Supplementary Material
Re-GAN: Data-Efficient GANs Training via Architectural Reconfiguration

This document presents the supplementary materials omitted from the main paper due to space limitation.

1. Comparison to state-of-the-art

In this section, we show the qualitative results (see Figure 1) and Precision-Recall curve for StyleGAN2 and Re-StyleGAN2 on CIFAR-10 full training dataset (see Figure 2). Instead of having a scalar for generative quality, Precision-Recall (PR) curve for GANs distinguish between mode-collapse (low recall) and bad quality (low precision). Results shows that Re-StyleGAN2 is performing well, i.e., generating high-quality images with more variation. F_8 and $F_{1/8}$ results in Table 1 shows that similar to StyleGAN2, Re-StyleGAN2 is also weighing precision higher than recall.

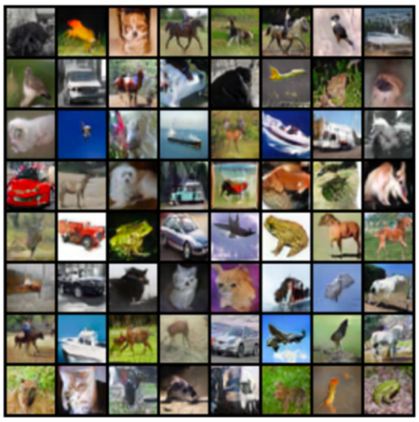

Figure 1: Qualitative result of Re-StyleGAN2 for CIFAR-10 dataset

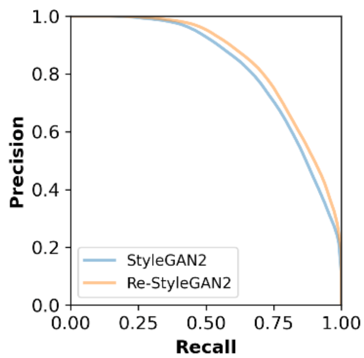

Figure 2: Precision and Recall curve[1]. Area under curve (AUC) of Re-StyleGAN2 for CIFAR-10 dataset is larger than dense one, which means Re-StyleGAN2 improves the performance of StyleGAN2. We use 50k generated images for both models and take the whole training set of CIFAR-10 as reference distribution.

Precision: It is the probability that one image from generated distribution falls into the support of the real distribution.
Recall: It is the probability that one image from real distribution falls into the support of the generated distribution.
F-score: A F-score is the harmonic average of precision and recall, which is defined by

$$F_\beta score = \frac{(1+\beta^2)precision \times recall}{\beta^2 precision + recall}$$

We follow the similar settings as in the original paper [1], $\beta = 8$, 1/8.

Table 1: F_8 and $F_{1/8}$ score of dense StyleGAN2 and Re-StyleGAN2@90. The best performance is highlighted.

Models	F_8	$F_{1/8}$
StyleGAN2	0.939	0.957
Re-StyleGAN2@90	**0.948**	**0.962**

[1] https://github.com/DrLSimon/precision-recall-distributions-icml19

Sparse training phases act as a regularizer, preventing the dense model from memorizing corrupted samples. Generally, it is said that memorization causes the quality degradation of image generation. However, it helps generator in learning true distribution without missing modes. Therefore it is important to show Re-GAN can capture all modes too. We generate 25 mixtures of Gaussian. We set update schedule to 5000 and sparsity 10% for Re-GAN. We use 2500 generated points for modes and high quality computation. Results in Table 2 show that Re-GAN can capture more number of modes and generate high quality images than baselines.

Table 2: Mode collapse and high quality study

Methods	Modes	High quality
GANs [2]	23±2.57	0.27±0.13
Re-GAN	**24.4±1.02**	**0.30±0.06**
ALI [3]	23.4±1.6	0.35±0.06
Re-ALI	**23.8±0.98**	**0.36±0.08**
GDPPGAN [5]	23.8±1.17	0.27±0.06
Re-GDPPGAN	**24.8±0.40**	**0.39±0.04**

2. Few-shot generation

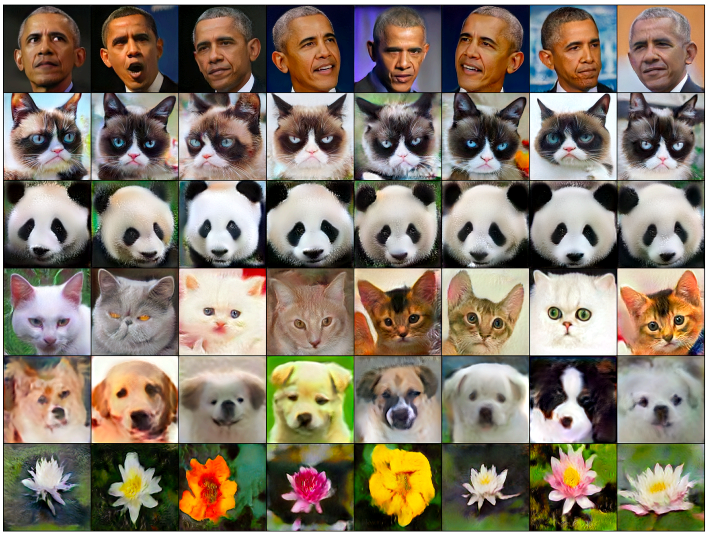

Figure 3: Few-shot generalization results of Re-StyleGAN2 on (256×256) 100-shot Obama, Grumpy Cat, Panda, A-Cat/Dog and Oxford flowers, respectively. Samples are randomly drawn without cherry picking.

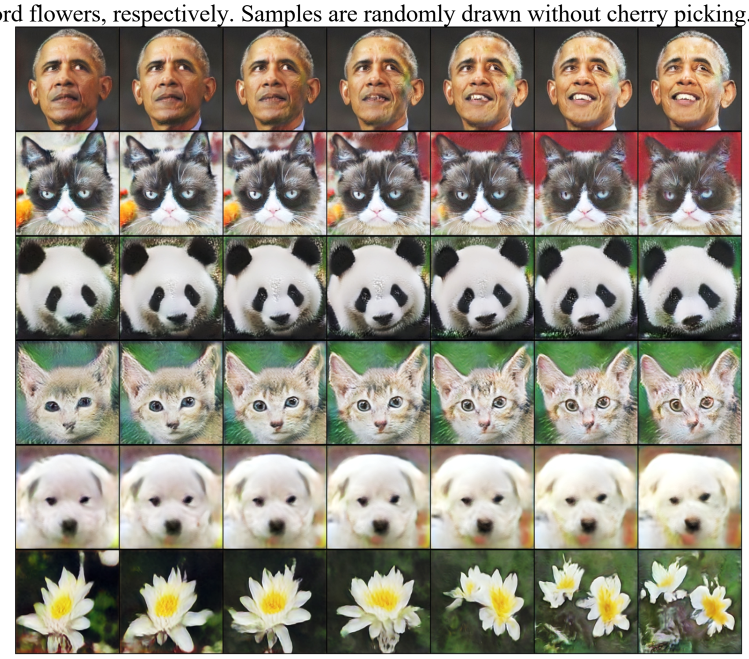

Figure 4: Style interpolations of Re-StyleGAN2 on 100-shot Obama, Grumpy Cat, Panda, A-Cat/Dog and Oxford flowers, respectively.

Results in Table 3 shows that Re-StyleGAN2 is not only achieving better FID score, it is also minimizing training time and training FLOPs for all 1k, 5k and 10k datasets.

Table 3: FID comparison on few-shot datasets at 1024×1024 resolution. FID and IS are calculated using 50k randomly generated samples, with the training data (70k) serving as the reference distribution. The best performance is highlighted.

Models	10k				5k				1k			
	Flops x10^{17}	#RI (in M)	Training Time (in hours)	FID↓	Flops x10^{17}	#RI (in M)	Training Time (in hours)	FID↓	Flops x10^{17}	#RI (in M)	Training Time (in hours)	FID↓
StyleGAN2	5.81	1.08	130	13.06	2.58	0.48	56	21.76	1.45	0.27	31	40.24
Re-StyleGAN2	**4.94**	**0.94**	**111**	**11.22**	**1.85**	**0.45**	**52**	**19.13**	**1.03**	**0.21**	**24**	**36.3**

To further demonstrate the effectiveness of our Re-GAN, we apply it to other recent methods, FastGAN[2] [4] for several few-shot datasets. The results in Table 4 shows that our Re-GAN is consistently outperforming than recent few-shot generation method, FastGAN. Re-FastGAN is not only achieving better FID score, it is minimizing training time and training FLOPs for all challenging few-shot datasets, Obama, Grumpy-Cat (G-Cat), Panda, Animal-Face (Cat and Dog) (named, A-Cat and A-Dog).

Table 4: FID comparison on few-shot datasets at 256×256 resolution. FID and IS are calculated using 5k randomly generated samples, with the training data serving as the reference distribution. The best performance is highlighted.

Datasets	# of Images	Models	#Real Images (RI)(in k)	Training Time (in hours)	FLOPs $\times 10^{15}$	FID↓
Obama	100	FastGAN	360	2.8	7.14	40.96
		Re-FastGAN	**320**	**2.5**	**6.27**	**39.84**
G-Cat	100	FastGAN	400	3.1	7.93	25.71
		Re-FastGAN	**280**	**2.2**	**5.41**	**24.86**
Panda	100	FastGAN	400	3.1	7.93	10.58
		Re-FastGAN	**360**	**2.8**	**6.86**	**10.26**
A-Cat	389	FastGAN	400	3.1	7.93	35.54
		Re-FastGAN	**280**	**2.2**	**5.41**	**34.13**
A-Dog	180	FastGAN	720	5.6	14.27	53.28
		Re-FastGAN	**640**	**5**	**12.4**	**52.16**

References

[1] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. "Assessing generative models via precision and recall." Advances in neural information processing systems, 31, 2018.

[2] I. J. Goodfellow et al., "Generative adversarial nets," in Advances in Neural Information Processing Systems, 2014, vol. 3, no. January, pp. 2672–2680.s

[3] Dumoulin, Vincent, Ishmael Belghazi, Ben Poole, Olivier Mastropietro, Alex Lamb, Martin Arjovsky, and Aaron Courville. "Adversarially learned inference." arXiv preprint arXiv:1606.00704 (2016).

[4] Liu, Bingchen, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal. "Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis." In International Conference on Learning Representations. 2020.

[5] M. Elfeki, C. Couprie, M. Riviere, and M. Elhoseiny. "Gdpp: Learning diverse generations using determinantal point processes." In ICML, pp. 1774-1783, 2019.

[2] https://github.com/odegeasslbc/FastGAN-pytorch