

From VAE to Diffusion Model

- From VAE to Diffusion Model
 - Variational Autoencoder (VAE)
 - Denoise Diffusion Probabilistic Model (DDPM)
 - Forward (diffusion) and Backward (denoise) process
 - Variational EM to optimize θ
 - Interpretation of the Training and Sampling Process
 - Denoising perspective

Varational Autoencoder (VAE)

In the normal auto-encoder (AE) model, for a data distribution $p(\mathbf{x})$, we first encode \mathbf{x} using $q(\mathbf{z}|\mathbf{x})$ and decode it using $p(\mathbf{x}|\mathbf{z})$. We need to optimize the loglikelihood of \mathbf{x} for a given encoding function $q(\mathbf{z}|\mathbf{x})$. This gives the following loss function:

$$\mathcal{L}_{ae}(\mathbf{x}) = \mathbf{E}_{q(\mathbf{z}|\mathbf{x})}[p(\mathbf{x}|\mathbf{z})] + P_\alpha$$

where P_α is regularization term. However it is unclear how to devise such regularization term in principle.

Based on AE, variational AE (VAE) derived the loss function in a probabilistic manner. We starts from $\log p(\mathbf{x})$:

$$\log p(\mathbf{x}) = \log \int p(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z} \tag{1}$$

$$= \log \int \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})}q(\mathbf{z}|\mathbf{x})d\mathbf{z} \tag{2}$$

$$\geq \int \log \frac{p(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})}q(\mathbf{z}|\mathbf{x})d\mathbf{z} \tag{3}$$

$$= \mathbf{E}_q[\log p(\mathbf{x}|\mathbf{z})] - \int \log \frac{p(\mathbf{z})}{q(\mathbf{z}|\mathbf{x})}q(\mathbf{z}|\mathbf{x})d\mathbf{z} \tag{4}$$

$$= \mathbf{E}_q[\log p(\mathbf{x}|\mathbf{z})] + \text{KL}(q(\mathbf{z}|\mathbf{x})||p(\mathbf{z})) \tag{5}$$

The first term in the above equation is the log-likelihood of decoder output, while the second term minimize the KL divergence between encoder output and the target encoder distribution. Now the P_α in the Eq (1) has a probabilistic definition.

Denoise Diffusion Probabilistic Model (DDPM)

Forward (diffusion) and Backward (denoise) process

In the DDPM, we start from \mathbf{x}_0 , whose distribution is unknown. At each step t , a diffusion process is used:

$$\mathbf{x}_t = \alpha_t \mathbf{x}_{t-1} + \beta_t \boldsymbol{\epsilon}_t$$

where $\boldsymbol{\varepsilon}_t$ draws from zero-mean unit-variance Gaussian distribution. Additionally, $\alpha_t^2 + \beta_t^2 = 1$. We have the following attributes regarding \boldsymbol{x}_t :

1. **Forward Process:** given \boldsymbol{x}_{t-1} , it is straightforward to know that $p(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})$ follow a Gaussian distribution with $\alpha_t \boldsymbol{x}_{t-1}$ as mean and $\beta_t^2 \mathbf{I}$ as the variance.
2. **Fast Forward Process:** A nice property of DDPM is that the conditional distribution of \boldsymbol{x}_t given \boldsymbol{x}_0 can be calculated explicitly without going through the recursive process, i.e.,

$$\begin{aligned}\boldsymbol{x}_t &= \alpha_t \boldsymbol{x}_{t-1} + \beta_t \boldsymbol{\varepsilon}_t \\ &= \alpha_t \alpha_{t-1} \dots \alpha_1 \boldsymbol{x}_0 + (\alpha_t \dots \alpha_2) \beta_1 \boldsymbol{\varepsilon}_1 + \dots + \beta_t \boldsymbol{\varepsilon}_t\end{aligned}\tag{6}$$

Except for the first term in Eq. (6), each term is a zero-mean, unit-variance Gaussian noise, therefore, Eq (6) can be also written as:

$$\boldsymbol{x}_t = \bar{\alpha}_t \boldsymbol{x}_0 + \bar{\beta}_t \bar{\boldsymbol{\varepsilon}}_t\tag{7}$$

where $\bar{\alpha}_t = \prod_{\tau=1}^t \alpha_\tau$, $\bar{\beta}_t = \sqrt{1 - \bar{\alpha}_t^2}$ and $\bar{\boldsymbol{\varepsilon}}_t$ is again a zero-mean, uni-variance Gaussian.

3. **Reverse Process:** However, we don't know the conditional distribution of \boldsymbol{x}_{t-1} given \boldsymbol{x}_t . We only know that for small enough β_t , it is still a Gaussian distribution. We use neural network (with parameter θ) to estimate the mean and variance, given \boldsymbol{x}_t and t , i.e.,

$$p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t) = \mathcal{N}(\boldsymbol{x}_{t-1}; \boldsymbol{\mu}_\theta(\boldsymbol{x}_t, t), \boldsymbol{\Sigma}_\theta(\boldsymbol{x}_t, t))$$

4. **Conditional Reverse Process:** Though we don't know the explicit form of reverse probability, a nice property of diffusion model is that $p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0)$ is a Gaussian distribution. This can be proved by the following deduction:

$$p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0) = \frac{p(\boldsymbol{x}_t, \boldsymbol{x}_{t-1}|\boldsymbol{x}_0)}{p(\boldsymbol{x}_t|\boldsymbol{x}_0)} = \frac{p(\boldsymbol{x}_t|\boldsymbol{x}_{t-1})p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_0)}{p(\boldsymbol{x}_t|\boldsymbol{x}_0)}$$

Since we know that $p(\boldsymbol{x}_\tau|\boldsymbol{x}_0)$ is a Gaussian distribution $\mathcal{N}(\boldsymbol{x}_\tau; \bar{\alpha}_\tau \boldsymbol{x}_0, \bar{\beta}_\tau \mathbf{I})$ for $\tau > 1$ from **Forward Process**, we then have the following equations:

$$p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0) \sim \exp\left(-\frac{1}{2\beta_t^2}(\boldsymbol{x}_t - \alpha_t \boldsymbol{x}_{t-1})^2\right) \cdot \exp\left(-\frac{1}{2\bar{\beta}_{t-1}^2}(\boldsymbol{x}_{t-1} - \bar{\alpha}_{t-1} \boldsymbol{x}_0)^2\right) \cdot \exp\left(\frac{1}{2\beta_t^2}(\boldsymbol{x}_t - \bar{\alpha}_t \boldsymbol{x}_0)^2\right)$$

Note that the above equation can be written as:

$$\begin{aligned}p(\boldsymbol{x}_{t-1}|\boldsymbol{x}_t, \boldsymbol{x}_0) &\sim \exp\left(-\frac{1}{2}(a\boldsymbol{x}_{t-1}^2 - 2b\boldsymbol{x}_{t-1} + c)\right) \\ &\sim \exp\left(-\frac{1}{2 \cdot 1/a}(\boldsymbol{x}_{t-1} - \frac{b}{a})^2\right)\end{aligned}$$

with (note that $\alpha_t^2 + \beta_t^2 = 1$, $\bar{\alpha}_t^2 + \bar{\beta}_t^2 = 1$ and $\bar{\alpha}_t = \bar{\alpha}_{t-1} \alpha_t$)

$$\begin{aligned}
a &= \frac{\alpha_t^2}{\beta_t^2} + \frac{1}{\beta_{t-1}^2} = \frac{\alpha_t^2(1 - \bar{\alpha}_{t-1}^2) + \beta_t^2}{\beta_t^2(1 - \bar{\alpha}_{t-1}^2)} = \frac{1 - \bar{\alpha}_t^2}{\beta_t^2(1 - \bar{\alpha}_{t-1}^2)} = \frac{\bar{\beta}_t^2}{\beta_t^2 \bar{\beta}_{t-1}^2} = \left(\frac{\bar{\beta}_t}{\beta_t \bar{\beta}_{t-1}}\right)^2 \\
b &= \frac{\alpha_t}{\beta_t^2} \mathbf{x}_t + \frac{\bar{\alpha}_{t-1}}{\beta_{t-1}^2} \mathbf{x}_0 \\
\frac{b}{a} &= \frac{\alpha_t(1 - \bar{\alpha}_{t-1}^2)}{1 - \bar{\alpha}_t^2} \mathbf{x}_t + \frac{\beta_t^2 \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t^2} \mathbf{x}_0 \\
&= \frac{\alpha_t \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2} \mathbf{x}_t + \frac{\bar{\alpha}_{t-1} \beta_t^2}{\bar{\beta}_t^2} \mathbf{x}_0
\end{aligned}$$

Therefore, we can see that $p(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$ is again Gaussian with $\frac{b}{a}$ as its mean and $\sqrt{\frac{1}{a}}$ as its variance. In other word, we can re-parameterize $\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0$ by

$$\mathbf{x}_{t-1} = \frac{\alpha_t \bar{\beta}_{t-1}^2}{\bar{\beta}_t^2} \mathbf{x}_t + \frac{\beta_t^2 \bar{\alpha}_{t-1}}{\bar{\beta}_t^2} \mathbf{x}_0 + \frac{\beta_t \bar{\beta}_{t-1}}{\bar{\beta}_t} \boldsymbol{\varepsilon}$$

with $\boldsymbol{\varepsilon}$ a sample from zero-mean, unit-variance Gaussian.

Using the **Fast forward** property, we know that

$$\mathbf{x}_0 = \frac{1}{\alpha_t} (\mathbf{x}_t - \bar{\beta}_t \bar{\boldsymbol{\varepsilon}}_t)$$

By combing the above 2 equations, we have:

$$\mathbf{x}_{t-1} = \frac{1}{\alpha_t} \left(\mathbf{x}_t - \frac{1 - \alpha_t^2}{\bar{\beta}_t} \bar{\boldsymbol{\varepsilon}}_t \right) + \frac{\beta_t \bar{\beta}_{t-1}}{\bar{\beta}_t} \boldsymbol{\varepsilon}$$

Variational EM to optimize θ

With the above properties, we can now derive the variational EM algorithm to maximize data distribution $p_\theta(\mathbf{x}_0)$ with respect to θ . Since only \mathbf{x}_0 is observed, and $\mathbf{x}_1, \dots, \mathbf{x}_T$ are latent, they can be treated as the latent variable \mathbf{z} in Eq. (5). Using $q = p(\mathbf{x}_1, \dots, \mathbf{x}_T | \mathbf{x}_0)$ and follow Eq. (3), we can have the following:

$$\log p_\theta(\mathbf{x}_0) \geq \mathbf{E}_q \left[\log \frac{p(\mathbf{x}_0, \dots, \mathbf{x}_T)}{q(\mathbf{x}_1, \dots, \mathbf{x}_T | \mathbf{x}_0)} \right] \quad (8)$$

At the same time, we can use chain rule of probability to factorize $p(\mathbf{x}_0, \dots, \mathbf{x}_T)$ in the following form:

$$p(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) p(\mathbf{x}_{0:T-1} | \mathbf{x}_T) \quad (9)$$

$$= p(\mathbf{x}_T) p(\mathbf{x}_{T-1} | \mathbf{x}_T) p(\mathbf{x}_{0:T-2} | \mathbf{x}_T, \mathbf{x}_{T-1}) \quad (10)$$

$$= p(\mathbf{x}_T) p(\mathbf{x}_{T-1} | \mathbf{x}_T) p(\mathbf{x}_{0:T-2} | \mathbf{x}_{T-1}) \quad (11)$$

$$= p(\mathbf{x}_T) \prod_{t=1}^T p(\mathbf{x}_{t-1} | \mathbf{x}_t) \quad (12)$$

Note that Eq (10) is possible because given \mathbf{x}_{T-1} , $\mathbf{x}_{0:T-2}$ is independent of \mathbf{x}_T .

Similarly, we have:

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) = \prod_{t=1}^T p(\mathbf{x}_t|\mathbf{x}_{t-1}) \quad (13)$$

we also note that for $t > 1$, $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)$ because \mathbf{x}_t is conditional independent of \mathbf{x}_0 given \mathbf{x}_{t-1} . We can further reforulate $q(\mathbf{x}_t|\mathbf{x}_{t-1})$, $\forall t > 1$ by

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) = q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0) = \frac{q(\mathbf{x}_t, \mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_{t-1}|\mathbf{x}_0)} = \frac{q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)q(\mathbf{x}_t|\mathbf{x}_0)}{q(\mathbf{x}_{t-1}|\mathbf{x}_0)} \quad (14)$$

Cominging Eq. (12 - 14), we have:

$$\begin{aligned} \log \frac{p(\mathbf{x}_0, \dots, \mathbf{x}_T)}{q(\mathbf{x}_1, \dots, \mathbf{x}_T|\mathbf{x}_0)} &= \log p(\mathbf{x}_T) + \sum_{t=1}^T \log p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) \\ &\quad - \log p(\mathbf{x}_1|\mathbf{x}_0) - \sum_{t=2}^T (\log q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) + \log q(\mathbf{x}_t|\mathbf{x}_0) - \log q(\mathbf{x}_{t-1}|\mathbf{x}_0)) \\ &= \log p(\mathbf{x}_T) + \sum_{t=1}^T \log p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) - \log q(\mathbf{x}_T|\mathbf{x}_0) - \sum_{t=2}^T \log q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \\ &= \log \frac{p(\mathbf{x}_T)}{q(\mathbf{x}_T|\mathbf{x}_0)} + \log p_\theta(\mathbf{x}_0|\mathbf{x}_1) + \sum_{t=2}^T \log \frac{p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t)}{q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)} \end{aligned} \quad (15)$$

Therefore,

$$\begin{aligned} \log p_\theta(\mathbf{x}_0) &\geq \mathbf{E}_q[\log \frac{p(\mathbf{x}_0, \dots, \mathbf{x}_T)}{q(\mathbf{x}_1, \dots, \mathbf{x}_T|\mathbf{x}_0)}] \\ &= -\mathcal{D}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T)) - \sum_{t=2}^T \mathcal{D}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p(\mathbf{x}_{t-1}|\mathbf{x}_t)) + \mathbf{E}_q[\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)] \end{aligned} \quad (16)$$

where $\mathcal{D}(p||q)$ is the KL-diveragence between p and q . Let's denote:

- $L_0 = -\mathbf{E}_q[\log p_\theta(\mathbf{x}_0|\mathbf{x}_1)]$
- $L_{t-1} = \mathcal{D}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)||p(\mathbf{x}_{t-1}|\mathbf{x}_t))$, $t = 2, \dots, T$
- $L_T = \mathcal{D}(q(\mathbf{x}_T|\mathbf{x}_0)||p(\mathbf{x}_T))$

It is also noted that L_T is not a function of θ . Then the loss function of θ becomes:

$$\mathcal{L}(\theta) = L_0 + \sum_{t=1}^{T-1} L_t \quad (17)$$

Recall that:

- $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \frac{1}{\alpha_t}(\mathbf{x}_t - \frac{\beta_t^2}{\beta_t} \bar{\boldsymbol{\epsilon}}_t), \frac{\beta_t \bar{\beta}_{t-1}}{\beta_t} \mathbf{I});$
- $p(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \boldsymbol{\Sigma}_\theta(\mathbf{x}_t, t)).$ For simplicity of derivation, we can reparameterize $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{1}{\alpha_t}(\mathbf{x}_t - \frac{\beta_t^2}{\beta_t} \boldsymbol{\epsilon}(\mathbf{x}_t, t)).$
- The KL distance between two Gaussian distributions has a closed form, i.e.,

$$\text{KL}(\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) || \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2)) = \frac{1}{2} \left[\log \frac{\det(\boldsymbol{\Sigma}_2)}{\det \boldsymbol{\Sigma}_1} + \text{tr}(\boldsymbol{\Sigma}_2^{-1} \boldsymbol{\Sigma}_1) + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \boldsymbol{\Sigma}_2^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) \right]$$

Therefore, we have

$$L_t(\theta) = \frac{\beta_t^4}{2\alpha_t^2 \bar{\beta}_t^2 \det(\boldsymbol{\Sigma}_\theta)} \cdot \mathbf{E}_{\mathbf{x}_0, \boldsymbol{\epsilon}_t} (\bar{\boldsymbol{\epsilon}}_t - \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t))^2 \quad (19)$$

$$= \frac{\beta_t^4}{2\alpha_t^2 \bar{\beta}_t^2 \det(\boldsymbol{\Sigma}_\theta)} \cdot \mathbf{E}_{\mathbf{x}_0, \boldsymbol{\epsilon}_t} (\bar{\boldsymbol{\epsilon}}_t - \boldsymbol{\epsilon}_\theta(\bar{\alpha}_t \mathbf{x}_0 + \bar{\beta}_t \bar{\boldsymbol{\epsilon}}_t, t))^2 \quad (20)$$

Note the above expectation is taken over \mathbf{x}_0 which is drawn from the data distribution and $\bar{\boldsymbol{\epsilon}}_t$ which is a zero-mean, unit-variance Gaussian distribution. Empirically, it is found that training the following simplified loss function yield better results:

$$\mathcal{L}_{\text{simple}} = \mathbf{E}_{t, \mathbf{x}_0, \bar{\boldsymbol{\epsilon}}_t} (\bar{\boldsymbol{\epsilon}}_t - \boldsymbol{\epsilon}_\theta(\bar{\alpha}_t \mathbf{x}_0 + \bar{\beta}_t \bar{\boldsymbol{\epsilon}}_t, t))^2 \quad (21)$$

Based on this, the following training and inference algorithm can be derived:

Algorithm 1 Training

- 1: **repeat**
 - 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
 - 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
 - 4: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 5: Take gradient descent step on $\nabla_\theta \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, t)\|^2$
 - 6: **until** converged
-

Algorithm 2 Sampling

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **do**
 - 3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 - 4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(\mathbf{x}_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t) \right) + \sigma_t \mathbf{z}$
 - 5: **end for**
 - 6: **return** \mathbf{x}_0
-

Interpretation of the Training and Sampling Process

In the previous section, we derive the training and sampling process in a mathematical rigorous way. On the other hand, it may not be easy to understand the algorithms. Here we provide a few approaches to interpret how the training and sampling algorithm is derived in an intuitive manner.

Denoising perspective

The first method approaches to the problem from the denoising perspective. By the definition of reverse process, $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$ is to recover the mean of \mathbf{x}_{t-1} , therefore a reasonable loss function to optimize is thus:

$$\mathcal{L}_{\text{denoise}} = \mathbf{E}_{t, \mathbf{x}_{t-1}, \mathbf{x}_t} \|\mathbf{x}_{t-1} - \boldsymbol{\mu}_\theta(\mathbf{x}_t, t)\|^2 \quad (22)$$

Since we don't know the distribution of \mathbf{x}_t or \mathbf{x}_{t-1} and their joint distribution, we cannot determine the above loss function. Instead, we know that (from the forward process):

$$\mathbf{x}_{t-1} = \frac{1}{\alpha_t} (\mathbf{x}_t - \beta_t \boldsymbol{\varepsilon}_t) \quad (23)$$

Accordingly, we re-parameterize $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$ as

$$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{1}{\alpha_t} (\mathbf{x}_t - \beta_t \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, t)) \quad (24)$$

Then the loss function becomes

$$\mathcal{L}_{\text{denoise}} = \mathbf{E}_{t, \mathbf{x}_t} \frac{\beta_t^2}{\alpha_t^2} \|\boldsymbol{\varepsilon}_t - \boldsymbol{\varepsilon}_\theta(\mathbf{x}_t, t)\|^2 \quad (25)$$

To sample \mathbf{x}_t , recall that (by the fast-forward process),

$$\mathbf{x}_t = \bar{\alpha}_t \mathbf{x}_0 + \bar{\beta}_t \bar{\boldsymbol{\varepsilon}}_t \quad (26)$$

$$= \alpha_t (\bar{\alpha}_{t-1} \mathbf{x}_0 + \bar{\beta}_{t-1} \bar{\boldsymbol{\varepsilon}}_{t-1}) + \beta_t \boldsymbol{\varepsilon}_t \quad (27)$$

Plugging Eq. (27) into the loss function Eq. (22), (note that we cannot plug Eq. (26) into Eq. (22), because $\bar{\boldsymbol{\varepsilon}}_t$ is a function of $\boldsymbol{\varepsilon}_t$, so they cannot be sampled independently), we got

$$\mathcal{L}_{\text{denoise}} = \mathbf{E}_{t, \boldsymbol{\varepsilon}_t, \bar{\boldsymbol{\varepsilon}}_{t-1}} \frac{\beta_t^2}{\alpha_t^2} \|\boldsymbol{\varepsilon}_t - \boldsymbol{\varepsilon}_\theta(\bar{\alpha}_t \mathbf{x}_0 + \alpha_t \bar{\beta}_{t-1} \bar{\boldsymbol{\varepsilon}}_{t-1} + \beta_t \boldsymbol{\varepsilon}_t, t)\|^2 \quad (28)$$

We further noting that:

$$\bar{\beta}_t \bar{\boldsymbol{\varepsilon}}_t = \alpha_t \bar{\beta}_{t-1} \bar{\boldsymbol{\varepsilon}}_{t-1} + \beta_t \boldsymbol{\varepsilon}_t \quad (29)$$

is independent of $\bar{\beta}_t \mathbf{w} = \beta_t \bar{\boldsymbol{\varepsilon}}_{t-1} - \alpha_t \bar{\beta}_{t-1} \boldsymbol{\varepsilon}_t$ and

$$\boldsymbol{\varepsilon}_t = \frac{\beta_t \bar{\boldsymbol{\varepsilon}}_t - \alpha_t \bar{\beta}_{t-1} \mathbf{w}}{\bar{\beta}_t} \quad (30)$$

With the above changing of variable, the loss function can then be written as:

$$\mathcal{L}_{\text{denoise}} = \mathbf{E}_{t, \bar{\boldsymbol{\varepsilon}}_t} \frac{\beta_t^2}{\alpha_t^2} \left\| \frac{\beta_t}{\bar{\beta}_t} \bar{\boldsymbol{\varepsilon}}_t - \boldsymbol{\varepsilon}_\theta(\bar{\alpha}_t \mathbf{x}_0 + \bar{\beta}_t \bar{\boldsymbol{\varepsilon}}_t, t) \right\|^2 + C \quad (31)$$

where C is a constant (only related to \mathbf{w}). Also note that after the changing of variables, $\boldsymbol{\varepsilon}_\theta(\cdot, \cdot)$ in Eq (28) and (31) are different functions (one is trying to denoise one step noise $\boldsymbol{\varepsilon}_t$, and the other is trying to denoise cumulative noise $\bar{\boldsymbol{\varepsilon}}_t$)