
Supplementary of "Semi-Implicit Denoising Diffusion Models (SIDDMs)"

Anonymous Author(s)
 Affiliation
 Address
 email

1 S I. Derivation of Training Objective

2 Before getting the final training objective, we formulate the forward posterior following [1, 2]. Via
 3 Bayes' rule, we can rewrite the forward posterior given \mathbf{x}_0 :

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

4 where all the components of the very right equation are forward diffusion and follow Gaussian
 5 distribution. Thus the forward posterior can be rewritten as Gaussian with mean and standard
 6 deviation as follows:

$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0), \tilde{\boldsymbol{\beta}}_t \mathbf{I}).$$

7 Here we do not give redundant derivation and give the form of the forward posterior given x_t, x_0 .

$$\tilde{\boldsymbol{\mu}}_t(\mathbf{x}_t, \mathbf{x}_0) := \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}\mathbf{x}_0 + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}\mathbf{x}_t \quad \text{and} \quad \tilde{\boldsymbol{\beta}}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t}\beta_t.$$

8 And we parameterize our denoised x'_{t-1} given the predicted x'_0 and the input x_t via simply replacing
 9 the x_0 with the predicted x'_0 in the above equation. Then, given the parameterized x'_{t-1} .

10 To get our final training objective, We can rewrite the distribution matching objective of Equation (7)
 11 as:

$$\begin{aligned} & \min_{\theta} \max_{D_{adv}} \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} \left[D_{adv}(q(x_{t-1})||p_{\theta}(x_{t-1})) \right. \\ & \quad \left. + \lambda_{AFD}[-H(p_{\theta}(x_t|x_{t-1})) + H(p_{\theta}(x_t|x_{t-1}), q(x_t|x_{t-1}))] \right] \\ & = \min_{\theta} \max_{D_{adv}, \psi} \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} \left[D_{adv}(q(x_{t-1})||p_{\theta}(x_{t-1})) \right. \\ & \quad \left. + \lambda_{AFD}[H(p_{\theta}(x_t|x_{t-1}), q(x_t|x_{t-1})) - H(p_{\theta}(x_t|x_{t-1}), p_{\psi}(x_t|x_{t-1}))] \right], \end{aligned}$$

12 where the first GAN matching objective can be written as:

$$\min_{\theta} \max_{D_{\phi}} \sum_{t>0} \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} [-\log(D_{\phi}(x_{t-1}, t))] + [-\log(1 - D_{\phi}(x'_{t-1}, t))].$$

13 In the first cross-entropy of our distribution matching objective, the $q(x_t|x_{t-1})$ is the forward diffusion
 14 with the mean $\sqrt{1 - \beta_t}x_{t-1}$ and variance $\beta_t \mathbf{I}$. Thus the likelihood can be written as:

$$H(p_{\theta}(x_t|x_{t-1}), q(x_t|x_{t-1})) = \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} \frac{(1 - \beta_t) \|x'_{t-1} - x_{t-1}\|^2}{\beta_t},$$

15 To solve the second cross-entropy between the denoised distribution and the parameterized regression
 16 model, we define $p_\psi(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I})$ as forward diffusion for the regression
 17 model. And we also define x'_t are sampled from the x'_{t-1} via the forward diffusion. Similar to the
 18 above likelihood of cross-entropy, we can write the following likelihood for the second cross-entropy
 19 as follows:

$$H(p_\theta(x_t|x_{t-1}), p_\psi(x_t|x_{t-1})) = \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} \frac{\|C_\psi(x'_{t-1}) - x'_t\|^2}{\beta_t},$$

20 Finally, we can get the final training objective of our proposed method.

$$\min_{\theta} \max_{D_\phi, C_\psi} \sum_{t>0} \mathbb{E}_{q(x_0)q(x_{t-1}|x_0)q(x_t|x_{t-1})} \left[[-\log(D_\phi(x_{t-1}, t))] + [-\log(1 - D_\phi(x'_{t-1}, t))] \right. \\ \left. + \lambda_{AFD} \frac{(1 - \beta_t) \|x'_{t-1} - x_{t-1}\|^2 - \|C_\psi(x'_{t-1}) - x'_t\|^2}{\beta_t} \right],$$

21 In the main paper formulation, we mistakenly exchange the position of the β_t and $1 - \beta_t$, it is a typo,
 22 we will correct it later.

23 S II. Derivation of Theorem 1

24 For simplicity, we denote q via Q , p_θ via P and x_{t-1}, x_t via X, Y . According to the triangle
 25 inequality of total variation (TV) distance, we have

$$d_{TV}(Q_{XY}, P_{XY}) \leq d_{TV}(Q_{XY}, Q_{Y|X}P_X) + d_{TV}(Q_{Y|X}P_X, P_{XY}). \quad (\text{E11})$$

26 Using the definition of TV distance, we have

$$d_{TV}(Q_{Y|X}Q_X, Q_{Y|X}P_X) = \frac{1}{2} \int |Q_{Y|X}(y|x)Q_X(x) - Q_{Y|X}(y|x)P_X(x)|\mu(x, y) \\ \stackrel{(a)}{\leq} \frac{1}{2} \int |Q_{Y|X}(y|x)|\mu(x, y) \int |Q_X(x) - P_X(x)|\mu(x) \\ \leq c_1 d_{TV}(Q_X, P_X), \quad (\text{E12})$$

27 where P and Q are densities, μ is a (σ -finite) measure, c_1 is an upper bound of
 28 $\frac{1}{2} \int |Q_{Y|X}(y|x)|\mu(x, y)$, and (a) follows from the Hölder inequality.

29 Similarly, we have

$$d_{TV}(Q_{Y|X}P_X, P_{Y|X}P_X) \leq c_2 d_{TV}(Q_{Y|X}, P_{Y|X}), \quad (\text{E13})$$

30 where c_2 is an upper bound of $\frac{1}{2} \int |P_X(x)|\mu(x)$. Combining (E11), (E12), and (E13), we have

$$d_{TV}(Q_{XY}, P_{XY}) \leq c_1 d_{TV}(Q_X, P_X) + c_2 d_{TV}(Q_{Y|X}, P_{Y|X}) \quad (\text{E14})$$

31 According to the Pinsker inequality $d_{TV}(P, Q) \leq \sqrt{\frac{KL(P||Q)}{2}}$ [3], and the relation between TV and
 32 JSD, i.e., $\frac{1}{2}d_{TV}(P, Q)^2 \leq JSD(P, Q) \leq 2d_{TV}(P, Q)$ [4], we can rewrite (E14) as

$$JSD(Q_{XY}, P_{XY}) \leq 2c_1 \sqrt{2JSD(Q_X, P_X)} + 2c_2 \sqrt{2KL(P_{Y|X}||Q_{Y|X})}. \quad (\text{E15})$$

33 S III. Societal impact

34 With the increasing utilization of generative models, our proposed SSIDMs will improve the diffusion-
 35 based generative model while maintaining the highest level of generative quality. The incorporation
 36 of SSIDMs enhances the capabilities of generative models, particularly in the domain of text-to-image
 37 generation and editing. By integrating SSIDMs into the existing generative model framework, we
 38 could unlock new possibilities for generating realistic and visually coherent images from textual
 39 descriptions. One of the key advantages of our SSIDMs is their ability to accelerate the inference
 40 process, even though our model takes more time and more resources to train because of the additional
 41 adversarial training objectives. With faster inference, we eliminate the time-consuming barriers
 42 previously associated with text-to-image generation. As a result, real-time applications of generative
 43 models become feasible, enabling on-the-fly image generation or instant editing.

44 **S IV. More Implementation Details**

45 For the time steps, we apply the continuous time setup with the cosine noise schedule for all the
 46 experiments. We also apply a similar network structure as [5] and the downsampling trick as [6],
 47 where we put the downsampling layer at the beginning of each ResBlock. As mentioned, we design
 48 the discriminator as UNet, which adopts the symmetric network structure as the generator. For
 49 the regression model C_ψ and the discriminator regularizer, we share most of the layers with the
 50 discriminator except that we put a different linear head for the marginal, conditional and regularizer
 51 outputs. To be notified, the C_ψ only works on the denoised data, and the regularizer only works on
 52 the sampled x_{t-1} from the real data via forward diffusion. By this design, our model does not bring
 53 obvious extra overhead than our baseline DDGANs, which only has two more linear head in the final
 54 output for the discriminator network. We also describe the detailed model hyperparameters in the
 55 following table. We train all the models until they converges to the best FID score.

	CIFAR10 32	CelebA-HQ 256	ImageNet 64
Resolution	32	256	64
Conditional on labels	False	False	True
Diffusion steps	4	2	4
Noise Schedule	cosine	cosine	cosine
Channels	256	192	256
Depth	2	2	2
Channels multiple	2,2,2	1,1,2,3,4	1,2,3,4
Heads	4	4	4
Heads Channels	64	64	64
Attention resolution	16,8	32,16,8	32,16,8
Dropout	0.1	0.1	0.1
Batch size	256	128	2048
Learning Rate of G	2e-4	2e-4	2e-4
Learning Rate of D	1e-4	1e-4	1e-4
EMA Rate of G	0.9999	0.9999	0.9999

Table 1: Hyperparameters for our SSIDMs on different datasets.

56 **S V. More Generated Results**

57 **References**

- 58 [1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In
 59 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural*
 60 *Information Processing Systems*, volume 33, pages 6840–6851. Curran Associates, Inc., 2020.
- 61 [2] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with
 62 denoising diffusion GANs. In *International Conference on Learning Representations, 2022*.
- 63 [3] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Com-
 64 pany, Incorporated, 1st edition, 2008.
- 65 [4] Kiran K Thekumparampil, Ashish Khetan, Zinan Lin, and Sewoong Oh. Robustness of condi-
 66 tional gans to noisy labels. In *NeurIPS*, pages 10271–10282. Curran Associates, Inc., 2018.
- 67 [5] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis.
 68 In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural*
 69 *Information Processing Systems, 2021*.
- 70 [6] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed
 71 Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans,
 72 Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion
 73 models with deep language understanding. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave,
 74 and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems, 2022*.



Figure 1: Randomly generated samples from our model. We randomly sample 768 images from the generated images from CIFAR10, which we used to produce our paper results



Figure 2: Randomly generated samples from our model. We randomly sample 192 images from the generated images from CelebHQ 256, which we used to produce our paper results



Figure 3: Randomly generated samples from our model. We randomly sample 768 images from the generated images from Imagenet 64, which we used to produce our paper results