Deceive D: Adaptive Pseudo Augmentation for GAN Training with Limited Data

Liming Jiang[1] Bo Dai[1] Wayne Wu[2] Chen Change Loy[1]
[1]S-Lab, Nanyang Technological University [2]SenseTime Research

{liming002, bo.dai, ccloy}@ntu.edu.sg wuwenyan@sensetime.com

NEURAL INFORMATION PROCESSING SYSTEMS

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE | S-LAB FOR ADVANCED INTELLIGENCE | 商汤 sensetime

Motivation

- Poor quality under limited data
- Sufficient data is sometimes infeasible
- Practical deployment of GANs

Underlying Cause: *Discriminator Overfitting*

NANYANG TECHNOLOGICAL UNIVERSITY SINGAPORE | S-LAB FOR ADVANCED INTELLIGENCE | 商汤 sensetime

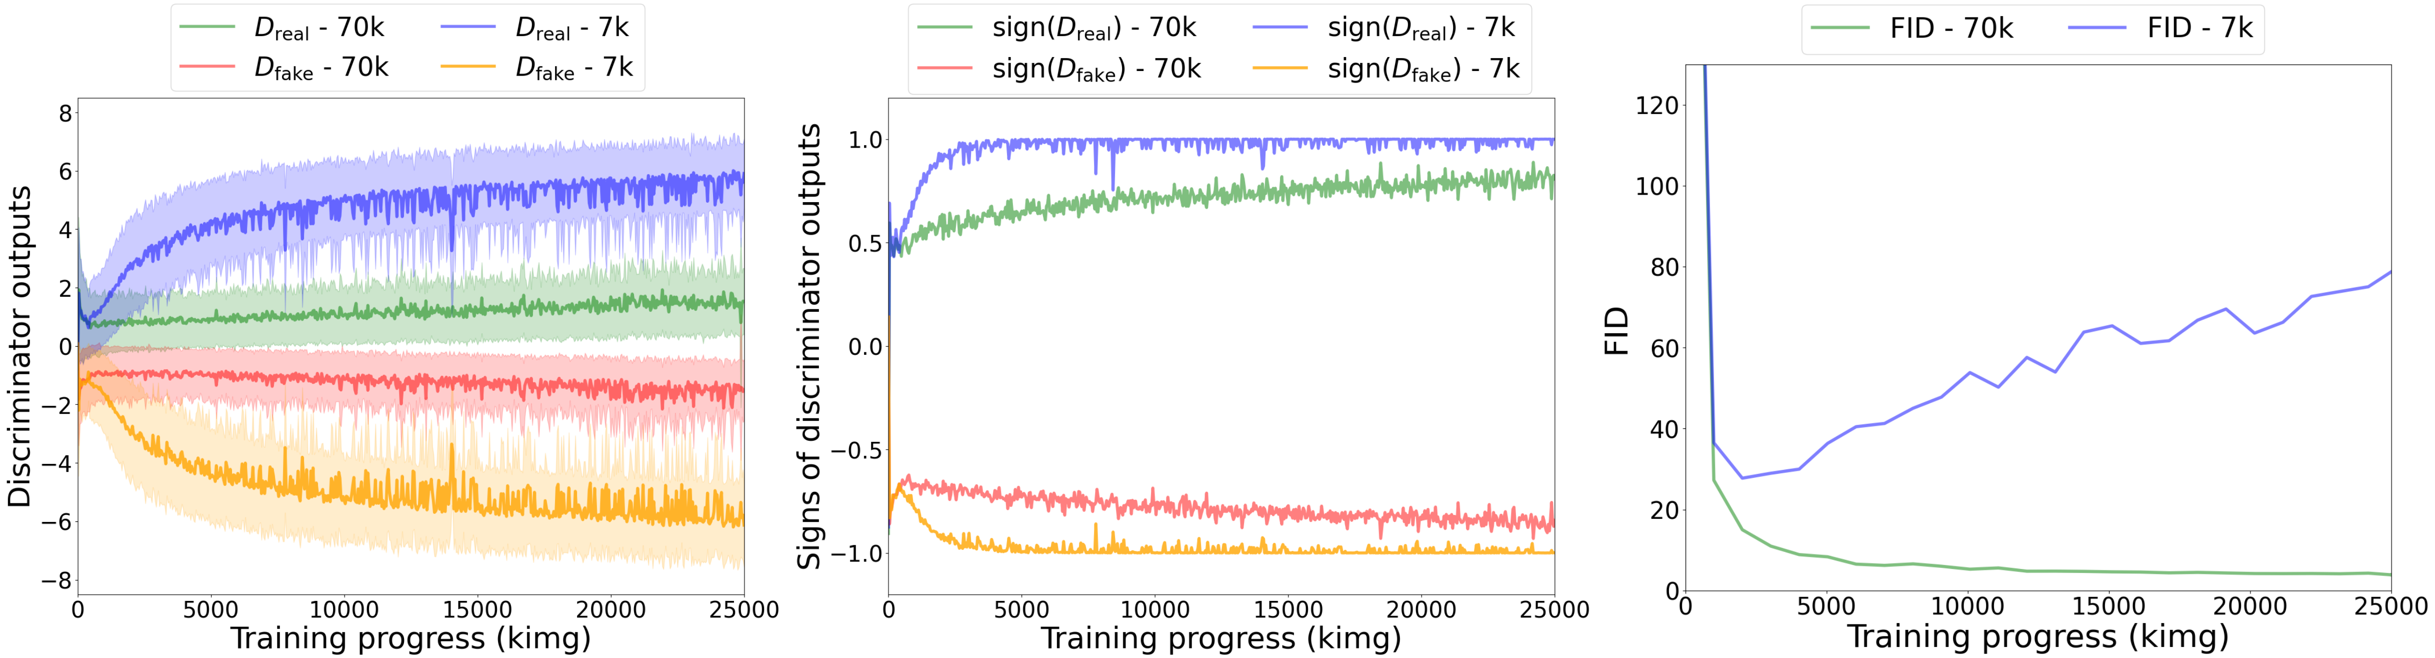

(a) Discriminator raw output logits **(b) Signs of discriminator outputs** **(c) Training convergence measured by FID**

- The discriminator predictions diverge much more rapidly with limited training data, indicating quick overfitting.
- The overfitting of discriminator impedes the generator’s convergence, rendering severe instability of training dynamics.
- The less informative feedback to the generator leads it to converge to an inferior point, compromising the quality of synthesized images.

Previous Solutions

- Standard Data Augmentations (*e.g.,* ADA)

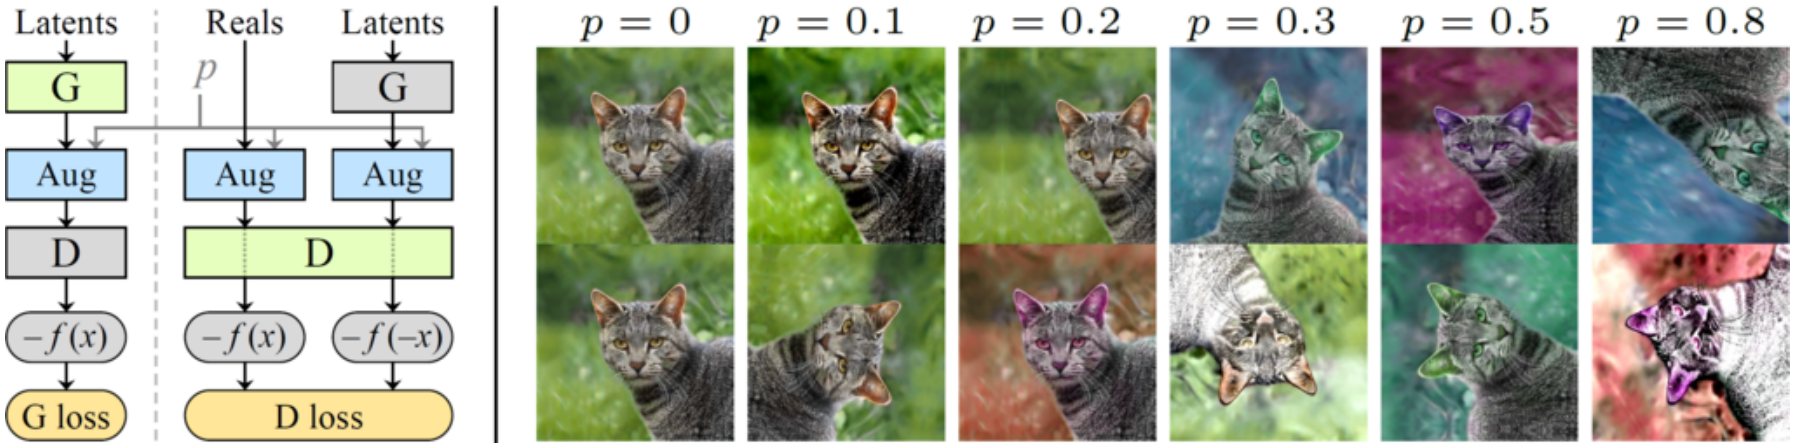

- Model Regularization (*e.g.,* LC-regularization)

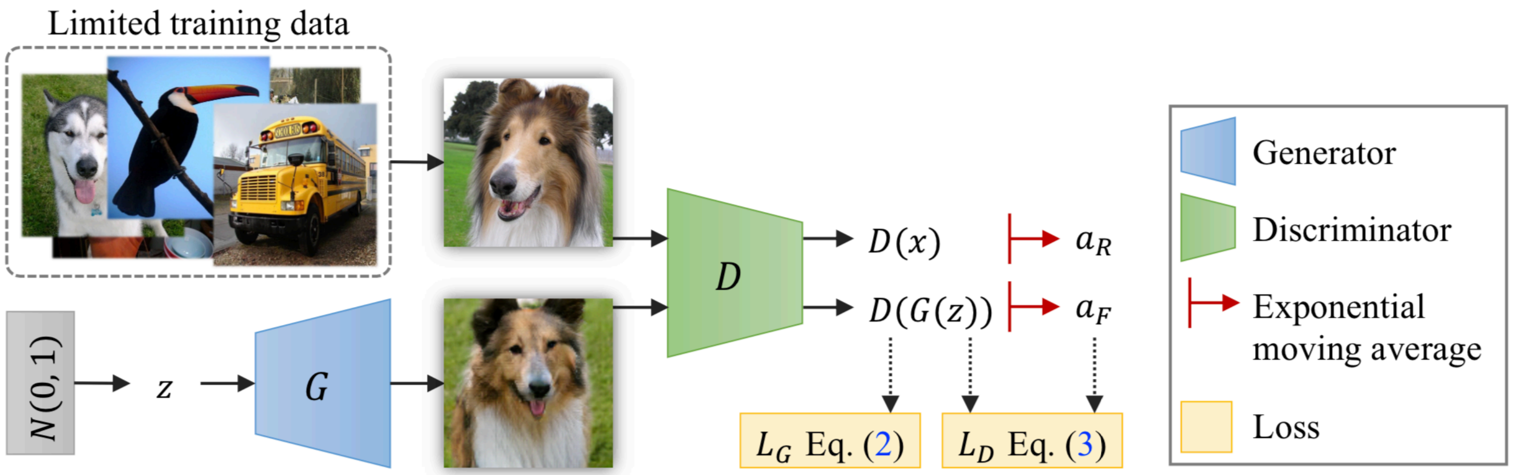

$$R_{\mathrm{LC}} = \mathop{\mathbb{E}}_{\boldsymbol{x}\sim\mathcal{T}}\left[\|D(\boldsymbol{x})-\alpha_F\|^2\right] + \mathop{\mathbb{E}}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}}\left[\|D(G(\boldsymbol{z}))-\alpha_R\|^2\right].$$

$$\min_D L_D, L_D = -V_D + \lambda R_{\mathrm{LC}}(D),$$

Methodology

- Adaptive Pseudo Augmentation (APA)

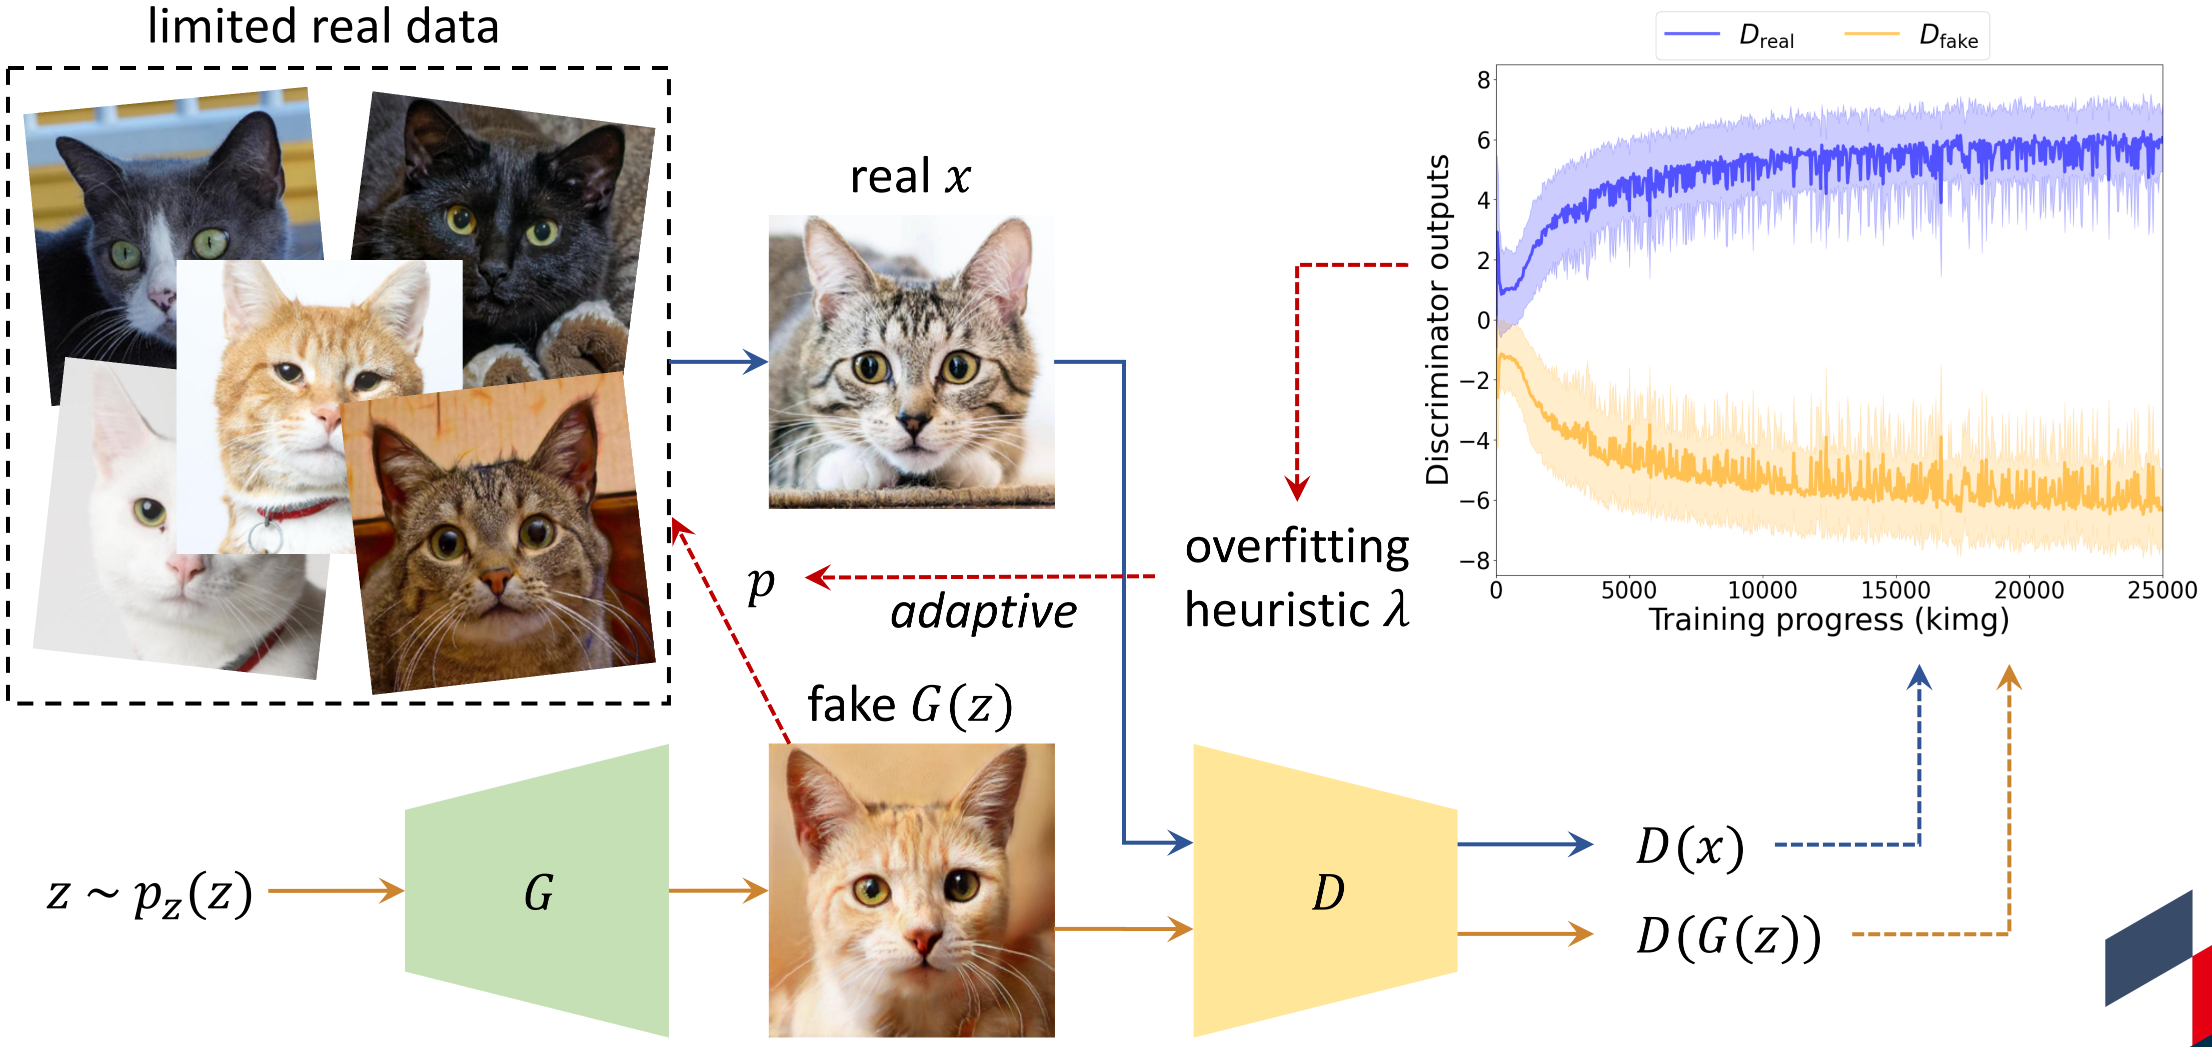

Methodology

- Overfitting Heuristic λ: *three variants*

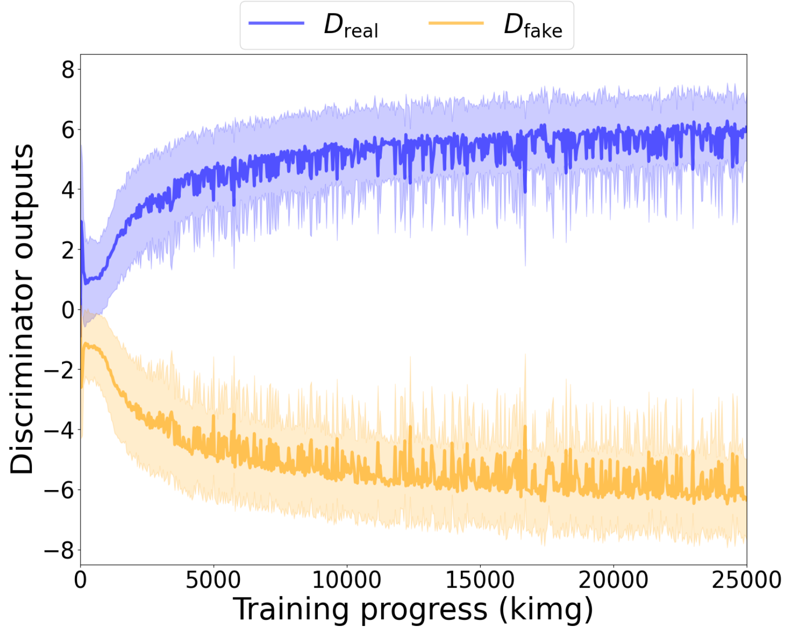

$$\lambda_r = \mathbb{E}\left(\operatorname{sign}\left(D_{\text{real}}\right)\right), \quad \lambda_f = -\mathbb{E}\left(\operatorname{sign}\left(D_{\text{fake}}\right)\right), \quad \lambda_{rf} = \frac{\mathbb{E}\left(\operatorname{sign}\left(D_{\text{real}}\right)\right) - \mathbb{E}\left(\operatorname{sign}\left(D_{\text{fake}}\right)\right)}{2},$$

where D_{real} and D_{fake} are defined as

$$D_{\text{real}} = \operatorname{logit}\left(D\left(x\right)\right), \quad D_{\text{fake}} = \operatorname{logit}\left(D\left(G\left(z\right)\right)\right),$$

** For all these heuristics, $\lambda = 0$ represents no overfitting, and $\lambda = 1$ means complete overfitting. The deception probability p is adjusted adaptively according to the overfitting heuristic λ (using λ_r by default).*

Results and Analysis

- Effectiveness of APA: *on various datasets*

StyleGAN2 | **APA (Ours)**

AFHQ-Cat-5k (5,153 img, 100% data)

StyleGAN2 | **APA (Ours)**

FFHQ-5k (5,000 img, ~7% data)

StyleGAN2 | **APA (Ours)**

Anime-5k (5,000 img, ~2% data)

StyleGAN2 | **APA (Ours)**

CUB-12k (11,788 img, 100% data)

Results and Analysis

- Effectiveness of APA: *given different data amounts*

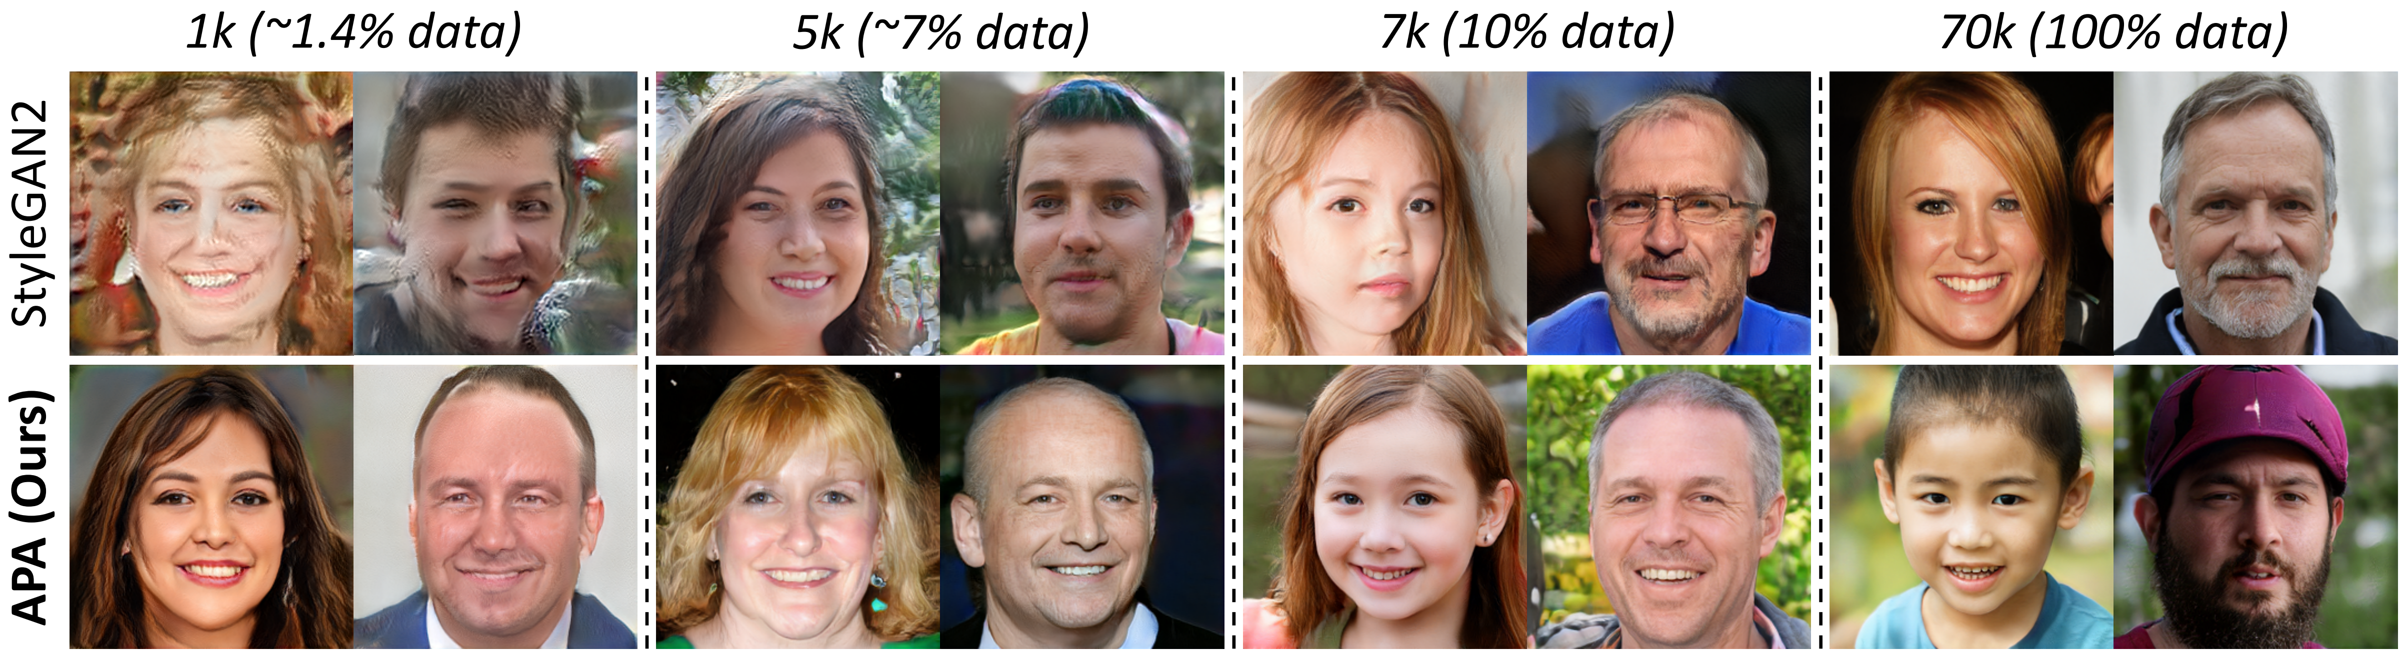

Different subsets of FFHQ

Results and Analysis

- Effectiveness of APA: *overfitting and convergence analysis*

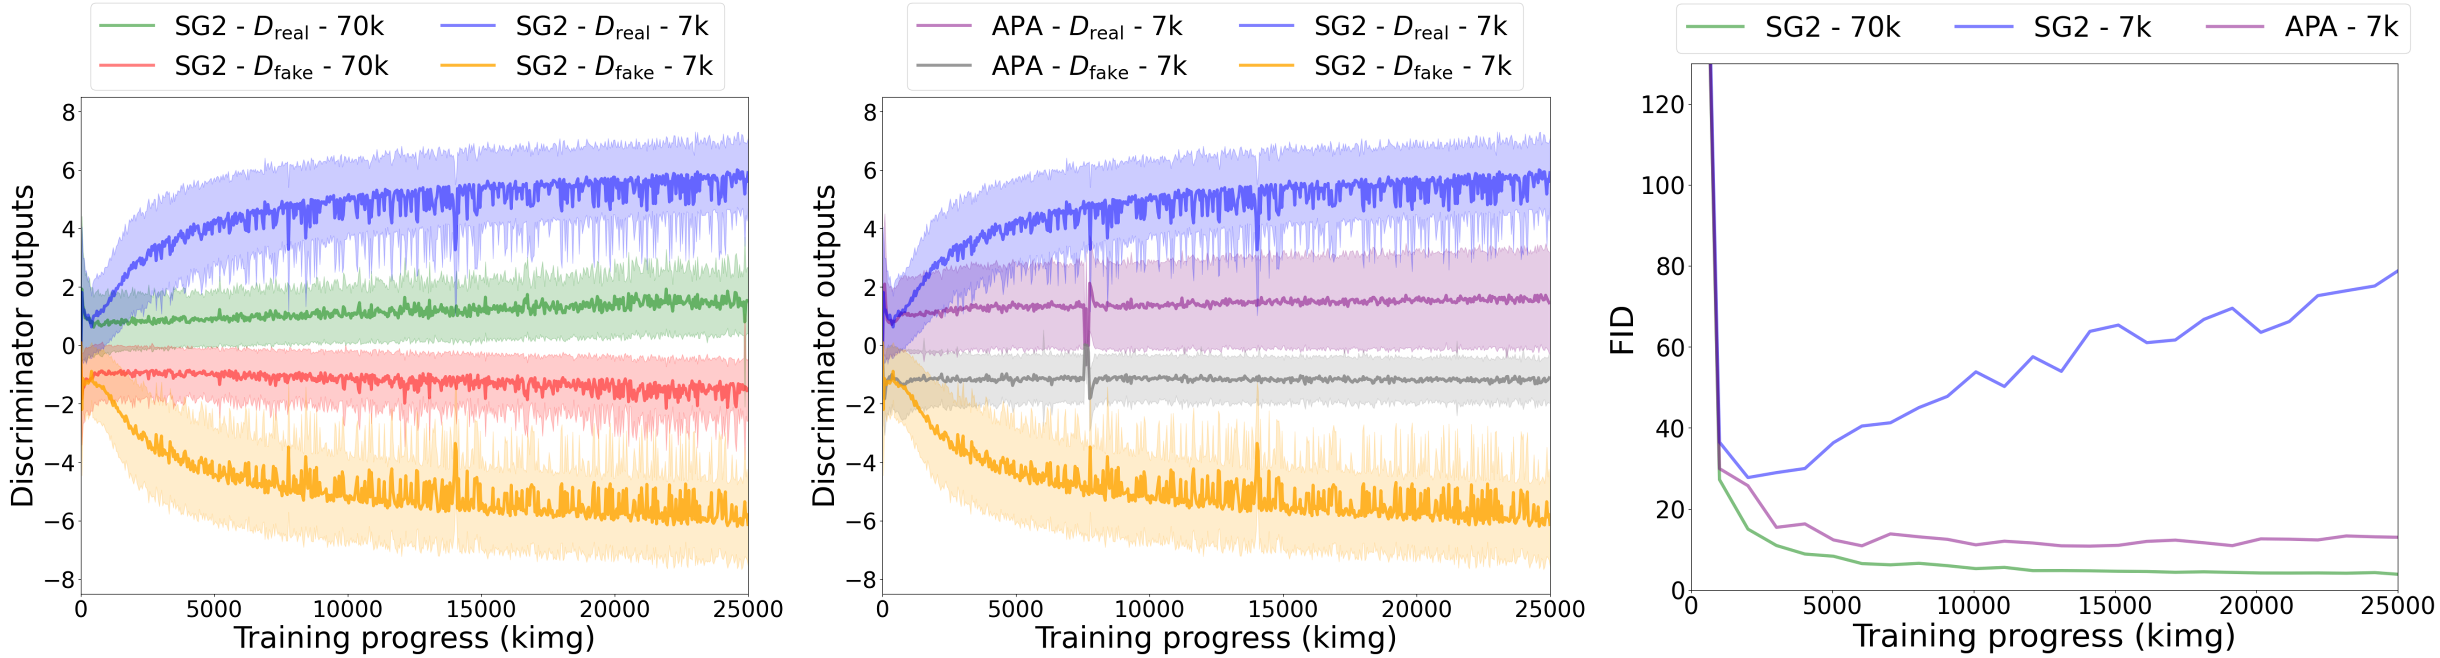

(a) Discriminator raw output logits of StyleGAN2 on the full (70k) or limited (7k) datasets

(b) Discriminator raw output logits of StyleGAN2 and APA on the limited (7k) dataset

(c) Training convergence measured by FID

Overfitting and convergence status of APA compared to StyleGAN2 (SG2) on FFHQ

Results and Analysis

- Comparison with Other State-of-the-Art Solutions

Training cost: average training time

- StyleGAN2: (4.740 ± 0.100) sec/kimg
- ADA: (5.327 ± 0.116) sec/kimg
- **APA (Ours):** (4.789 ± 0.078) sec/kimg

** Negligible computational cost of APA*

Results and Analysis

- Higher-Resolution Examples (1024 × 1024): *on FFHQ-5k (5,000 images, ~7% of full data)*

StyleGAN2

APA (Ours)

Results and Analysis

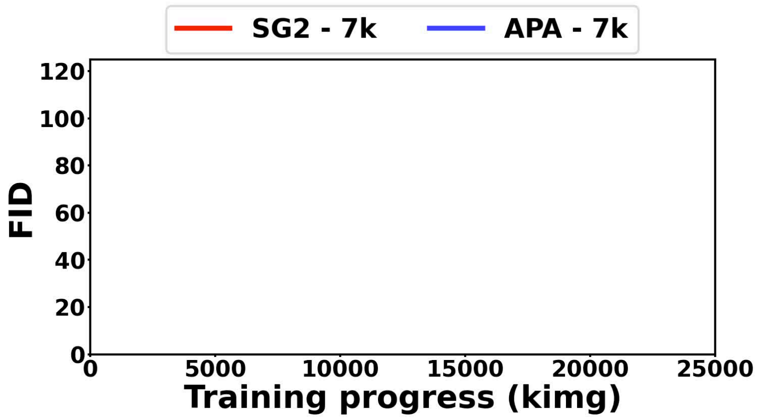

- Additional Training Convergence Visualizations: *on FFHQ-7k (7,000 images, 10% of full data)*

StyleGAN2

APA (Ours)

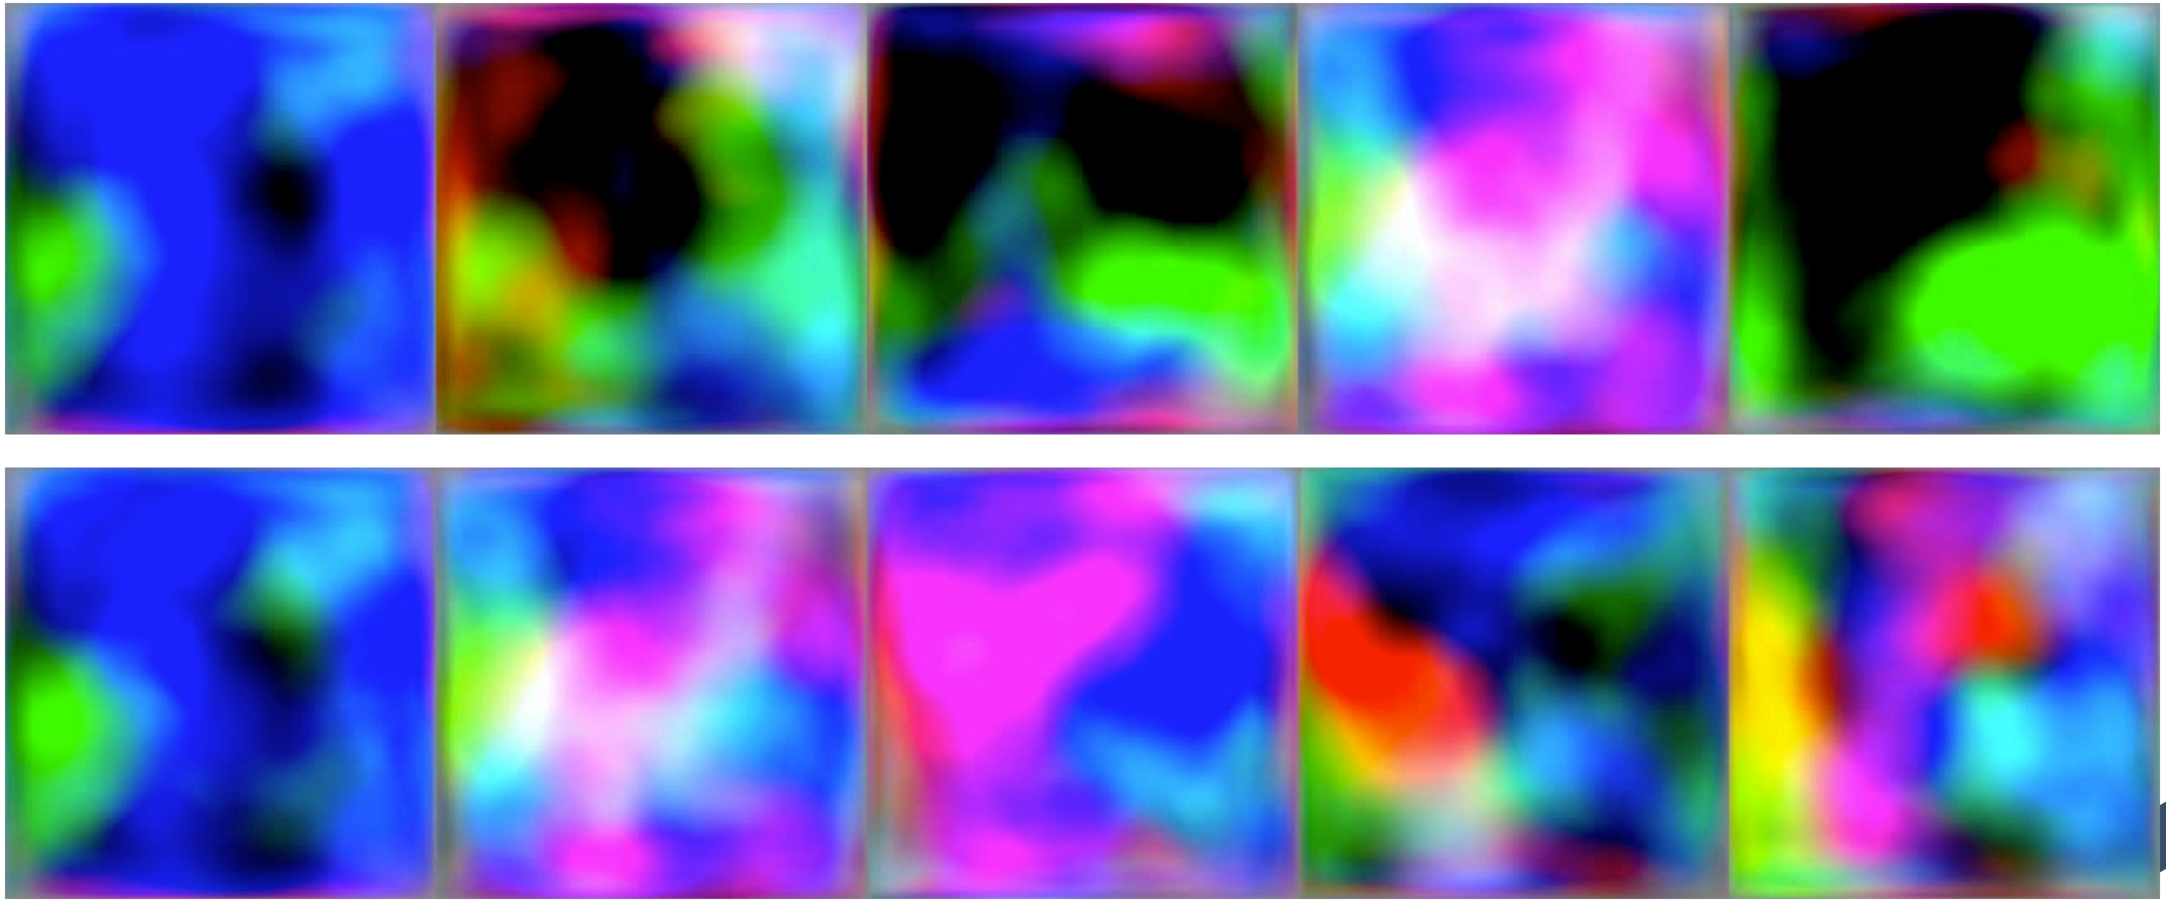

** For more details (e.g., quantitative results, other baselines like BigGAN, ablation studies), please refer to our paper.*

NEURAL INFORMATION PROCESSING SYSTEMS

Deceive D: Adaptive Pseudo Augmentation for GAN Training with Limited Data

Thanks!

Project Page

https://www.mmlab-ntu.com/project/apa/index.html