a noise distribution into a data distribution  $x_0 \sim q(x_0)$  (Ho et al., 2020). DDPM consists of a forward process  $q$  that iteratively adds a noise on the data distribution, and a reverse process  $p$  that iteratively denoises a noise distribution toward a final data distribution. The forward process adds a Gaussian noise to  $x_t$  using a Markov process according to a variance schedule  $\{\beta_t\}_{t=1}^T$ :

$$q(x_{1:T}|x_0) := \prod_{t=1}^T q(x_t|x_{t-1}), \quad q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I) \quad (4)$$

Ho et al. state that it is possible to sample  $x_t$  from  $x_0$  directly, using the notation  $\alpha_t := 1 - \beta_t$  and  $\bar{\alpha}_t := \prod_{s=0}^t \alpha_s$ :

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I) \quad (5)$$

$$= \sqrt{\bar{\alpha}_t}x_0 + \epsilon\sqrt{1 - \bar{\alpha}_t}, \quad \epsilon \sim \mathcal{N}(0, I) \quad (6)$$

Using Bayes theorem, posterior  $q(x_{t-1}|x_t, x_0)$  is also a Gaussian distribution with mean  $\tilde{\mu}_t(x_t, x_0)$  and variance  $\tilde{\beta}_t$ :

$$q(x_{t-1}|x_t, x_0) = \mathcal{N}(x_{t-1}; \tilde{\mu}_t(x_t, x_0), \tilde{\beta}_t I), \quad (7)$$

$$\text{where } \tilde{\mu}_t(x_t, x_0) := \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}x_0 + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}x_t \quad \text{and} \quad \tilde{\beta}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}\beta_t \quad (8)$$

With sufficiently large  $T$  and a well defined  $\beta_t$ , the latent  $x_T$  becomes nearly an isotropic Gaussian distribution. Assuming this, to sample from the data distribution  $q(x_0)$ , we can first sample from an isotropic Gaussian distribution and then iteratively apply  $q(x_{t-1}|x_t)$  to obtain  $x_0$ . However,  $q(x_{t-1}|x_t)$  depends on the entire data distribution so it is hard to exactly compute when the data distribution is unknown. As a result, we train a neural network to predict a mean  $\mu_\theta$  and a diagonal covariance matrix  $\Sigma_\theta$ :

$$p_\theta(x_{0:T}) := p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t), \quad p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (9)$$

The network is trained by optimizing the usual variational bound on negative log likelihood,  $L_{vlb}$ :

$$L_{vlb} := L_0 + L_1 + \dots + L_{T-1} + L_T \quad (10)$$

$$L_0 := -\log p_\theta(x_0|x_1) \quad (11)$$

$$L_{t-1} := D_{KL}(q(x_{t-1}|x_t, x_0) || p_\theta(x_{t-1}|x_t)) \quad (12)$$

$$L_T := D_{KL}(q(x_T|x_0) || p(x_T)) \quad (13)$$

Ho et al. identify that training the model to predict  $\epsilon$  in Eq. 6 improves sample quality than directly predicting  $\mu_\theta(x_t, t)$ . Therefore,  $L_{vlb}$  is simplified to:

$$L_{\text{simple}} = E_{t, x_0, \epsilon} [||\epsilon - \epsilon_\theta(x_t, t)||^2] \quad (14)$$

When the training is done, we can sample from the data distribution by inserting the predicted  $\epsilon_\theta(x_t, t)$  to the equation:

$$x_{t-1} = \frac{1}{\sqrt{1 - \beta_t}} \left( x_t - \frac{\beta_t}{\sqrt{1 - \alpha_t}} \epsilon_\theta(x_t) \right) + \sigma_t z_t, \quad (15)$$

where  $z_t \sim \mathcal{N}(0, I)$  and  $\sigma_t^2$  is a variance which is set to  $\sigma_t^2 = \beta_t$ .

DDPM shows a powerful performance on image generation but it has a severe drawback of significantly slow sampling speed. To sample one image, it should feedforward a neural network for each denoising step, total  $T$  times. DDIM (Song et al., 2020) accelerates the sampling speed of DDPM (Appendix A.2).

## A.2 DENOISING DIFFUSION IMPLICIT MODEL (DDIM)

DDIM generalizes DDPM as a class of non-Markovian diffusion processes (Song et al., 2020):

$$q_\sigma(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\sqrt{\alpha_{t-1}}\mathbf{x}_0 + \sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \frac{\mathbf{x}_t - \sqrt{\alpha_t}\mathbf{x}_0}{\sqrt{1 - \alpha_t}}, \sigma_t^2 \mathbf{I}) \quad (16)$$

Consequently, the reverse process becomes

$$\mathbf{x}_{t-1} = \underbrace{\sqrt{\alpha_{t-1}} \left( \frac{\mathbf{x}_t - \sqrt{1 - \alpha_t} \epsilon_t(\mathbf{x}_t)}{\sqrt{\alpha_t}} \right)}_{\text{"predicted } \mathbf{x}_0 \text{"}} + \underbrace{\sqrt{1 - \alpha_{t-1} - \sigma_t^2} \cdot \epsilon_t(\mathbf{x}_t)}_{\text{"direction pointing to } \mathbf{x}_t \text{"}} + \underbrace{\sigma_t \mathbf{z}_t}_{\text{random noise}} \quad (17)$$

When  $\sigma_t = \sqrt{(1 - \alpha_{t-1}) / (1 - \alpha_t)} \sqrt{1 - \alpha_t / \alpha_{t-1}}$  for all  $t$ , the forward process becomes Markovian which means that the reverse process becomes a DDPM. When  $\sigma_t = 0$ , the forward process becomes deterministic and produces high quality samples much faster.

## B QUALITATIVE RESULTS



Figure A1: Image samples of MDM+ADM trained on CIFAR-10 dataset.







