How does Lipschitz Regularization Influence GAN Training? - Supplementary Material -

Yipeng Qin[1,3], Niloy Mitra[2], and Peter Wonka[3]

[1] Cardiff University, qiny16@cardiff.ac.uk
[2] UCL/Adobe Research, n.mitra@cs.ucl.ac.uk
[3] KAUST, pwonka@gmail.com

Abstract. As supplementary material to the main paper we provide the detailed network architectures (Sec. 1), additional quantitative results (Sec. 2) and additional qualitative results (Sec. 3).

1 Detailed Network Architectures

The detailed network architectures of the generators and the discriminators are shown in Table 1.

2 Additional Quantitative Results

Table 2 shows the FID scores against α on all the three datasets (MNIST, CIFAR10 and CelebA), which is complementary to Table 4 in the main paper. The discussion in the main paper also applies to the additional results in Table 2.

We further verified the influence of domain scaling on the CIFAR10 dataset with different Lipschitz regularizers, *i.e.* Spectral Normalization (SN) [3] and Stable Rank Normalization (SRN)+SN [4], using additional metrics, *i.e.* Inception Scores (IS) [1] and Neural Divergence (ND) [2]. Specifically, we set $k_{SN} = 1$ for SN, $c = 0.7$ for SRN and reported both IS and ND scores with different α. As Table 3 and Table 4 show,

- When α is small, all loss functions degenerate to linear ones like WGAN and have similar IS and ND scores.
- When α is large, different loss functions behave differently and have different IS, ND scores (mostly worse).

Thus, our conclusions still hold in these settings.

3 Additional Qualitative Results

In this section, we show the qualitative results (sample images) corresponding to the quantitative experiments in the main paper. The FID scores and line plots are shown together with the samples.

3.1 Samples of FID scores v.s. k_{SN} Experiment

This subsection corresponds to the FID scores v.s. k_{SN} experiment (Table 3 in the main paper).

With varying k_{SN} on the MNIST dataset,

- Figure 1 shows sample images of the **NS-GAN-SN**;
- Figure 2 shows sample images of the **LS-GAN-SN**;
- Figure 3 shows sample images of the **WGAN-SN**;
- Figure 4 shows sample images of the **COS-GAN-SN**;
- Figure 5 shows sample images of the **EXP-GAN-SN**.

With varying k_{SN} on the CIFAR10 dataset,

- Figure 6 shows sample images of the **NS-GAN-SN**;
- Figure 7 shows sample images of the **LS-GAN-SN**;
- Figure 8 shows sample images of the **WGAN-SN**;
- Figure 9 shows sample images of the **COS-GAN-SN**;
- Figure 10 shows sample images of the **EXP-GAN-SN**.

With varying k_{SN} on the CelebA dataset,

- Figure 11 shows sample images of the **NS-GAN-SN**;
- Figure 12 shows sample images of the **LS-GAN-SN**;
- Figure 13 shows sample images of the **WGAN-SN**;
- Figure 14 shows sample images of the **COS-GAN-SN**;
- Figure 15 shows sample images of the **EXP-GAN-SN**.

3.2 Samples of the FID scores v.s. α Experiment

This subsection corresponds to the FID scores v.s. α experiment (Table 2). The results show that instead of the restriction of the neural network gradients, the restriction of the loss function is the dominating factor of Lipschitz regularization.

Results for Table 2 (a) With varying α, $k_{SN} = 50.0$ and on MNIST dataset,

- Figure 16 shows sample images of the **NS-GAN-SN**;
- Figure 17 shows sample images of the **LS-GAN-SN**;
- Figure 18 shows sample images of the **LS-GAN$^{\#}$-SN**;
- Figure 19 shows sample images of the **EXP-GAN-SN**;
- Figure 20 shows sample images of the **COS-GAN-SN**.

With varying α, $k_{SN} = 50.0$ and on CIFAR10 dataset,

- Figure 21 shows sample images of the **NS-GAN-SN**;
- Figure 22 shows sample images of the **LS-GAN-SN**;
- Figure 23 shows sample images of the **LS-GAN$^{\#}$-SN**;
- Figure 24 shows sample images of the **EXP-GAN-SN**;

– Figure 25 shows sample images of the **COS-GAN-SN**.

With varying α, $k_{SN} = 50.0$ and on CelebA dataset,

– Figure 26 shows sample images of the **NS-GAN-SN**;
– Figure 27 shows sample images of the **LS-GAN-SN**;
– Figure 28 shows sample images of the **LS-GAN$^{\#}$-SN**;
– Figure 29 shows sample images of the **EXP-GAN-SN**;
– Figure 30 shows sample images of the **COS-GAN-SN**.

Results for Table 2 (b) With varying α, $k_{SN} = 1.0$ and on MNIST dataset,

– Figure 31 shows sample images of the **NS-GAN-SN**;
– Figure 32 shows sample images of the **LS-GAN-SN**;
– Figure 33 shows sample images of the **LS-GAN$^{\#}$-SN**;
– Figure 34 shows sample images of the **EXP-GAN-SN**;
– Figure 35 shows sample images of the **COS-GAN-SN**.

With varying α, $k_{SN} = 1.0$ and on CIFAR10 dataset,

– Figure 36 shows sample images of the **NS-GAN-SN**;
– Figure 37 shows sample images of the **LS-GAN-SN**;
– Figure 38 shows sample images of the **LS-GAN$^{\#}$-SN**;
– Figure 39 shows sample images of the **EXP-GAN-SN**;
– Figure 40 shows sample images of the **COS-GAN-SN**.

With varying α, $k_{SN} = 1.0$ and on CelebA dataset,

– Figure 41 shows sample images of the **NS-GAN-SN**;
– Figure 42 shows sample images of the **LS-GAN-SN**;
– Figure 43 shows sample images of the **LS-GAN$^{\#}$-SN**;
– Figure 44 shows sample images of the **EXP-GAN-SN**;
– Figure 45 shows sample images of the **COS-GAN-SN**.

3.3 Samples of the Degenerate Loss Function

Figure 46 shows the sample images of WGAN-SN and some extremely degenerate loss functions, which corresponds to Table 5 in the main paper. It can be observed that all loss functions have similar performance.

References

1. Barratt, S., Sharma, R.: A note on the inception score. arXiv preprint arXiv:1801.01973 (2018)
2. Gulrajani, I., Raffel, C., Metz, L.: Towards GAN benchmarks which require generalization. In: International Conference on Learning Representations (2019), https://openreview.net/forum?id=HkxKH2AcFm
3. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: International Conference on Learning Representations (2018), https://openreview.net/forum?id=B1QRgziT-
4. Sanyal, A., Torr, P.H., Dokania, P.K.: Stable rank normalization for improved generalization in neural networks and gans. In: International Conference on Learning Representations (2020), https://openreview.net/forum?id=H1enKkrFDB

Table 1: Network architectures. BN: batch normalization; SN: spectral normalization; GP: gradient penalty.

$z \in \mathbb{R}^{100} \sim \mathcal{N}(0,1)$
Reshape $\rightarrow 1 \times 1 \times 100$
4×4, stride=1, deconv. BN 512 ReLU
4×4, stride=2, deconv. BN 256 ReLU
4×4, stride=2, deconv. BN 128 ReLU
4×4, stride=2, deconv. BN 64 ReLU
4×4, stride=2, deconv. 3 Tanh

Generator

RGB image $x \in [-1,1]^{64\times64\times3}$
4×4, stride=2, conv 64 BN lReLU(0.2)
4×4, stride=2, conv 128 BN lReLU(0.2)
4×4, stride=2, conv 256 BN lReLU(0.2)
4×4, stride=2, conv 512 BN lReLU(0.2)
4×4, stride=1, conv 1

Discriminator

(a) Network architectures with gradient penalty (GP) regularizer (64×64 resolution).

$z \in \mathbb{R}^{100} \sim \mathcal{N}(0,1)$
Reshape $\rightarrow 1 \times 1 \times 100$
4×4, stride=1, deconv. BN 512 ReLU
4×4, stride=2, deconv. BN 256 ReLU
4×4, stride=2, deconv. BN 128 ReLU
4×4, stride=2, deconv. BN 64 ReLU
4×4, stride=2, deconv. 3 Tanh

Generator

RGB image $x \in [-1,1]^{64\times64\times3}$
4×4, stride=2, conv 64 SN lReLU(0.2)
4×4, stride=2, conv 128 SN lReLU(0.2)
4×4, stride=2, conv 256 SN lReLU(0.2)
4×4, stride=2, conv 512 SN lReLU(0.2)
4×4, stride=1, conv 1

Discriminator

(b) Network architectures with spectral normalization (SN) regularizer (64×64 resolution).

$z \in \mathbb{R}^{100} \sim \mathcal{N}(0,1)$
Reshape $\rightarrow 1 \times 1 \times 100$
4×4, stride=1, deconv. BN 256 ReLU
4×4, stride=2, deconv. BN 128 ReLU
4×4, stride=2, deconv. BN 64 ReLU
4×4, stride=2, deconv. 3 Tanh

Generator

Image $x \in [-1,1]^{32\times32\times c}$
4×4, stride=2, conv 64 SN lReLU(0.2)
4×4, stride=2, conv 128 SN lReLU(0.2)
4×4, stride=2, conv 256 SN lReLU(0.2)
4×4, stride=1, conv 1

Discriminator

(c) Network architectures with spectral normalization (SN) regularizer (32×32 resolution). For input $x \in [-1,1]^{32\times32\times c}$, $c = 3$ for CIFAR10 dataset and $c = 1$ for MNIST dataset.

Table 2: FID scores v.s. α. For the line plots, the x-axis shows α (in log scale) and the y-axis shows the FID scores. Lower FID scores are better. The FID scores of the datasets in **bold** are not included in the main paper.

Dataset	GANs	FID Scores $\alpha = 1e^{-11}$	$1e^{-9}$	$1e^{-7}$	$1e^{-5}$	$1e^{-3}$	$1e^{-1}$	Line Plot
MNIST	NS-GAN-SN	3.67	3.61	3.81	33.48	154.14	155.54	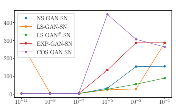
	LS-GAN-SN	281.82	6.48	3.74	24.93	29.31	285.46	
	LS-GAN$^{\#}$-SN	4.25	4.15	3.99	27.62	55.64	90.00	
	COS-GAN-SN	3.74	3.93	3.66	445.15	306.09	263.85	
	EXP-GAN-SN	4.14	4.01	3.54	134.76	286.96	286.96	
CIFAR10	NS-GAN-SN	19.58	22.46	18.73	24.57	49.56	43.42	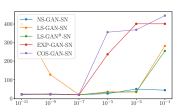
	LS-GAN-SN	400.91	127.38	18.68	34.78	33.17	282.11	
	LS-GAN$^{\#}$-SN	20.16	19.96	18.13	31.03	35.06	254.88	
	COS-GAN-SN	22.16	22.69	20.19	356.10	369.11	445.37	
	EXP-GAN-SN	21.93	21.70	18.50	236.77	401.24	401.24	
CelebA	NS-GAN-SN	9.08	7.05	7.84	18.51	18.41	242.64	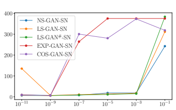
	LS-GAN-SN	135.17	6.57	10.67	13.39	17.42	311.93	
	LS-GAN$^{\#}$-SN	6.66	5.68	8.72	11.13	14.90	383.61	
	COS-GAN-SN	8.00	6.31	300.55	280.84	373.31	318.53	
	EXP-GAN-SN	8.85	6.09	264.49	375.32	375.32	375.32	

(a) $k_{SN} = 50.0$

Dataset	GANs	FID Scores $\alpha = 1e^{1}$	$1e^{3}$	$1e^{5}$	$1e^{7}$	$1e^{9}$	$1e^{11}$	Line Plot
MNIST	NS-GAN-SN	6.55	148.97	134.44	133.82	130.21	131.87	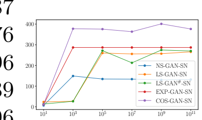
	LS-GAN-SN	23.37	26.96	260.05	255.73	256.96	265.76	
	LS-GAN$^{\#}$-SN	13.43	26.51	271.85	212.74	274.63	269.96	
	COS-GAN-SN	11.79	377.62	375.72	363.45	401.12	376.39	
	EXP-GAN-SN	11.02	286.96	286.96	286.96	286.96	286.96	
CIFAR10	NS-GAN-SN	17.63	47.31	46.85	45.44	45.67	39.90	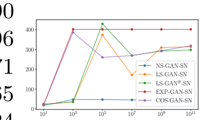
	LS-GAN-SN	25.55	34.44	373.07	171.18	309.55	312.96	
	LS-GAN$^{\#}$-SN	20.45	36.18	429.21	269.63	291.55	297.71	
	COS-GAN-SN	18.59	386.24	259.83	268.89	293.29	318.65	
	EXP-GAN-SN	21.56	401.24	401.24	401.24	401.24	401.24	
CelebA	NS-GAN-SN	5.88	16.14	17.75	17.67	16.87	18.81	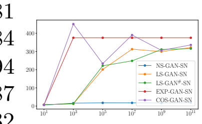
	LS-GAN-SN	8.41	12.09	201.22	312.83	299.30	321.84	
	LS-GAN$^{\#}$-SN	7.21	13.13	221.41	248.48	311.21	315.94	
	COS-GAN-SN	6.62	450.57	233.42	390.40	306.17	335.87	
	EXP-GAN-SN	6.91	375.32	375.32	375.32	375.32	375.32	

(b) $k_{SN} = 1.0$

Table 3: Inception Scores (IS) v.s. α on CIFAR10. Higher IS are better.

GANs	$\alpha = 1e^{-9}$	$1e^{-7}$	$1e^{-5}$	$1e^{-3}$	$1e^{-1}$	$1e^{0}$	$1e^{1}$	$1e^{3}$	$1e^{5}$	$1e^{7}$	$1e^{9}$
NS-GAN	4.32	4.04	4.03	4.00	4.47	4.48	4.64	3.60	3.23	3.31	3.49
LS-GAN	1.00	1.00	1.46	2.13	4.36	4.30	4.19	3.66	1.02	1.00	1.33
LS-GAN#	3.95	4.13	4.28	4.28	4.37	4.36	4.40	3.91	1.04	1.30	1.03
EXP-GAN	4.20	4.01	3.97	4.24	4.34	4.31	4.30	1.00	1.00	1.00	1.00
COS-GAN	4.20	4.10	4.24	4.13	4.32	4.35	4.40	2.00	1.00	1.05	1.35

(a) Lipschitz regularizer: SN

GANs	$\alpha = 1e^{-9}$	$1e^{-7}$	$1e^{-5}$	$1e^{-3}$	$1e^{-1}$	$1e^{0}$	$1e^{1}$	$1e^{3}$	$1e^{5}$	$1e^{7}$	$1e^{9}$
NS-GAN	3.80	4.03	4.00	4.29	4.37	4.29	4.48	3.36	3.51	3.55	3.52
LS-GAN	1.00	1.00	1.53	2.24	4.37	4.59	4.25	3.60	1.21	1.37	1.27
LS-GAN#	4.05	3.88	3.91	4.02	4.31	4.54	4.44	3.92	1.64	1.32	1.00
EXP-GAN	3.99	3.94	3.98	4.22	4.42	4.34	4.43	1.00	1.00	1.00	1.00
COS-GAN	3.98	4.21	3.98	3.71	4.25	4.60	4.43	1.04	1.61	1.29	1.66

(b) Lipschitz regularizer: SRN+SN

Table 4: Neural Divergence (ND) v.s. α on CIFAR10. Lower ND are better.

GANs	$\alpha = 1e^{-9}$	$1e^{-7}$	$1e^{-5}$	$1e^{-3}$	$1e^{-1}$	$1e^{0}$	$1e^{1}$	$1e^{3}$	$1e^{5}$	$1e^{7}$	$1e^{9}$
NS-GAN	44.76	36.04	39.01	37.31	36.39	32.98	24.47	52.44	46.27	52.33	51.26
LS-GAN	157.69	222.71	82.76	118.53	37.16	25.94	42.15	38.47	72.27	155.82	286.46
LS-GAN#	36.71	35.07	40.61	43.93	32.97	28.01	32.62	41.48	72.47	259.10	167.50
EXP-GAN	36.22	43.71	41.67	35.37	34.23	32.49	31.79	nan	nan	nan	nan
COS-GAN	40.53	35.66	36.33	39.48	36.39	31.73	33.21	477.07	61.40	69.79	81.15

(a) Lipschitz regularizer: SN

GANs	$\alpha = 1e^{-9}$	$1e^{-7}$	$1e^{-5}$	$1e^{-3}$	$1e^{-1}$	$1e^{0}$	$1e^{1}$	$1e^{3}$	$1e^{5}$	$1e^{7}$	$1e^{9}$
NS-GAN	39.28	37.12	31.60	37.82	33.56	23.33	26.57	34.30	49.68	45.96	52.72
LS-GAN	149.16	144.51	75.44	106.35	32.28	26.55	31.18	35.09	81.23	314.82	273.76
LS-GAN#	45.73	36.29	41.87	39.78	22.94	28.94	28.52	42.42	60.87	240.42	142.73
EXP-GAN	38.78	39.59	31.07	33.98	27.61	29.34	24.56	nan	nan	nan	nan
COS-GAN	35.85	40.54	31.28	41.60	30.50	39.72	25.62	104.24	107.67	263.35	158.75

(b) Lipschitz regularizer: SRN+SN

FID scores v.s. k_{SN} of Different Loss Functions

- MNIST -

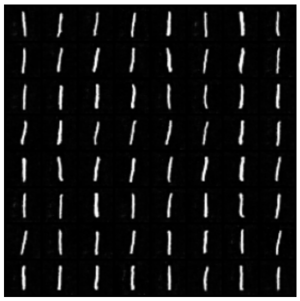

(a) k_{SN}=50.0, FID=155.41

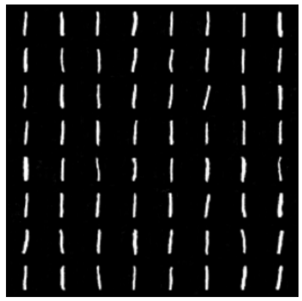

(b) k_{SN}=10.0, FID=156.60

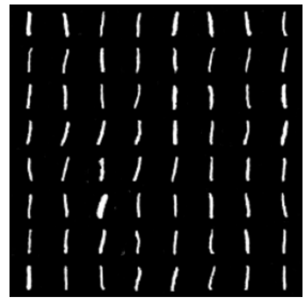

(c) k_{SN}=5.0, FID=144.28

(d) k_{SN}=1.0, FID=3.90

(e) k_{SN}=0.5, FID=4.20

(f) k_{SN}=0.25, FID=3.99

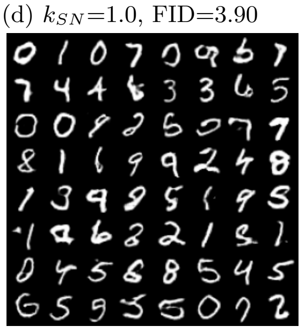

(g) k_{SN}=0.2, FID=5.41

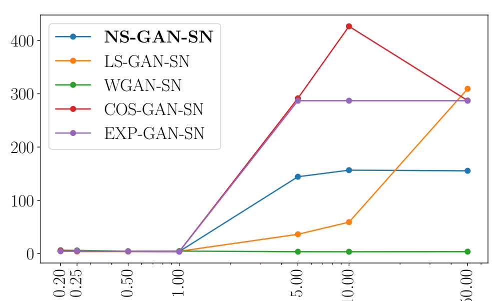

Fig. 1: Samples of randomly generated images with NS-GAN-SN of varying k_{SN} (MNIST). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

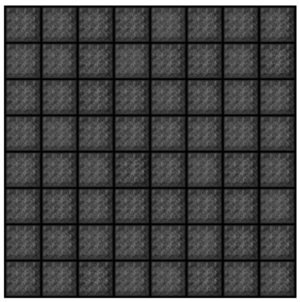

(a) k_{SN}=50.0, FID=309.35

(b) k_{SN}=10.0, FID=59.04

(c) k_{SN}=5.0, FID=36.26

(d) k_{SN}=1.0, FID=4.42

(e) k_{SN}=0.5, FID=3.90

(f) k_{SN}=0.25, FID=3.96

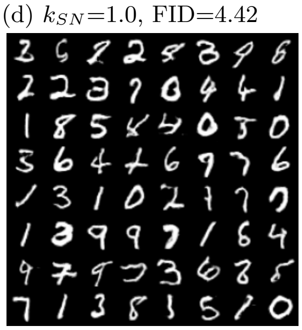

(g) k_{SN}=0.2, FID=5.14

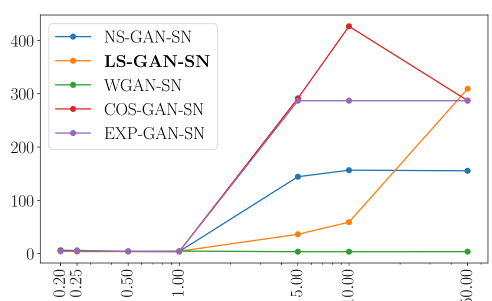

Fig. 2: Samples of randomly generated images with LS-GAN-SN of varying k_{SN} (MNIST). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

(a) $k_{SN} = 50.0$, FID = 3.71

(b) $k_{SN} = 10.0$, FID = 3.50

(c) $k_{SN} = 5.0$, FID = 3.58

(d) $k_{SN} = 1.0$, FID = 4.70

(e) $k_{SN} = 0.5$, FID = 4.44

(f) $k_{SN} = 0.25$, FID = 6.06

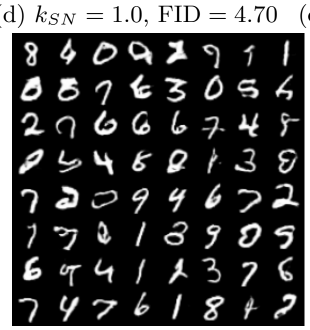

(g) $k_{SN} = 0.2$, FID = 6.35

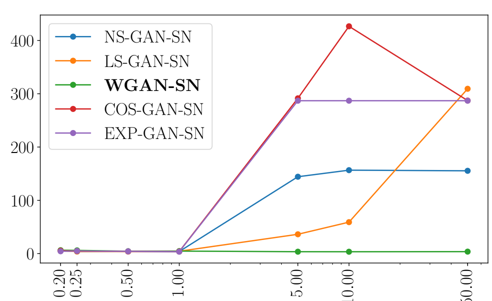

Fig. 3: Samples of randomly generated images with WGAN-SN of varying k_{SN} (MNIST). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

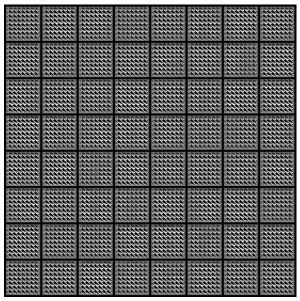

(a) k_{SN}=50.0, FID=287.23

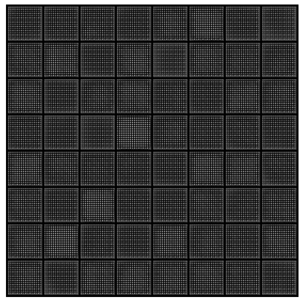

(b) k_{SN}=10.0, FID=426.62

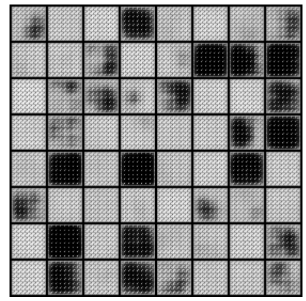

(c) k_{SN}=5.0, FID=291.44

(d) k_{SN}=1.0, FID=3.86

(e) k_{SN}=0.5, FID=4.05

(f) k_{SN}=0.25, FID=4.83

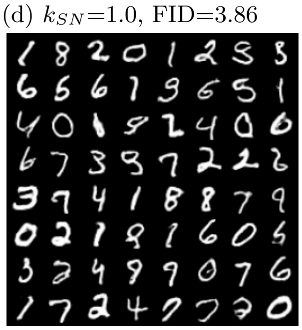

(g) k_{SN}=0.2, FID=5.41

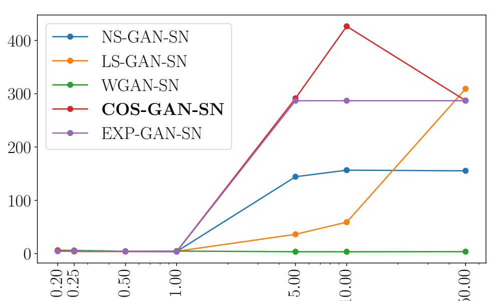

Fig. 4: Samples of randomly generated images with COS-GAN-SN of varying k_{SN} (MNIST). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

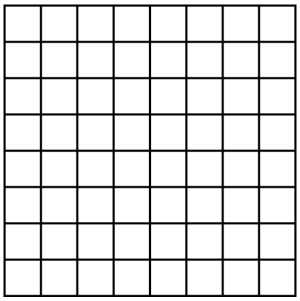
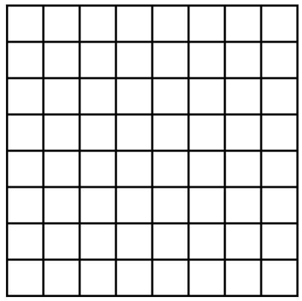
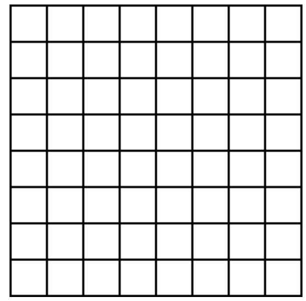

(a) k_{SN}=50.0, FID=286.96 (b) k_{SN}=10.0, FID=286.96 (c) k_{SN}=5.0, FID=286.96



(d) k_{SN}=1.0, FID=3.69 (e) k_{SN}=0.5, FID=4.25 (f) k_{SN}=0.25, FID=4.93

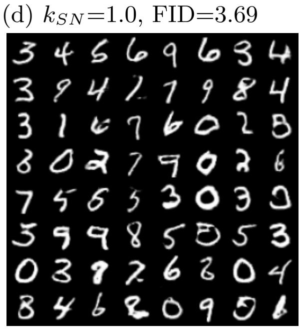

(g) k_{SN}=0.2, FID=4.38

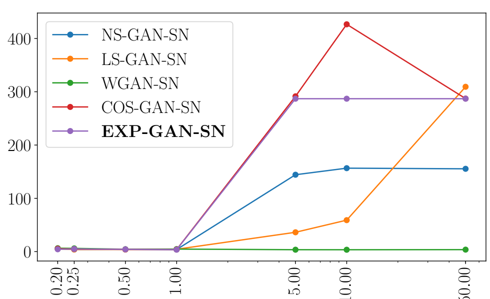

Fig. 5: Samples of randomly generated images with EXP-GAN-SN of varying k_{SN} (MNIST). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

FID scores v.s. k_{SN} of Different Loss Functions

- CIFAR10 -

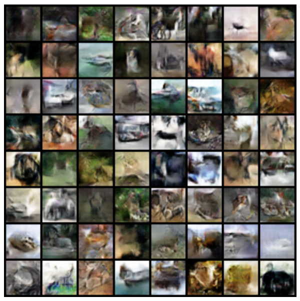

(a) k_{SN}=50.0, FID=48.03

(b) k_{SN}=10.0, FID=49.67

(c) k_{SN}=5.0, FID=41.04

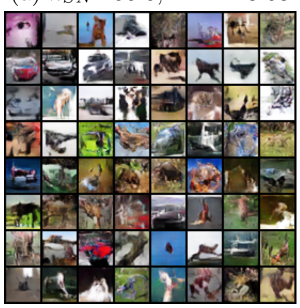

(d) k_{SN}=1.0, FID=15.81

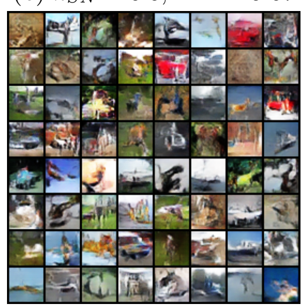

(e) k_{SN}=0.5, FID=23.29

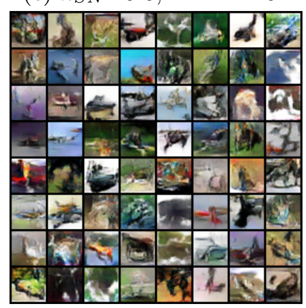

(f) k_{SN}=0.25, FID=24.37

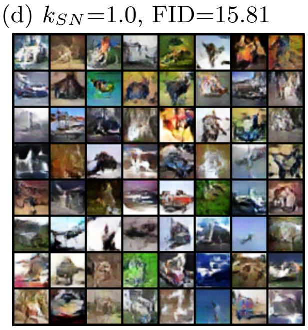

(g) k_{SN}=0.2, FID=29.23

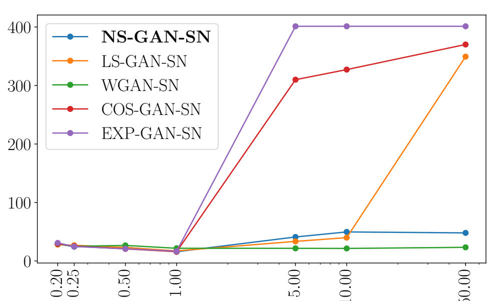

Fig. 6: Samples of randomly generated images with NS-GAN-SN of varying k_{SN} (CIFAR10). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

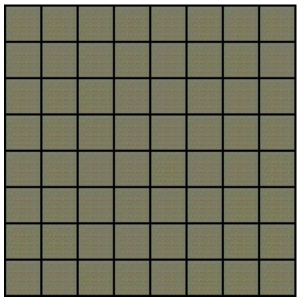

(a) k_{SN}=50.0, FID= 349.35

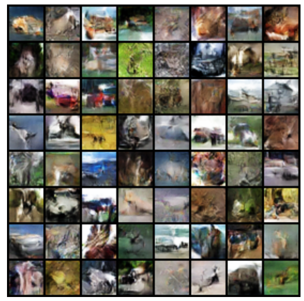

(b) k_{SN}=10.0, FID= 39.90

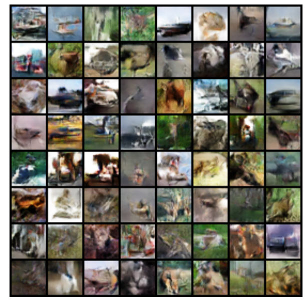

(c) k_{SN}=5.0, FID= 33.53

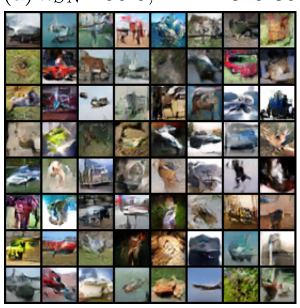

(d) k_{SN}=1.0, FID= 17.30

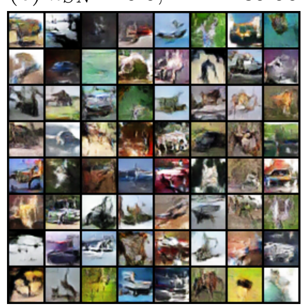

(e) k_{SN}=0.5, FID= 23.14

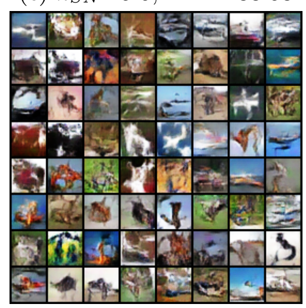

(f) k_{SN}=0.25, FID= 26.85

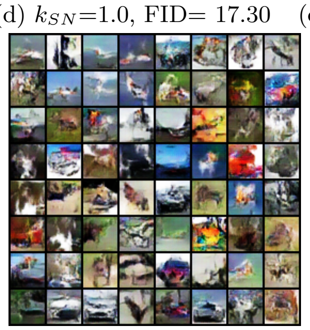

(g) k_{SN}=0.2, FID= 28.04

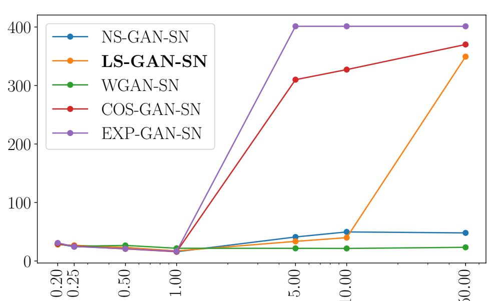

Fig. 7: Samples of randomly generated images with LS-GAN-SN of varying k_{SN} (CIFAR10). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

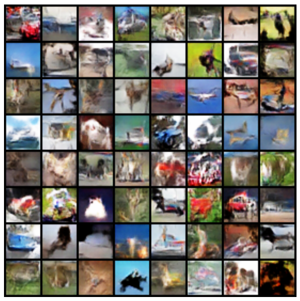

(a) k_{SN}=50.0, FID= 23.36

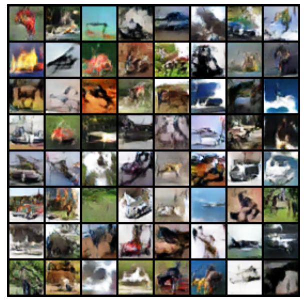

(b) k_{SN}=10.0, FID= 21.45

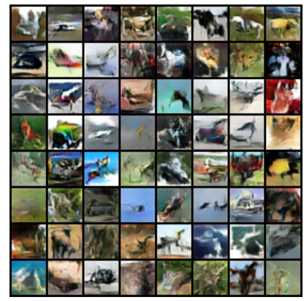

(c) k_{SN}=5.0, FID= 21.63

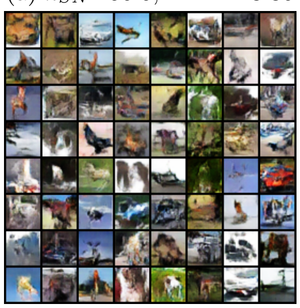

(d) k_{SN}=1.0, FID= 21.75

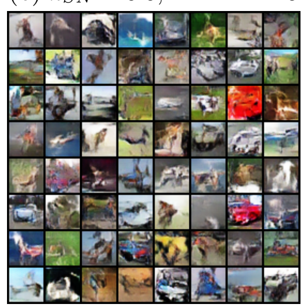

(e) k_{SN}=0.5, FID= 26.61

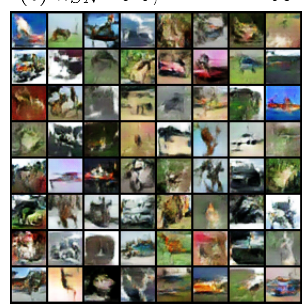

(f) k_{SN}=0.25, FID= 25.07

(g) k_{SN}=0.2, FID= 29.20

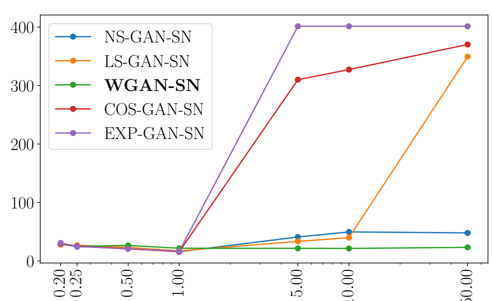

Fig. 8: Samples of randomly generated images with WGAN-SN of varying k_{SN} (CIFAR10). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

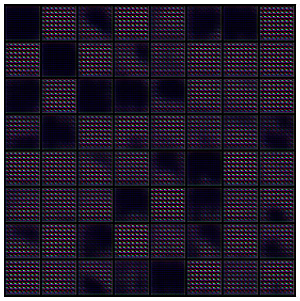
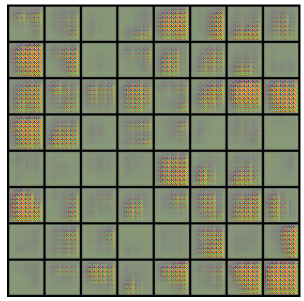
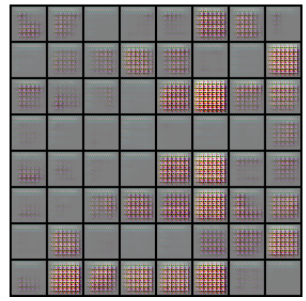

(a) k_{SN}=50.0, FID= 370.13 (b) k_{SN}=10.0, FID= 327.20 (c) k_{SN}=5.0, FID= 309.96

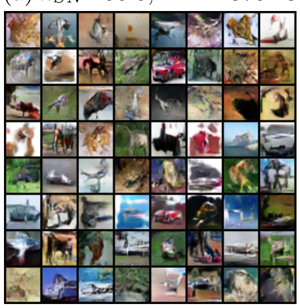
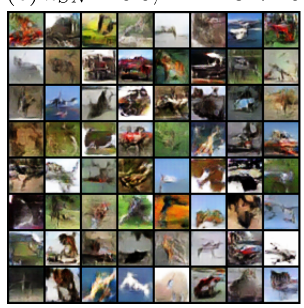
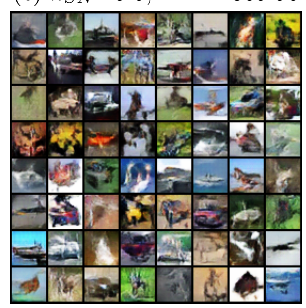

(d) k_{SN}=1.0, FID= 15.88 (e) k_{SN}=0.5, FID= 20.73 (f) k_{SN}=0.25, FID= 25.31

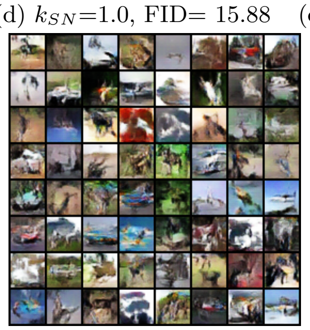

(g) k_{SN}=0.2, FID= 29.45

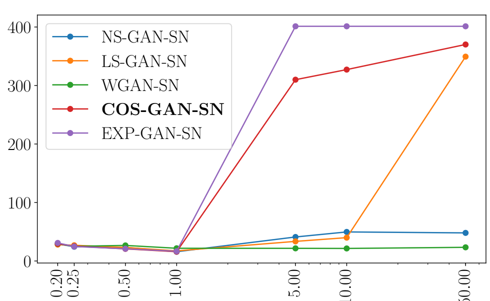

Fig. 9: Samples of randomly generated images with COS-GAN-SN of varying k_{SN} (CIFAR10). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

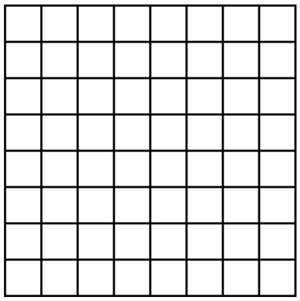
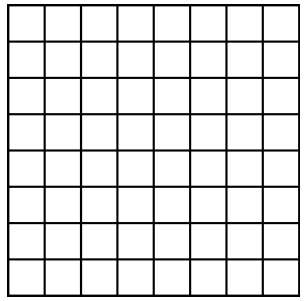
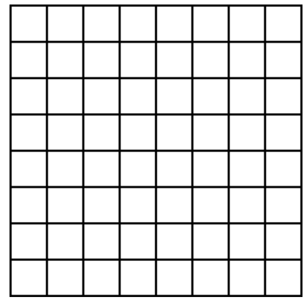

(a) k_{SN}=50.0, FID= 401.24 (b) k_{SN}=10.0, FID= 401.24 (c) k_{SN}=5.0, FID= 401.24

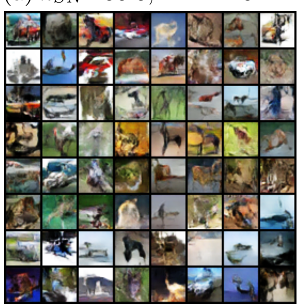
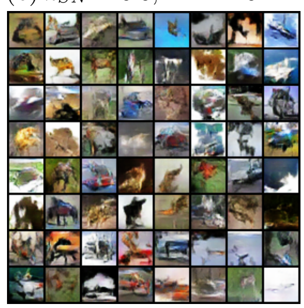
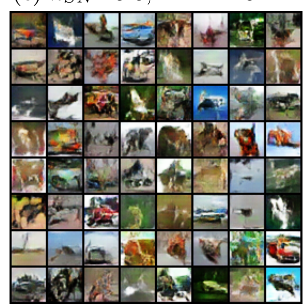

(d) k_{SN}=1.0, FID= 16.66 (e) k_{SN}=0.5, FID= 20.90 (f) k_{SN}=0.25, FID= 24.74

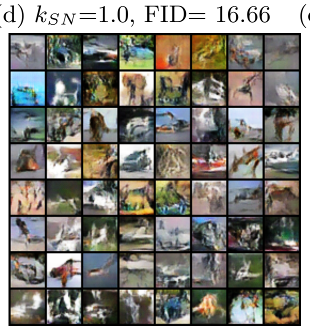

(g) k_{SN}=0.2, FID= 30.89

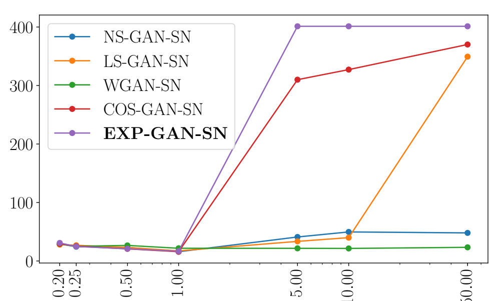

Fig. 10: Samples of randomly generated images with EXP-GAN-SN of varying k_{SN} (CIFAR10). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

FID scores v.s. k_{SN} of Different Loss Functions
- CelebA -

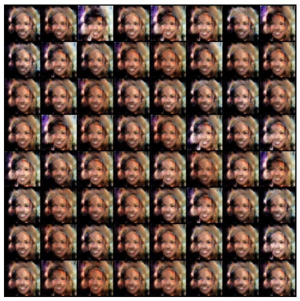

(a) k_{SN}=50.0, FID= 184.06

(b) k_{SN}=10.0, FID= 17.04

(c) k_{SN}=5.0, FID= 18.95

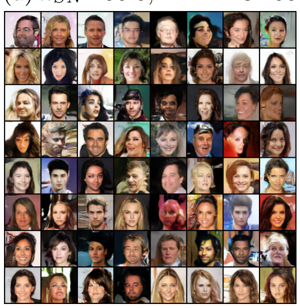

(d) k_{SN}=1.0, FID= 6.11

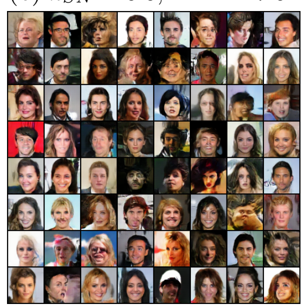

(e) k_{SN}=0.5, FID= 8.04

(f) k_{SN}=0.25, FID= 12.71

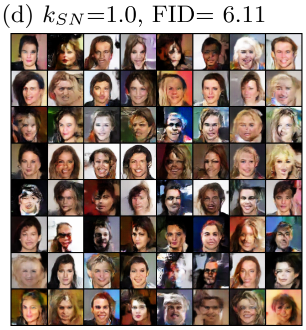

(g) k_{SN}=0.2, FID= 18.59

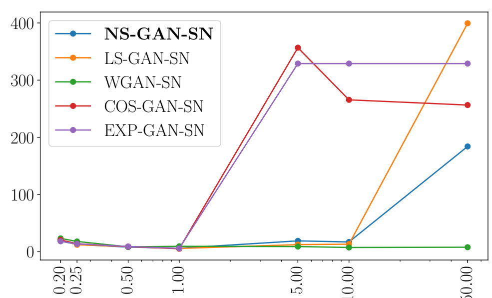

Fig. 11: Samples of randomly generated images with NS-GAN-SN of varying k_{SN} (CelebA). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

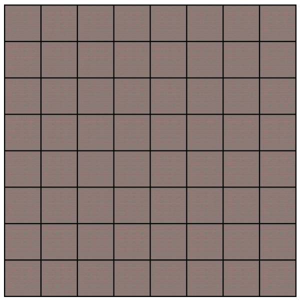

(a) k_{SN}=50.0, FID= 399.39

(b) k_{SN}=10.0, FID= 13.14

(c) k_{SN}=5.0, FID= 12.40

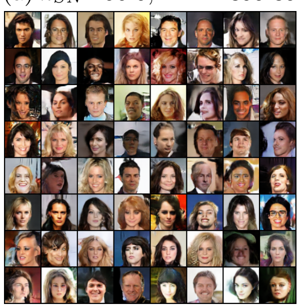

(d) k_{SN}=1.0, FID= 5.69

(e) k_{SN}=0.5, FID= 8.85

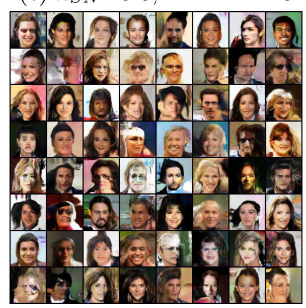

(f) k_{SN}=0.25, FID= 12.14

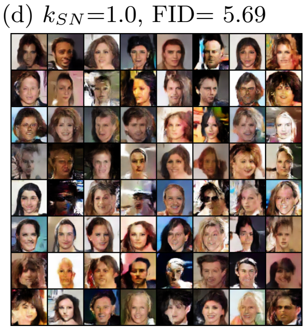

(g) k_{SN}=0.2, FID= 20.34

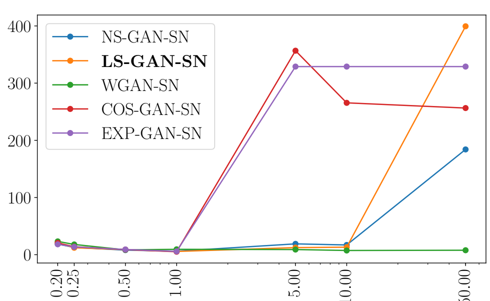

Fig. 12: Samples of randomly generated images with LS-GAN-SN of varying k_{SN} (CelebA). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

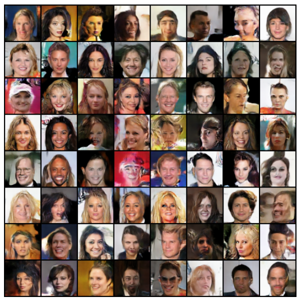

(a) k_{SN}=50.0, FID= 7.82

(b) k_{SN}=10.0, FID= 7.37

(c) k_{SN}=5.0, FID= 9.03

(d) k_{SN}=1.0, FID= 9.41

(e) k_{SN}=0.5, FID= 8.48

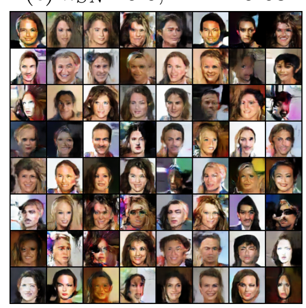

(f) k_{SN}=0.25, FID= 17.93

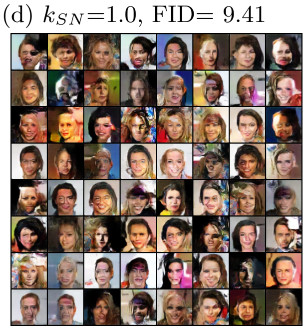

(g) k_{SN}=0.2, FID= 23.26

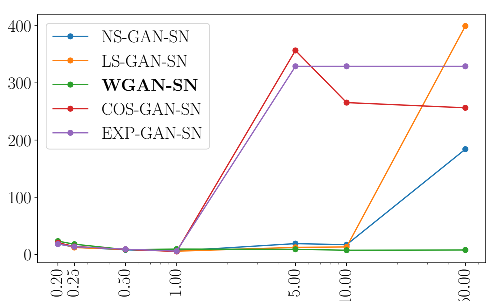

Fig. 13: Samples of randomly generated images with WGAN-SN of varying k_{SN} (CelebA). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

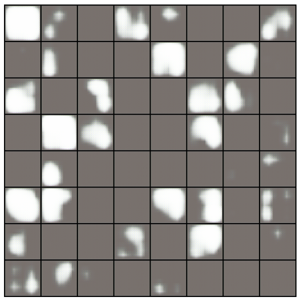

(a) k_{SN}=50.0, FID= 256.44

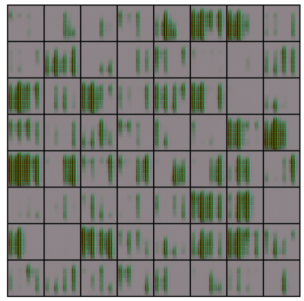

(b) k_{SN}=10.0, FID= 265.53

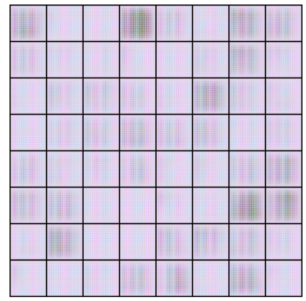

(c) k_{SN}=5.0, FID= 356.70

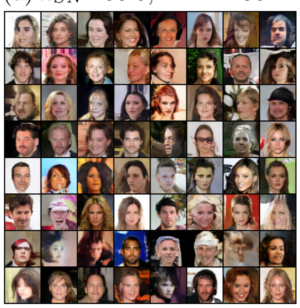

(d) k_{SN}=1.0, FID= 5.20

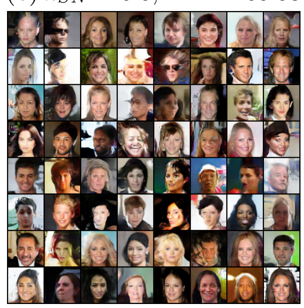

(e) k_{SN}=0.5, FID= 8.88

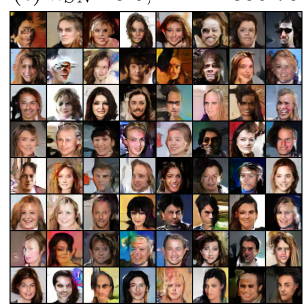

(f) k_{SN}=0.25, FID= 13.93

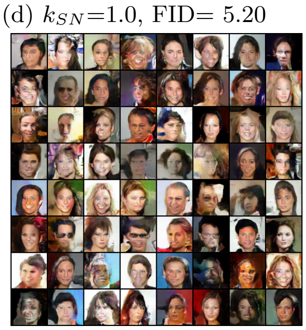

(g) k_{SN}=0.2, FID= 20.59

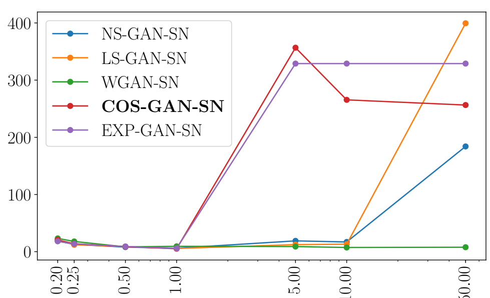

Fig. 14: Samples of randomly generated images with COS-GAN-SN of varying k_{SN} (CelebA). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

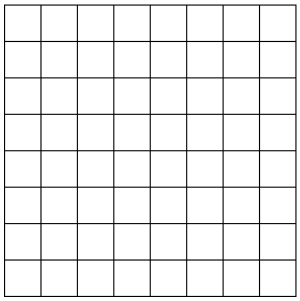

(a) k_{SN}=50.0, FID= 328.94

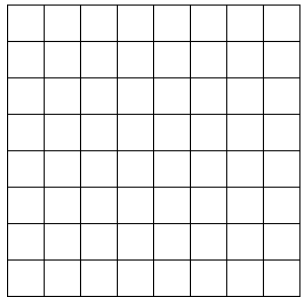

(b) k_{SN}=10.0, FID= 328.94

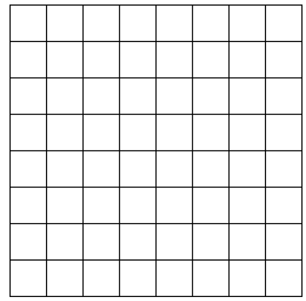

(c) k_{SN}=5.0, FID= 328.94

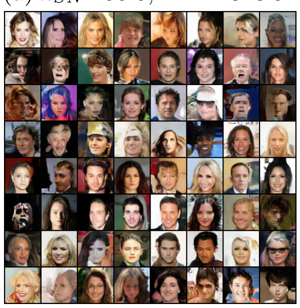

(d) k_{SN}=1.0, FID= 5.88

(e) k_{SN}=0.5, FID= 9.18

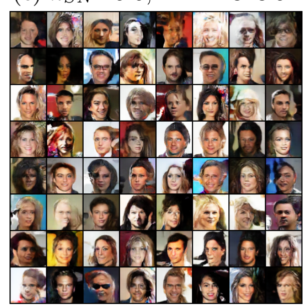

(f) k_{SN}=0.25, FID= 13.65

(g) k_{SN}=0.2, FID= 18.23

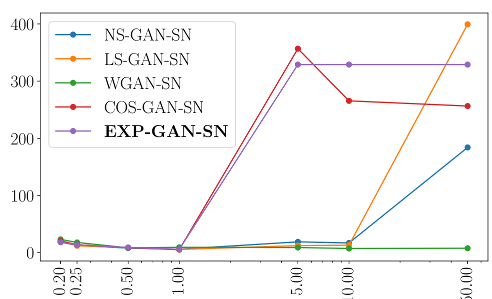

Fig. 15: Samples of randomly generated images with EXP-GAN-SN of varying k_{SN} (CelebA). For the line plot, x-axis shows k_{SN} (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 50.0$) of Different Loss Functions

- MNIST -

(a) $\alpha = 1e^{-11}$, FID = 3.67 (b) $\alpha = 1e^{-9}$, FID = 3.61 (c) $\alpha = 1e^{-7}$, FID = 3.81

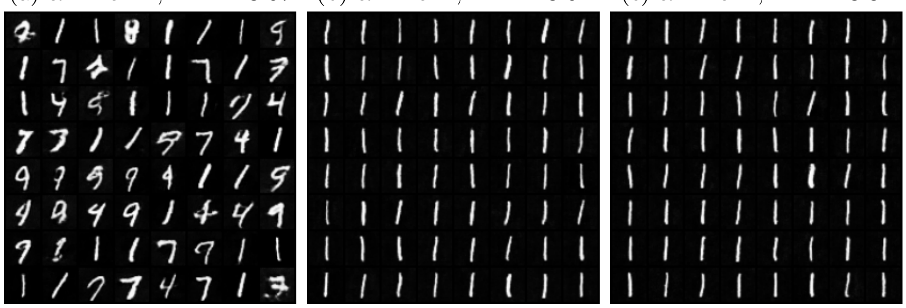

(d) $\alpha = 1e^{-5}$, FID = 33.48 (e) $\alpha = 1e^{-3}$, FID = 154.14 (f) $\alpha = 1e^{-1}$, FID = 155.54

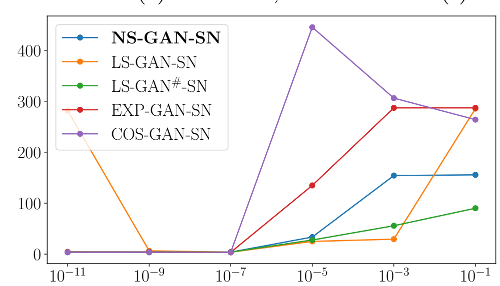

Fig. 16: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 50.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

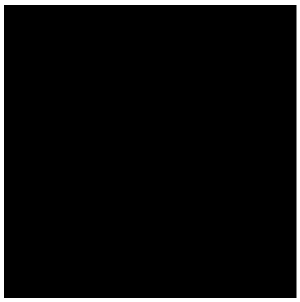

(a) $\alpha = 1e^{-11}$, FID=281.82 (b) $\alpha = 1e^{-9}$, FID = 6.48 (c) $\alpha = 1e^{-7}$, FID = 3.74



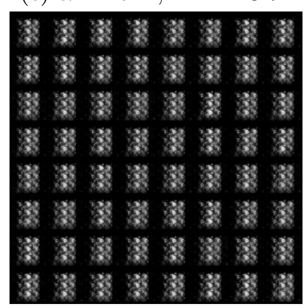

(d) $\alpha = 1e^{-5}$, FID = 24.93 (e) $\alpha = 1e^{-3}$, FID = 29.31 (f) $\alpha = 1e^{-1}$, FID = 285.46

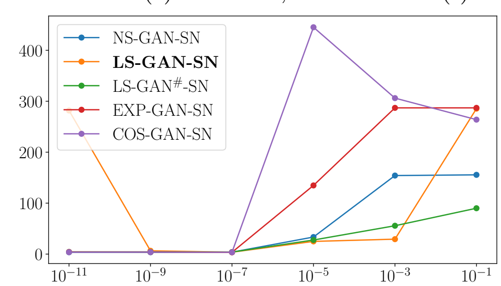

Fig. 17: Samples of randomly generated images with LS-GAN-SN of varying α ($k_{SN} = 50.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^{-11}$, FID = 4.25

(b) $\alpha = 1e^{-9}$, FID = 4.15

(c) $\alpha = 1e^{-7}$, FID = 3.99

(d) $\alpha = 1e^{-5}$, FID = 27.62

(e) $\alpha = 1e^{-3}$, FID = 55.64

(f) $\alpha = 1e^{-1}$, FID = 90.00

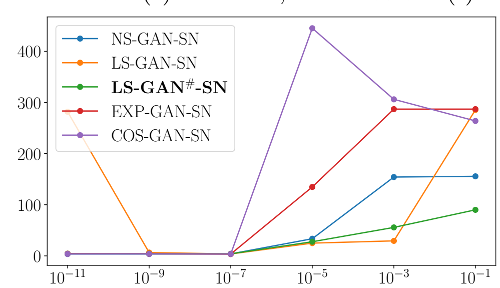

Fig. 18: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α ($k_{SN} = 50.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

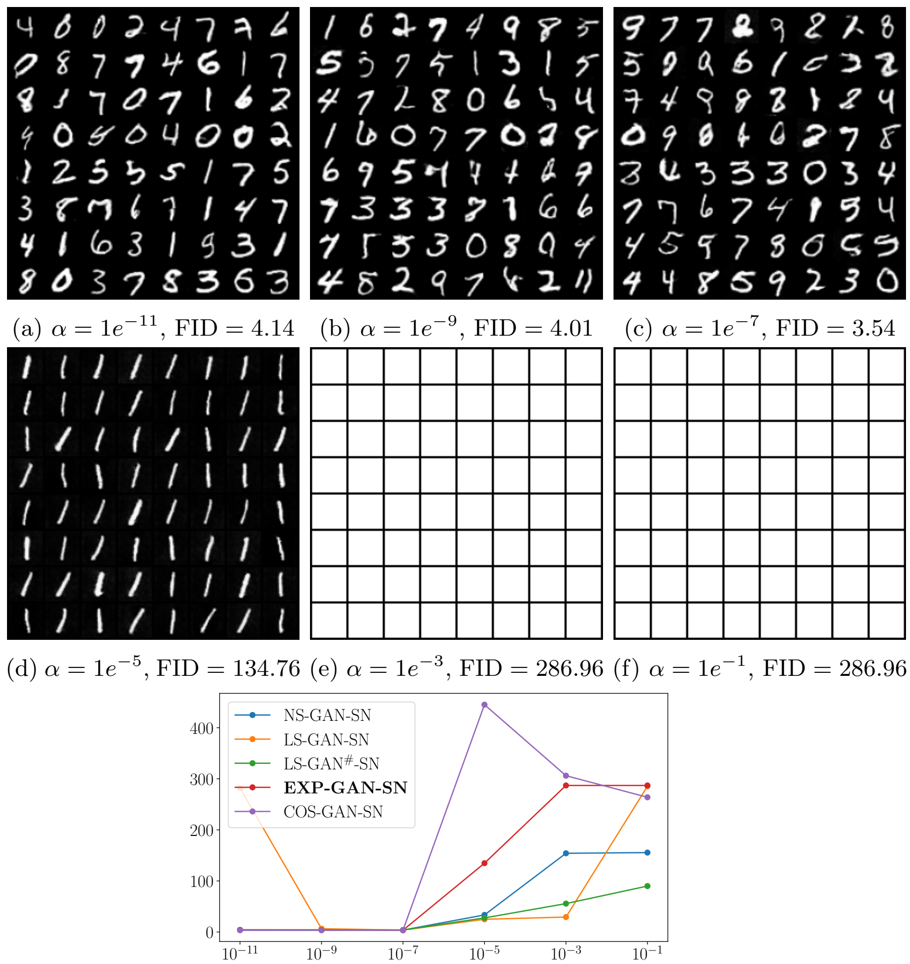

Fig. 19: Samples of randomly generated images with EXP-GAN-SN of varying α ($k_{SN} = 50.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^{-11}$, FID = 3.74 (b) $\alpha = 1e^{-9}$, FID = 3.93 (c) $\alpha = 1e^{-7}$, FID = 3.66

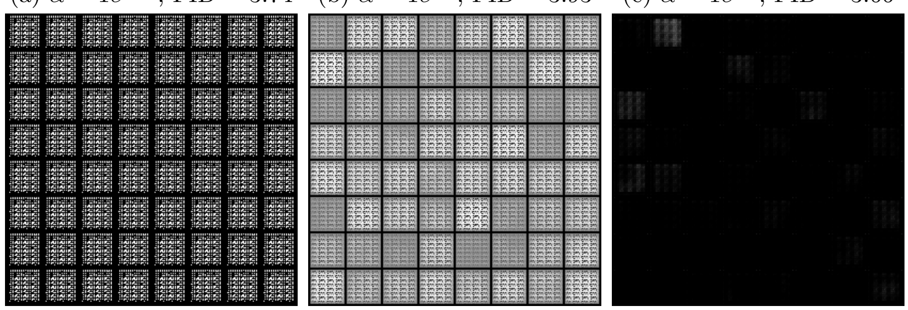

(d) $\alpha = 1e^{-5}$, FID = 445.15 (e) $\alpha = 1e^{-3}$, FID = 306.09 (f) $\alpha = 1e^{-1}$, FID = 263.85

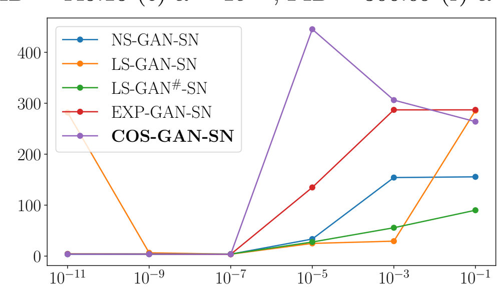

Fig. 20: Samples of randomly generated images with COS-GAN-SN of varying α ($k_{SN} = 50.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 50.0$) of Different Loss Functions

- CIFAR10 -

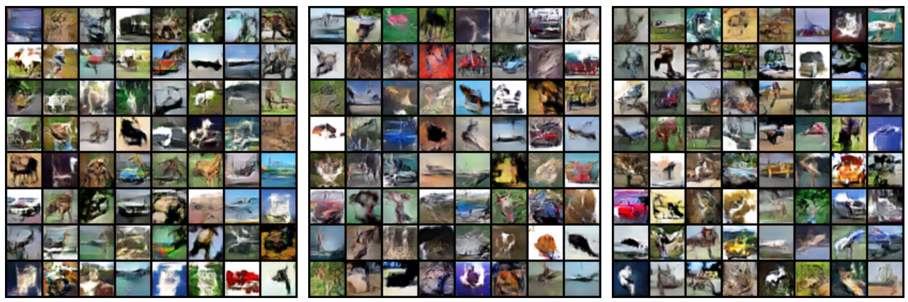

(a) $\alpha = 1e^{-11}$, FID = 19.58 (b) $\alpha = 1e^{-9}$, FID = 22.46 (c) $\alpha = 1e^{-7}$, FID = 18.73

(d) $\alpha = 1e^{-5}$, FID = 24.57 (e) $\alpha = 1e^{-3}$, FID = 49.56 (f) $\alpha = 1e^{-1}$, FID = 43.42

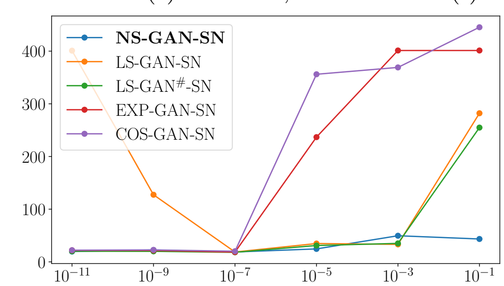

Fig. 21: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 50.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

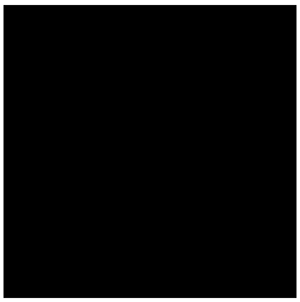
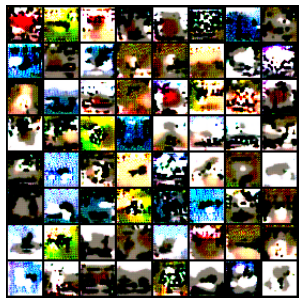
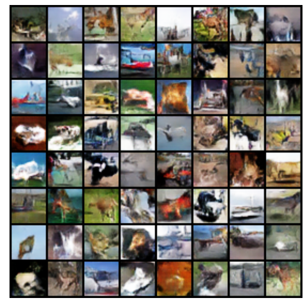

(a) $\alpha = 1e^{-11}$, FID=400.91 (b) $\alpha = 1e^{-9}$, FID = 127.38 (c) $\alpha = 1e^{-7}$, FID = 18.68

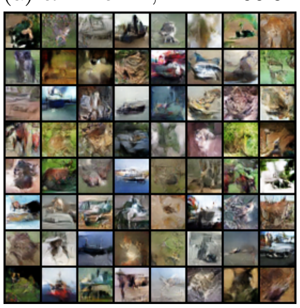
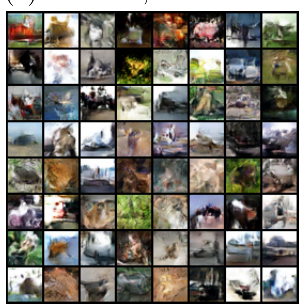
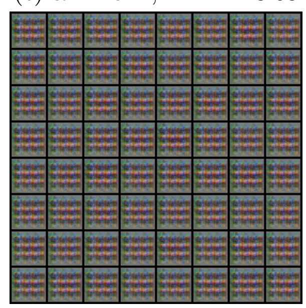

(d) $\alpha = 1e^{-5}$, FID = 34.78 (e) $\alpha = 1e^{-3}$, FID = 33.17 (f) $\alpha = 1e^{-1}$, FID = 282.11

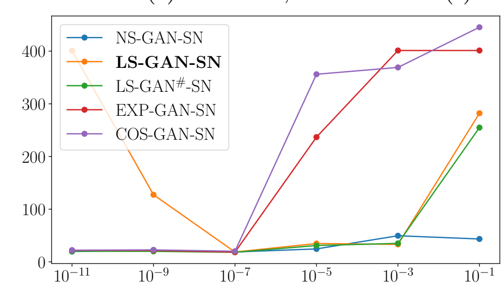

Fig. 22: Samples of randomly generated images with LS-GAN-SN of varying α (k_{SN} = 50.0, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

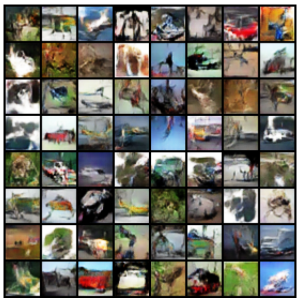

(a) $\alpha = 1e^{-11}$, FID = 20.16

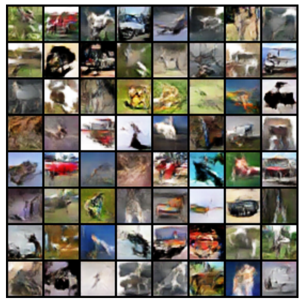

(b) $\alpha = 1e^{-9}$, FID = 19.96

(c) $\alpha = 1e^{-7}$, FID = 18.13

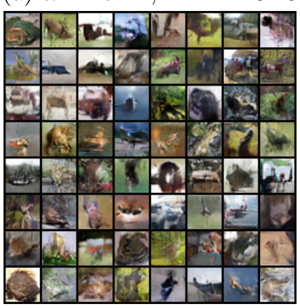

(d) $\alpha = 1e^{-5}$, FID = 31.03

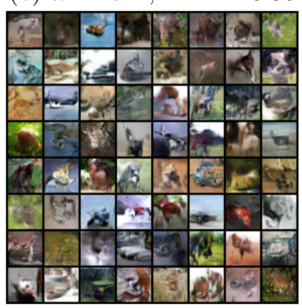

(e) $\alpha = 1e^{-3}$, FID = 35.06

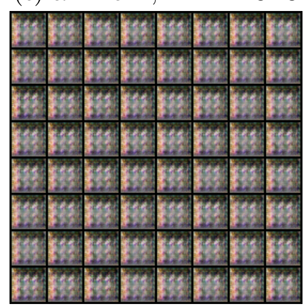

(f) $\alpha = 1e^{-1}$, FID = 254.88

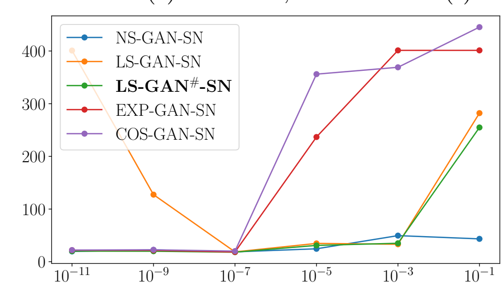

Fig. 23: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α ($k_{SN} = 50.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.


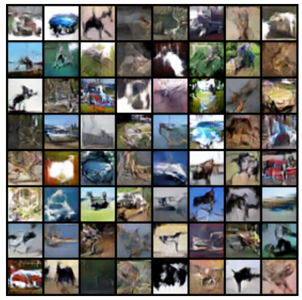
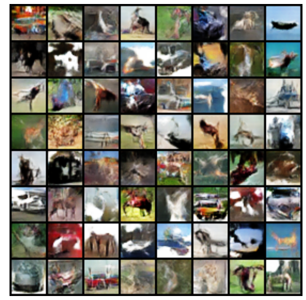

(a) $\alpha = 1e^{-11}$, FID = 21.93 (b) $\alpha = 1e^{-9}$, FID = 21.70 (c) $\alpha = 1e^{-7}$, FID = 18.50

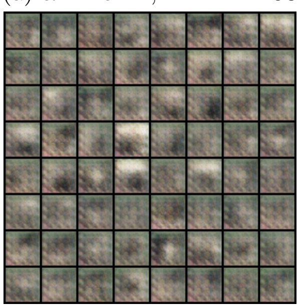
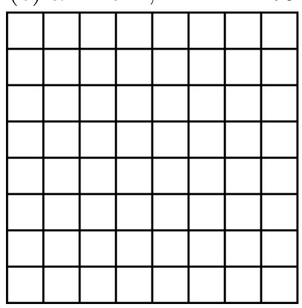
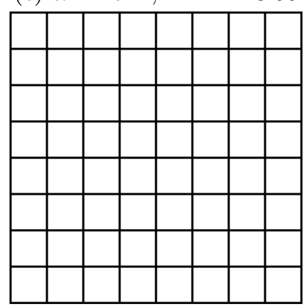

(d) $\alpha = 1e^{-5}$, FID = 236.77 (e) $\alpha = 1e^{-3}$, FID = 401.24 (f) $\alpha = 1e^{-1}$, FID = 401.24

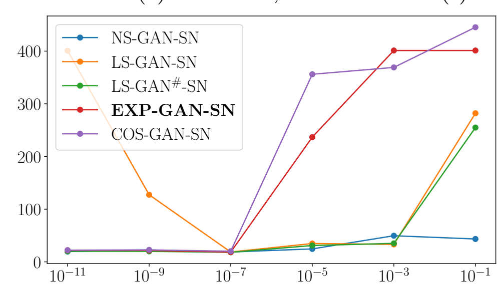

Fig. 24: Samples of randomly generated images with EXP-GAN-SN of varying α (k_{SN} = 50.0, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

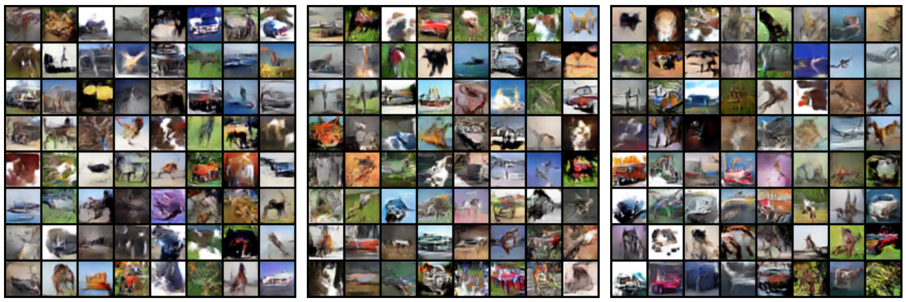

(a) $\alpha = 1e^{-11}$, FID = 22.16 (b) $\alpha = 1e^{-9}$, FID = 22.69 (c) $\alpha = 1e^{-7}$, FID = 20.19

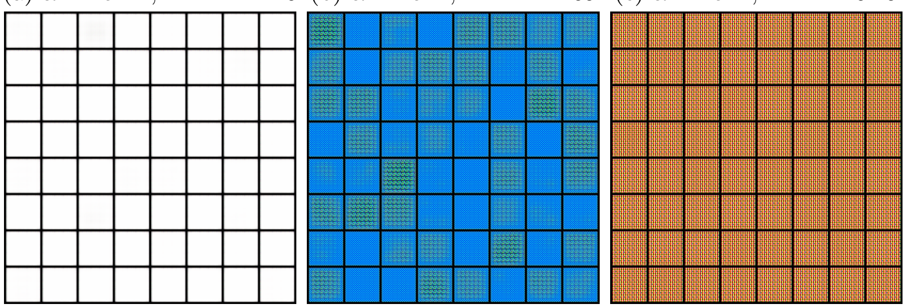

(d) $\alpha = 1e^{-5}$, FID = 356.10 (e) $\alpha = 1e^{-3}$, FID = 369.11 (f) $\alpha = 1e^{-1}$, FID = 445.37

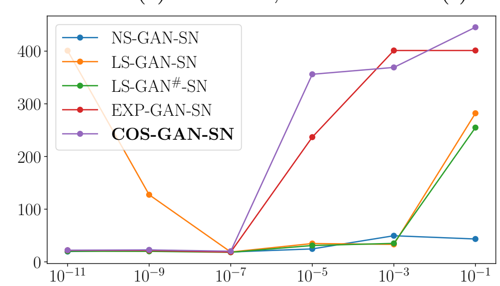

Fig. 25: Samples of randomly generated images with COS-GAN-SN of varying α ($k_{SN} = 50.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 50.0$) of Different Loss Functions

- CelebA -

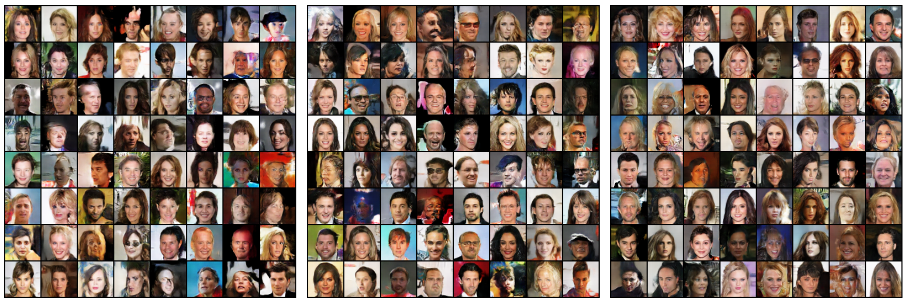

(a) $\alpha = 1e^{-11}$, FID = 9.08 (b) $\alpha = 1e^{-9}$, FID = 7.05 (c) $\alpha = 1e^{-7}$, FID = 7.84

(d) $\alpha = 1e^{-5}$, FID = 18.51 (e) $\alpha = 1e^{-3}$, FID = 18.41 (f) $\alpha = 1e^{-1}$, FID = 242.64

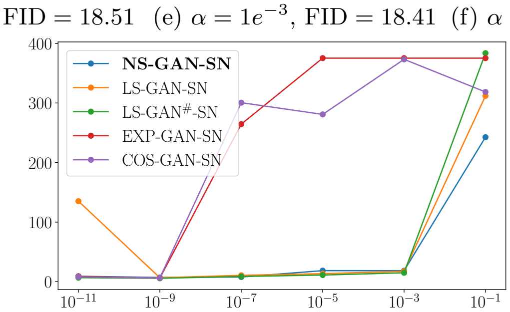

Fig. 26: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 50.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

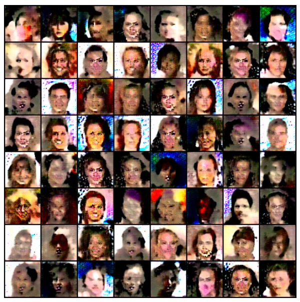

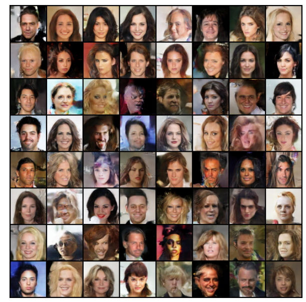

(a) $\alpha = 1e^{-11}$, FID=135.17 (b) $\alpha = 1e^{-9}$, FID = 6.57 (c) $\alpha = 1e^{-7}$, FID = 10.67

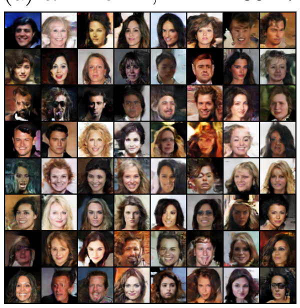

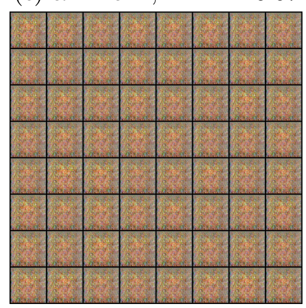

(d) $\alpha = 1e^{-5}$, FID = 13.39 (e) $\alpha = 1e^{-3}$, FID = 17.42 (f) $\alpha = 1e^{-1}$, FID = 311.93

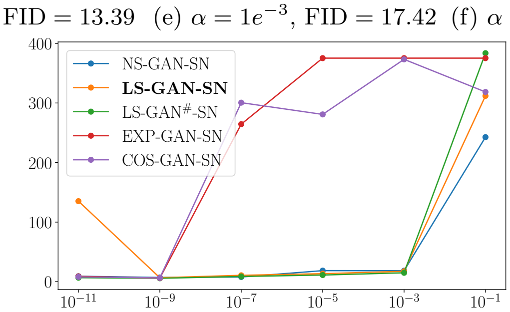

Fig. 27: Samples of randomly generated images with LS-GAN-SN of varying α ($k_{SN} = 50.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

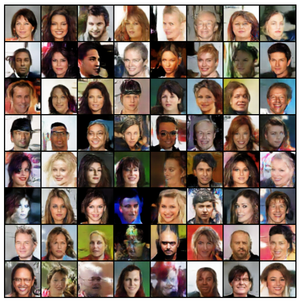
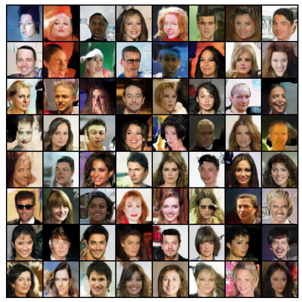


(a) $\alpha = 1e^{-11}$, FID = 6.66 (b) $\alpha = 1e^{-9}$, FID = 5.68 (c) $\alpha = 1e^{-7}$, FID = 8.72


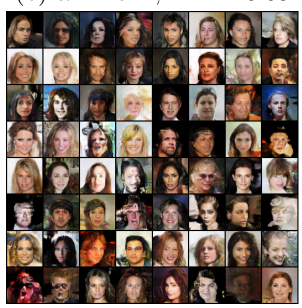
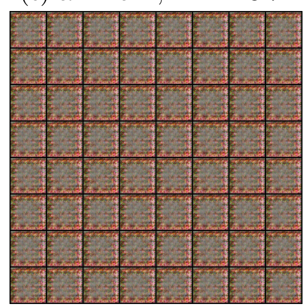

(d) $\alpha = 1e^{-5}$, FID = 11.13 (e) $\alpha = 1e^{-3}$, FID = 14.90 (f) $\alpha = 1e^{-1}$, FID = 383.61

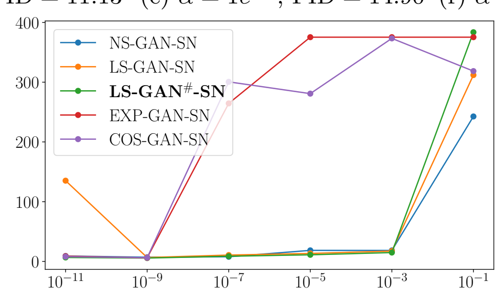

Fig. 28: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α ($k_{SN} = 50.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

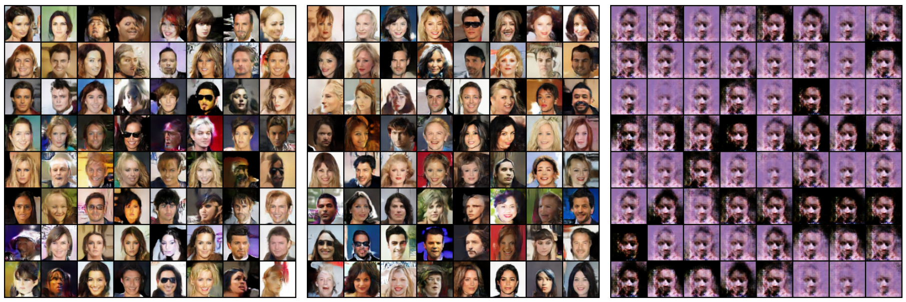

(a) $\alpha = 1e^{-11}$, FID $= 8.85$ (b) $\alpha = 1e^{-9}$, FID $= 6.09$ (c) $\alpha = 1e^{-7}$, FID $= 264.49$

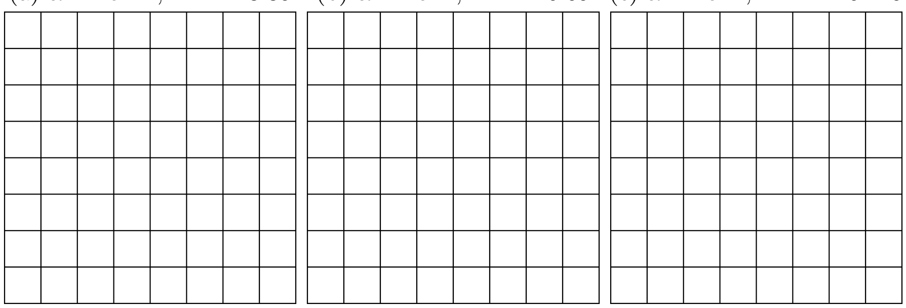

(d) $\alpha = 1e^{-5}$, FID $= 375.32$ (e) $\alpha = 1e^{-3}$, FID $= 375.32$ (f) $\alpha = 1e^{-1}$, FID $= 375.32$

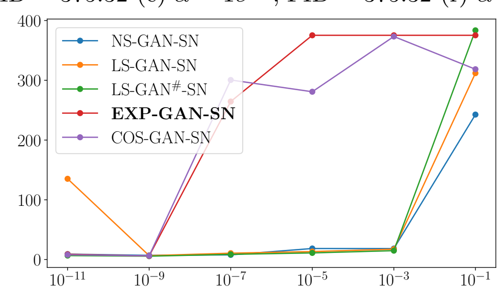

Fig. 29: Samples of randomly generated images with EXP-GAN-SN of varying α ($k_{SN} = 50.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.



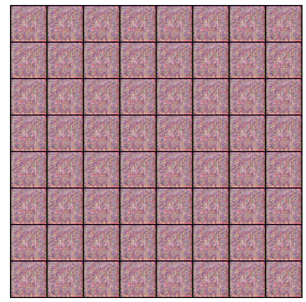

(a) $\alpha = 1e^{-11}$, FID = 8.00 (b) $\alpha = 1e^{-9}$, FID = 6.31 (c) $\alpha = 1e^{-7}$, FID = 300.55

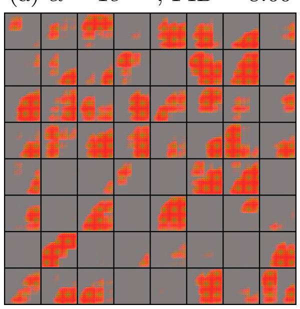
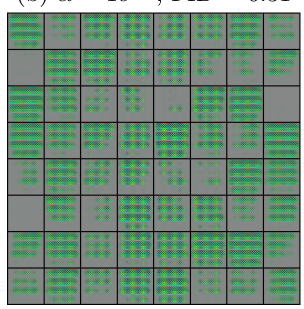
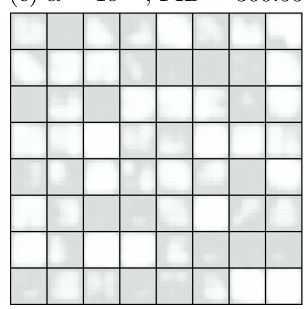

(d) $\alpha = 1e^{-5}$, FID = 280.84 (e) $\alpha = 1e^{-3}$, FID = 373.31 (f) $\alpha = 1e^{-1}$, FID = 318.53

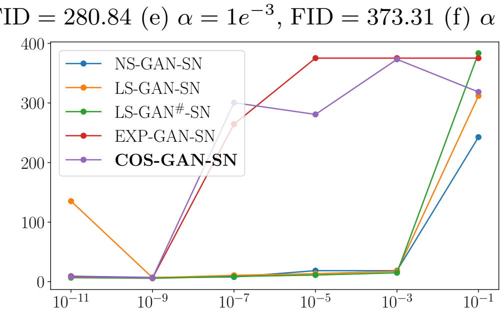

Fig. 30: Samples of randomly generated images with COS-GAN-SN of varying α ($k_{SN} = 50.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 1.0$) of Different Loss Functions

- MNIST -

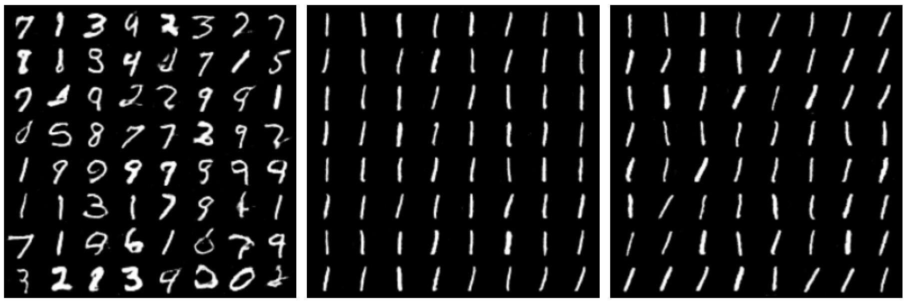

(a) $\alpha = 1e^1$, FID = 6.55 (b) $\alpha = 1e^3$, FID = 148.97 (c) $\alpha = 1e^5$, FID = 134.44

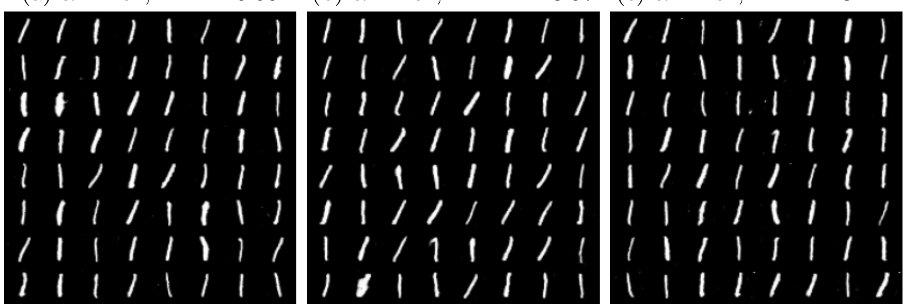

(d) $\alpha = 1e^7$, FID = 133.82 (e) $\alpha = 1e^9$, FID = 130.21 (f) $\alpha = 1e^{11}$, FID = 131.87

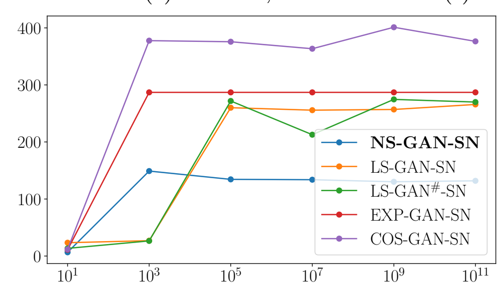

Fig. 31: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 1.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 23.37 (b) $\alpha = 1e^3$, FID = 26.96 (c) $\alpha = 1e^5$, FID = 260.05

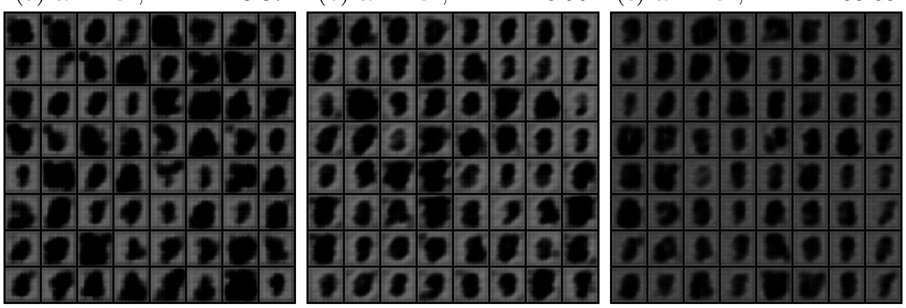

(d) $\alpha = 1e^7$, FID = 255.73 (e) $\alpha = 1e^9$, FID = 256.96 (f) $\alpha = 1e^{11}$, FID = 265.76

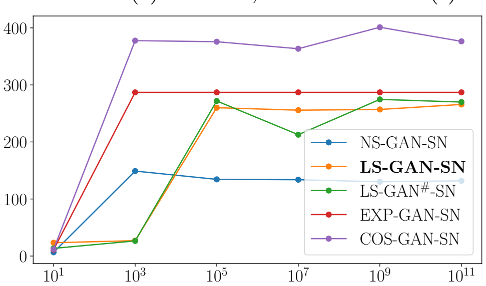

Fig. 32: Samples of randomly generated images with LS-GAN-SN of varying α ($k_{SN} = 1.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 13.43

(b) $\alpha = 1e^3$, FID = 26.51

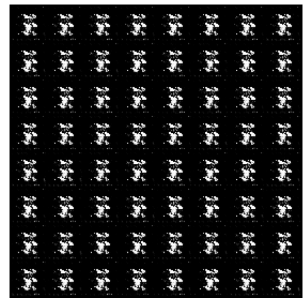

(c) $\alpha = 1e^5$, FID = 271.85

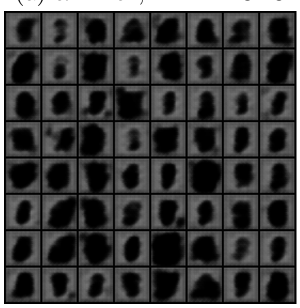

(d) $\alpha = 1e^7$, FID = 212.74

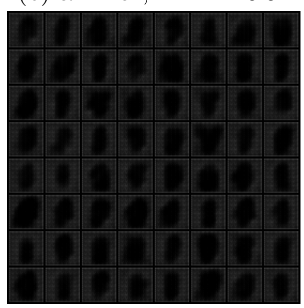

(e) $\alpha = 1e^9$, FID = 274.63

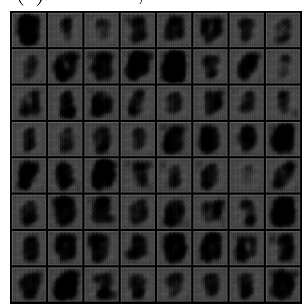

(f) $\alpha = 1e^{11}$, FID = 269.96

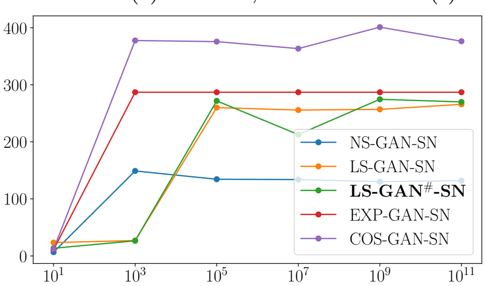

Fig. 33: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α ($k_{SN} = 1.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 11.02

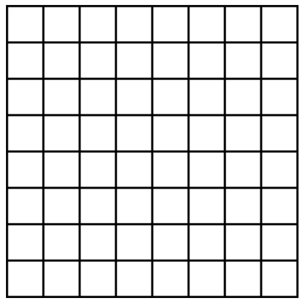

(b) $\alpha = 1e^3$, FID = 286.96

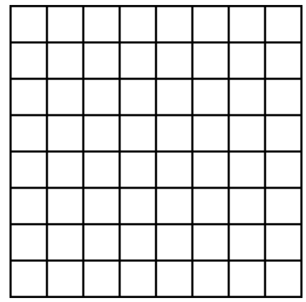

(c) $\alpha = 1e^5$, FID = 286.96

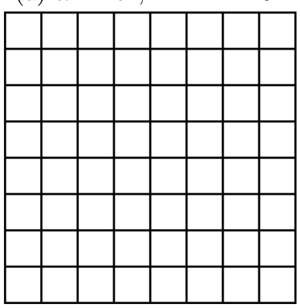

(d) $\alpha = 1e^7$, FID = 286.96

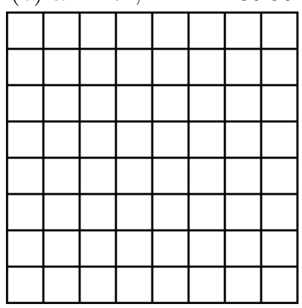

(e) $\alpha = 1e^9$, FID = 286.96

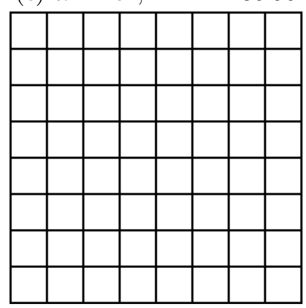

(f) $\alpha = 1e^{11}$, FID = 286.96

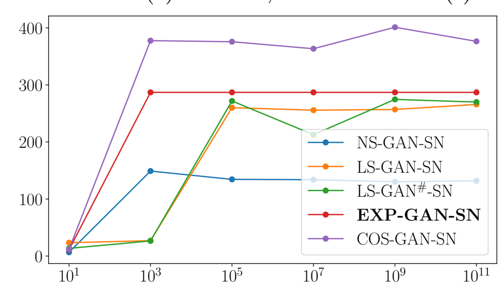

Fig. 34: Samples of randomly generated images with EXP-GAN-SN of varying α ($k_{SN} = 1.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

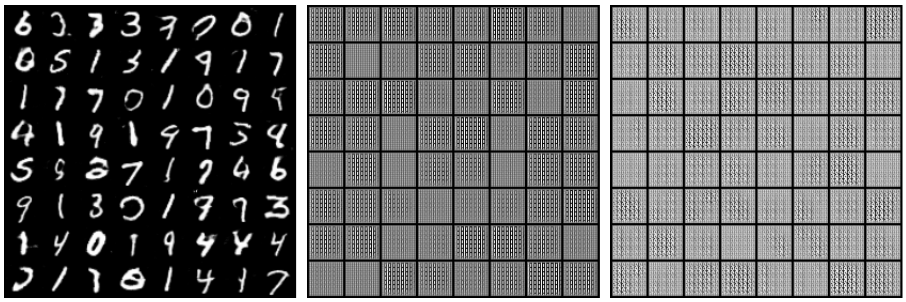

(a) $\alpha = 1e^1$, FID = 11.79 (b) $\alpha = 1e^3$, FID = 377.62 (c) $\alpha = 1e^5$, FID = 375.72

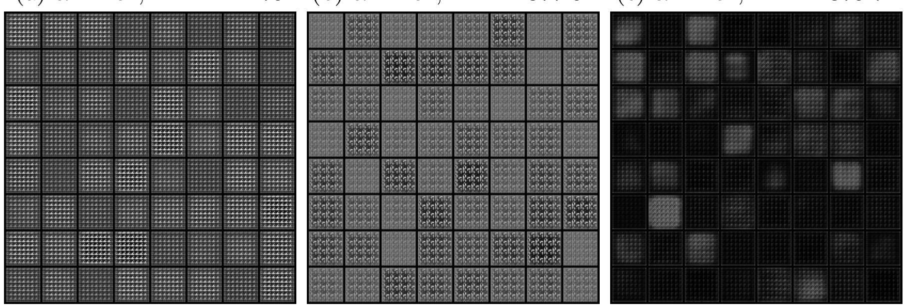

(d) $\alpha = 1e^7$, FID = 363.45 (e) $\alpha = 1e^9$, FID = 401.12 (f) $\alpha = 1e^{11}$, FID = 376.39

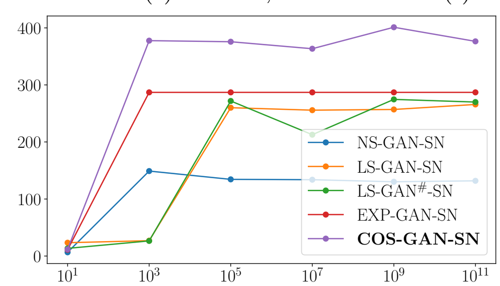

Fig. 35: Samples of randomly generated images with COS-GAN-SN of varying α ($k_{SN} = 1.0$, MNIST). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 1.0$) of Different Loss Functions

- CIFAR10 -

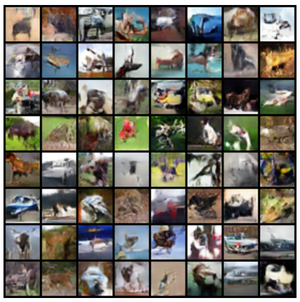

(a) $\alpha = 1e^1$, FID = 17.63 (b) $\alpha = 1e^3$, FID = 47.31 (c) $\alpha = 1e^5$, FID = 46.85

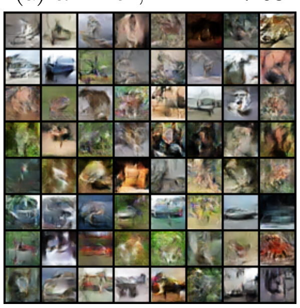

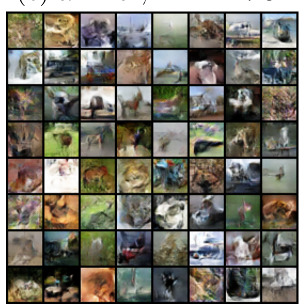

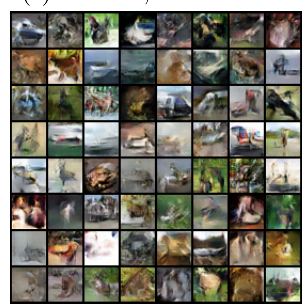

(d) $\alpha = 1e^7$, FID = 45.44 (e) $\alpha = 1e^9$, FID = 45.67 (f) $\alpha = 1e^{11}$, FID = 39.90

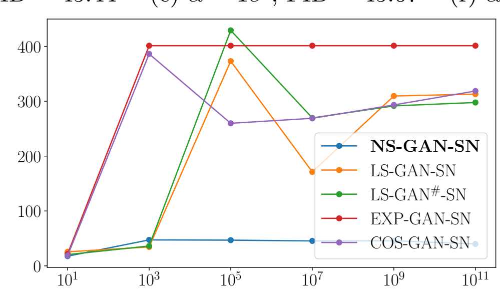

Fig. 36: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 1.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

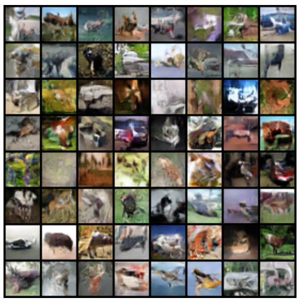

(a) $\alpha = 1e^1$, FID = 25.55

(b) $\alpha = 1e^3$, FID = 34.44

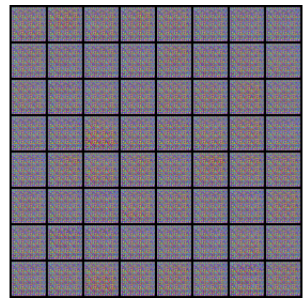

(c) $\alpha = 1e^5$, FID = 373.07

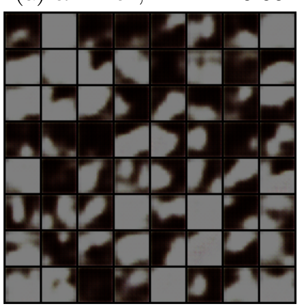

(d) $\alpha = 1e^7$, FID = 171.18

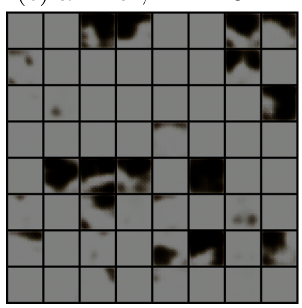

(e) $\alpha = 1e^9$, FID = 309.55

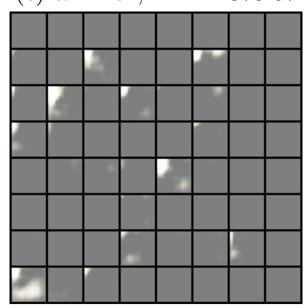

(f) $\alpha = 1e^{11}$, FID = 312.96

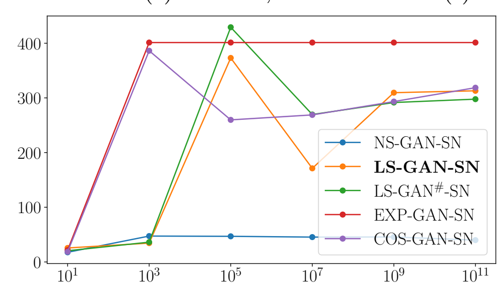

Fig. 37: Samples of randomly generated images with LS-GAN-SN of varying α ($k_{SN} = 1.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

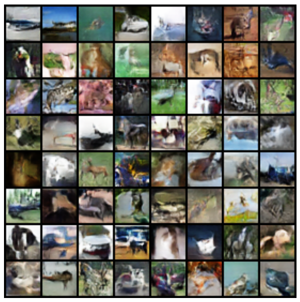

(a) $\alpha = 1e^1$, FID = 20.45

(b) $\alpha = 1e^3$, FID = 36.18

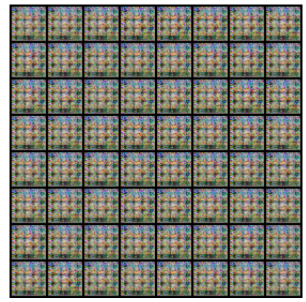

(c) $\alpha = 1e^5$, FID = 429.21

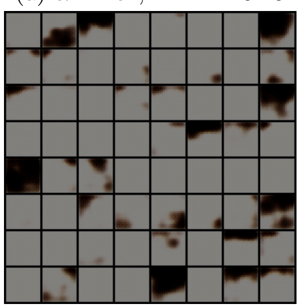

(d) $\alpha = 1e^7$, FID = 269.63

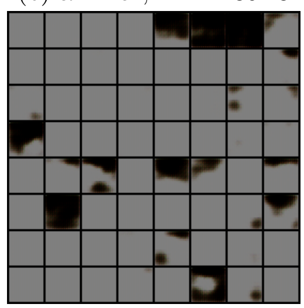

(e) $\alpha = 1e^9$, FID = 291.55

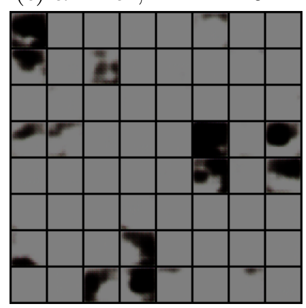

(f) $\alpha = 1e^{11}$, FID = 297.71

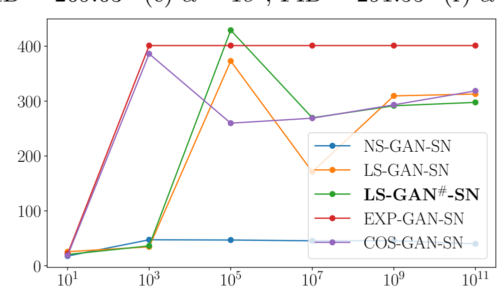

Fig. 38: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α (k_{SN} = 1.0, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

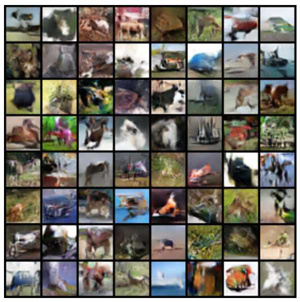
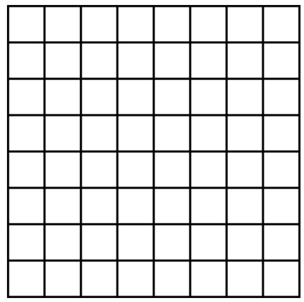
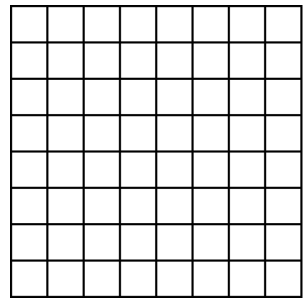

(a) $\alpha = 1e^1$, FID = 21.56 (b) $\alpha = 1e^3$, FID = 401.24 (c) $\alpha = 1e^5$, FID = 401.24

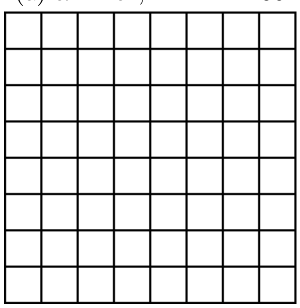
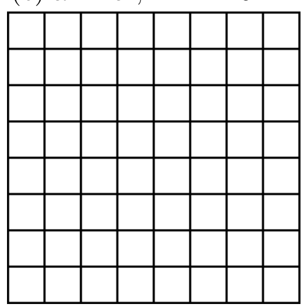
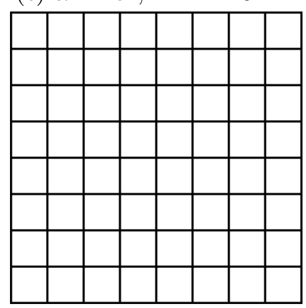

(d) $\alpha = 1e^7$, FID = 401.24 (e) $\alpha = 1e^9$, FID = 401.24 (f) $\alpha = 1e^{11}$, FID = 401.24

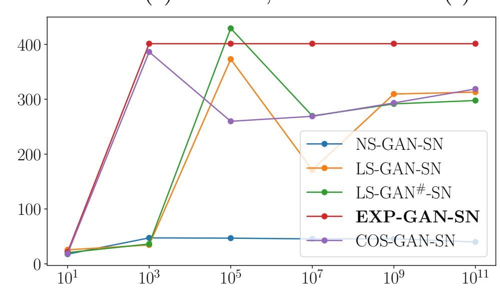

Fig. 39: Samples of randomly generated images with EXP-GAN-SN of varying α ($k_{SN} = 1.0$, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

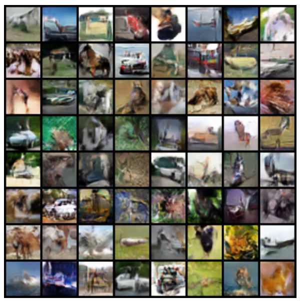

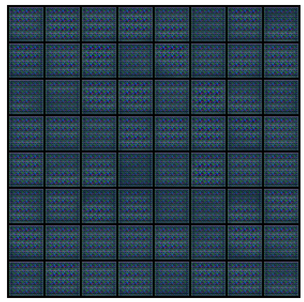

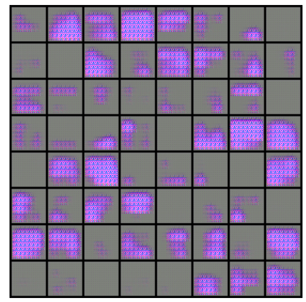

(a) $\alpha = 1e^1$, FID = 18.59 (b) $\alpha = 1e^3$, FID = 386.24 (c) $\alpha = 1e^5$, FID = 259.83

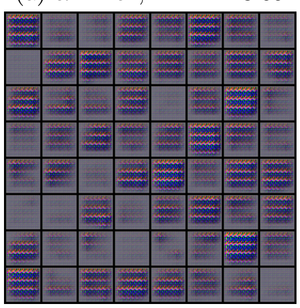

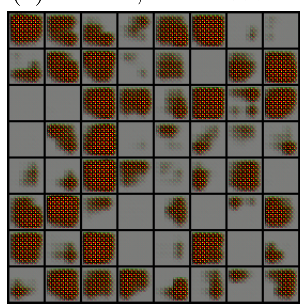

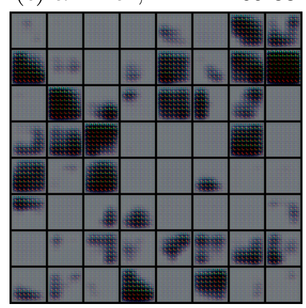

(d) $\alpha = 1e^7$, FID = 268.89 (e) $\alpha = 1e^9$, FID = 293.29 (f) $\alpha = 1e^{11}$, FID = 318.65

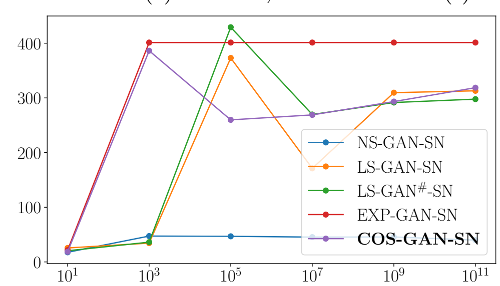

Fig. 40: Samples of randomly generated images with COS-GAN-SN of varying α (k_{SN} = 1.0, CIFAR10). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores v.s. α ($k_{SN} = 1.0$) of Different Loss Functions

- CelebA -

(a) $\alpha = 1e^1$, FID = 5.88

(b) $\alpha = 1e^3$, FID = 16.14

(c) $\alpha = 1e^5$, FID = 17.75

(d) $\alpha = 1e^7$, FID = 17.67

(e) $\alpha = 1e^9$, FID = 16.87

(f) $\alpha = 1e^{11}$, FID = 18.81

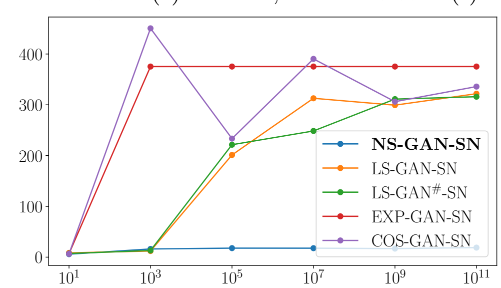

Fig. 41: Samples of randomly generated images with NS-GAN-SN of varying α ($k_{SN} = 1.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

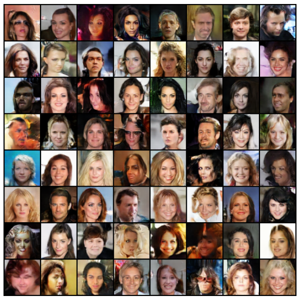

(a) $\alpha = 1e^1$, FID = 8.41

(b) $\alpha = 1e^3$, FID = 12.09

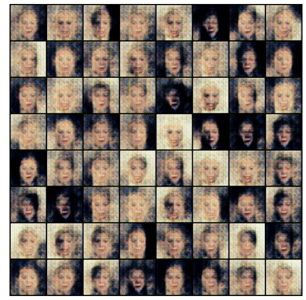

(c) $\alpha = 1e^5$, FID = 201.22

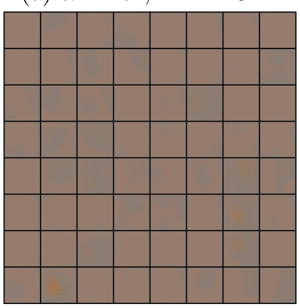

(d) $\alpha = 1e^7$, FID = 312.83

(e) $\alpha = 1e^9$, FID = 299.30

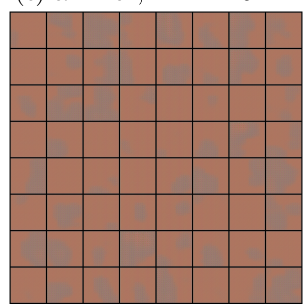

(f) $\alpha = 1e^{11}$, FID = 321.84

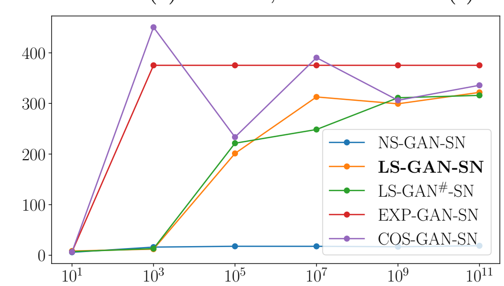

Fig. 42: Samples of randomly generated images with LS-GAN-SN of varying α ($k_{SN} = 1.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 7.21

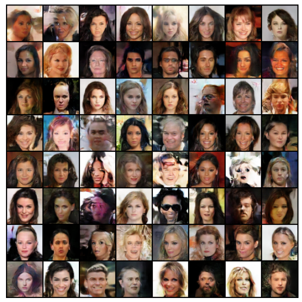

(b) $\alpha = 1e^3$, FID = 13.13

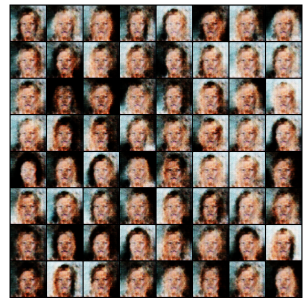

(c) $\alpha = 1e^5$, FID = 221.41

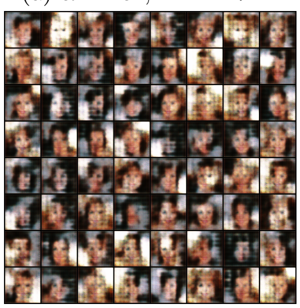

(d) $\alpha = 1e^7$, FID = 248.48

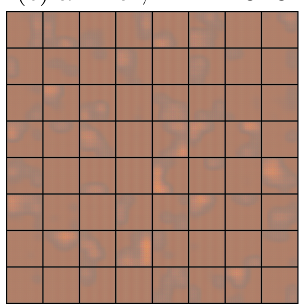

(e) $\alpha = 1e^9$, FID = 311.21

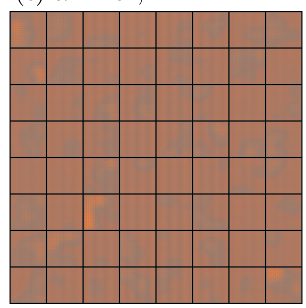

(f) $\alpha = 1e^{11}$, FID = 315.94

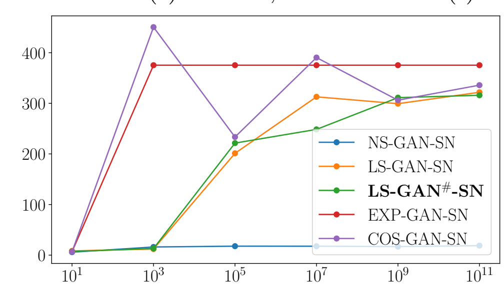

Fig. 43: Samples of randomly generated images with LS-GAN$^{\#}$-SN of varying α ($k_{SN} = 1.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 6.91

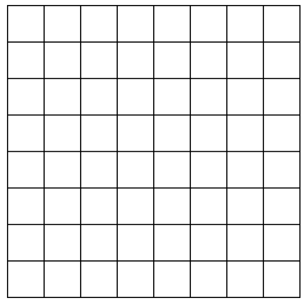

(b) $\alpha = 1e^3$, FID = 375.32

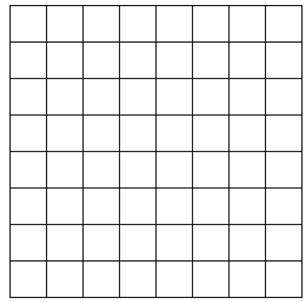

(c) $\alpha = 1e^5$, FID = 375.32

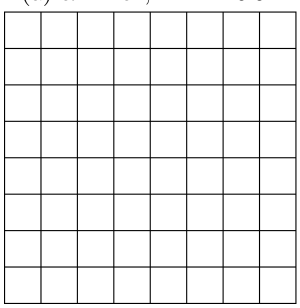

(d) $\alpha = 1e^7$, FID = 375.32

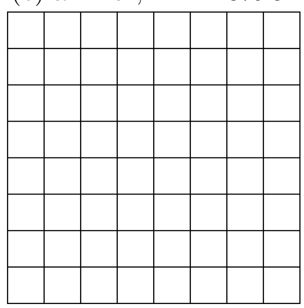

(e) $\alpha = 1e^9$, FID = 375.32

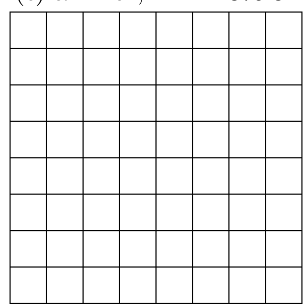

(f) $\alpha = 1e^{11}$, FID = 375.32

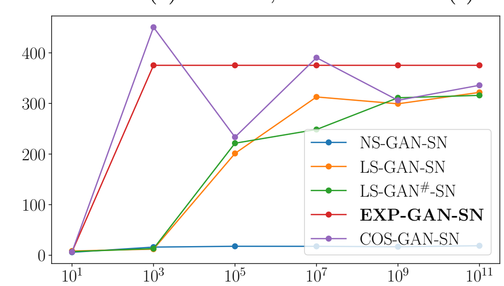

Fig. 44: Samples of randomly generated images with EXP-GAN-SN of varying α ($k_{SN} = 1.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

(a) $\alpha = 1e^1$, FID = 6.62

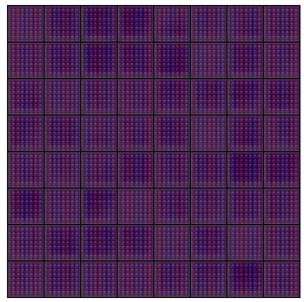

(b) $\alpha = 1e^3$, FID = 450.57

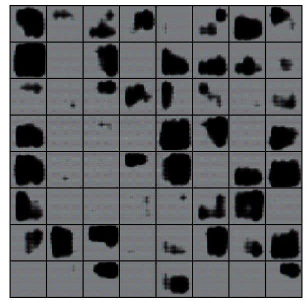

(c) $\alpha = 1e^5$, FID = 233.42

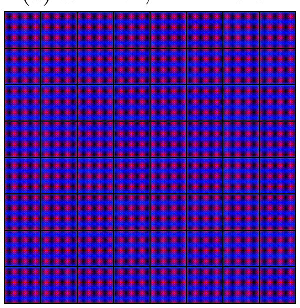

(d) $\alpha = 1e^7$, FID = 390.40

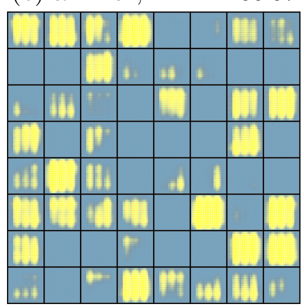

(e) $\alpha = 1e^9$, FID = 306.17

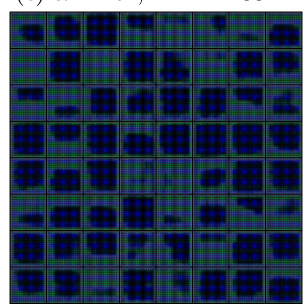

(f) $\alpha = 1e^{11}$, FID = 335.87

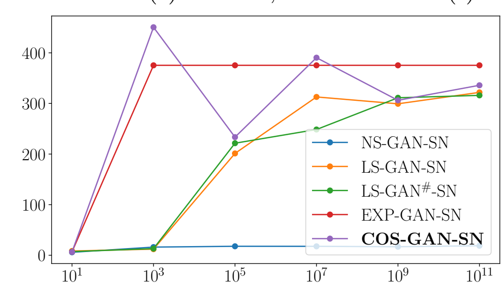

Fig. 45: Samples of randomly generated images with COS-GAN-SN of varying α ($k_{SN} = 1.0$, CelebA). For the line plot, x-axis shows α (in log scale) and y-axis shows the FID scores.

FID scores of WGAN-SN and some extremely degenerate loss functions ($\alpha = 1e^{-25}$) on different datasets

(a) WGAN, FID = 3.71

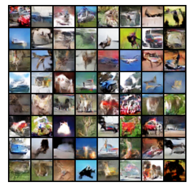

(b) WGAN, FID = 23.36

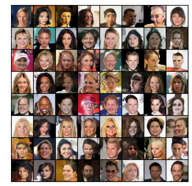

(c) WGAN, FID = 7.82

(d) NSGAN, FID = 3.74

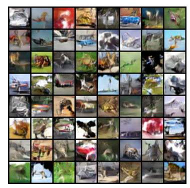

(e) NSGAN, FID = 21.92

(f) NSGAN, FID = 8.10

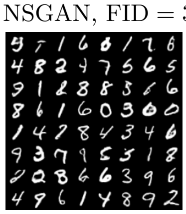

(g) LSGAN#, FID = 3.81

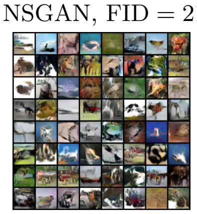

(h) LSGAN#, FID = 21.47

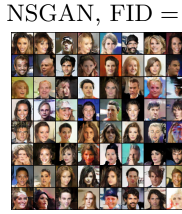

(i) LSGAN#, FID = 8.51

(j) COSGAN, FID = 3.96

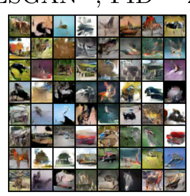

(k) COSGAN, FID = 23.65

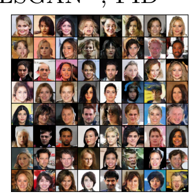

(l) COSGAN, FID = 8.30

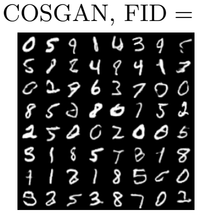

(m) EXPGAN, FID = 3.86

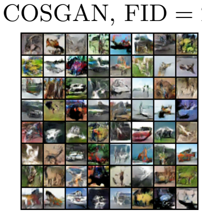

(n) EXPGAN, FID = 21.91

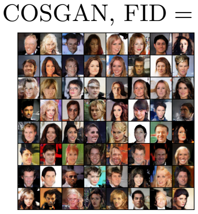

(o) EXPGAN, FID = 8.22

Fig. 46: Samples of randomly generated images with WGAN-SN and some extremely degenerate loss functions ($\alpha = 1e^{-25}$) on different datasets. We use $k_{SN} = 50$ for all our experiments.