A MULTI-PLAYER MINIMAX GAME FOR GENERATIVE ADVERSARIAL NETWORKS

Ying Jin, Yunbo Wang, Mingsheng Long, Jianmin Wang, Philip S. Yu, Jiaguang Sun

School of Software, BNRist, Tsinghua University, China
Research Center for Big Data, Tsinghua University, China
Beijing Key Laboratory for Industrial Big Data System and Application

Generative Adversarial Network (GAN)

A framework which has various applications.

Problems: Mode Collapse

Image-to-Image Translation

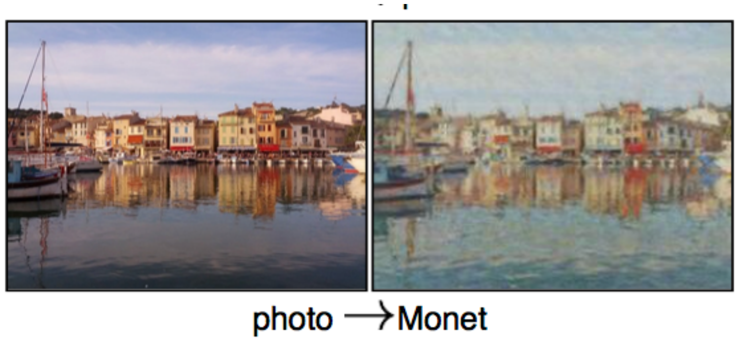

Pixel Prediction

Pathak et al. CVPR16
Zhu et al. NIPS17

GMAN – Generative Multi-Adversarial Networks

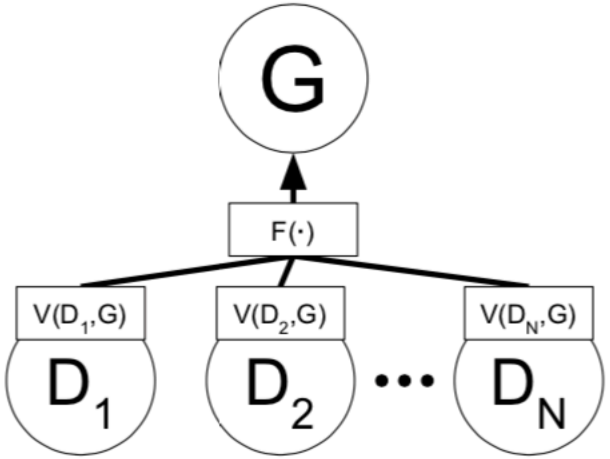

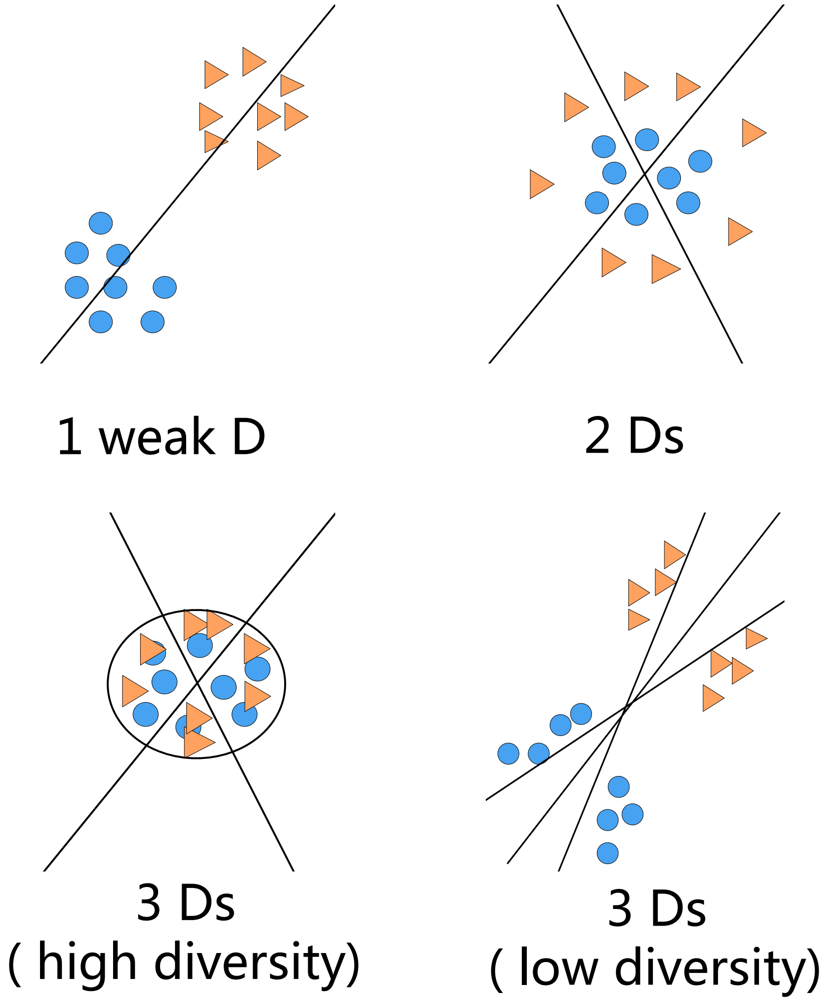

Durugkar et al. ICLR17

DDL – Discriminator Discrepancy Loss

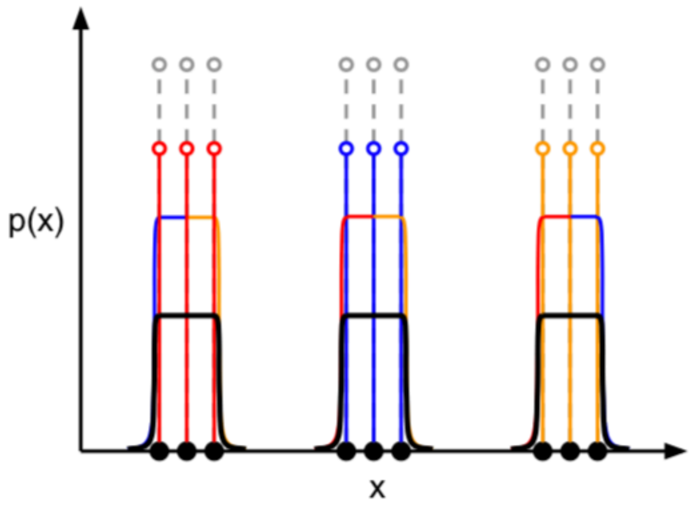

The ideal situation for GMAN: K discriminators excels in separate region

$$L_{\mathrm{DDL}}\left(x;\{D_k\}_{k=1}^{K}\right)=\frac{1}{K}\sum_{k=1}^{K}\left|\phi\left(D_k(x)\right)-\sum_{k'=1}^{K}\frac{\phi\left(D_{k'}(x)\right)}{K}\right|$$

GAN: $\Phi(x) = \log(x)$; WGAN: $\Phi(x) = x$

Larger DDL, More diversity

Durugkar et al. ICLR17

DDL Minimax Game

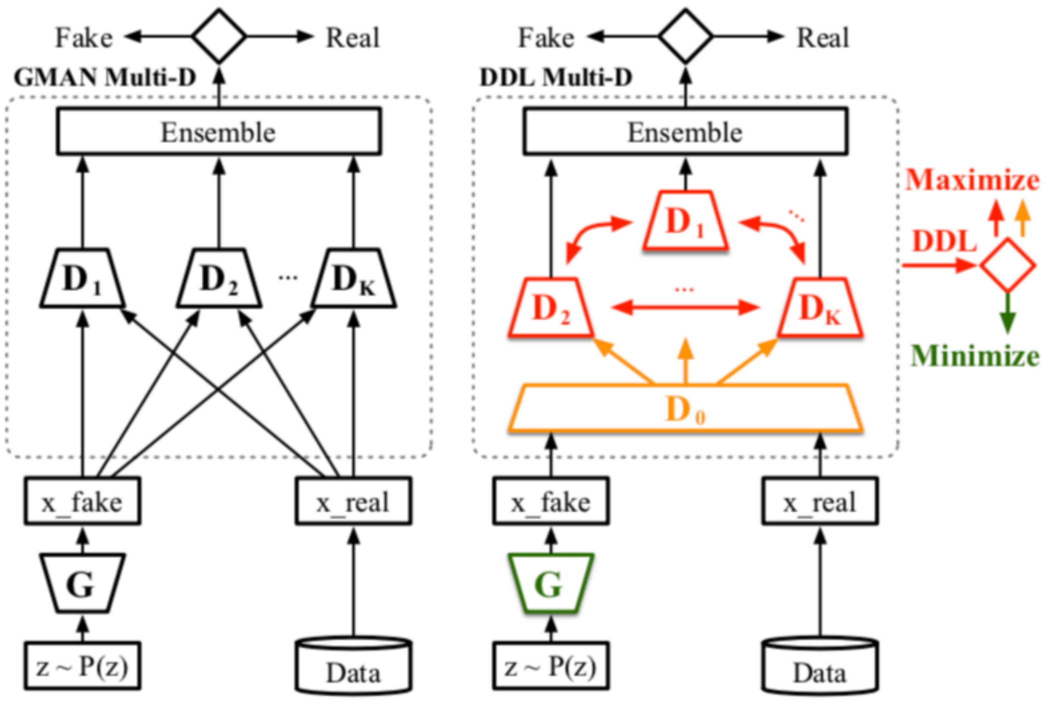

(a) GMAN (b) Our method

$$
\begin{aligned}
L(\theta_G, \{\theta_D^k\}_{k=1}^K) &= \mathbb{E}_{x \sim P_{\text{data}}} \sum_{k=1}^{K} \frac{\phi\left(D_k\left(x\right)\right)}{K} \\
&+ \mathbb{E}_{z \sim P_z} \sum_{k=1}^{K} \frac{\phi\left(1 - D_k\left(G(z)\right)\right)}{K} \\
&+ \lambda \mathbb{E}_{x \sim P_{\text{data}}} L_{\text{DDL}}\left(x; \{D_k\}_{k=1}^K\right) \\
&+ \lambda \mathbb{E}_{z \sim P_z} L_{\text{DDL}}\left(G(z); \{D_k\}_{k=1}^K\right)
\end{aligned}
$$

Layer Sharing

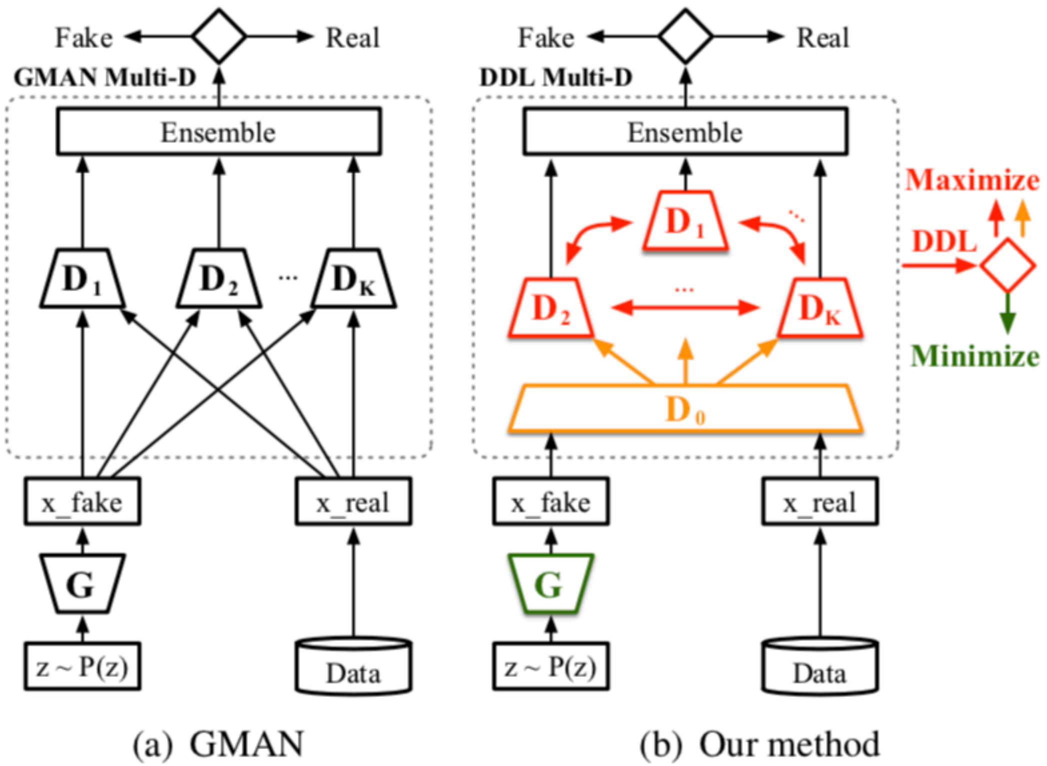

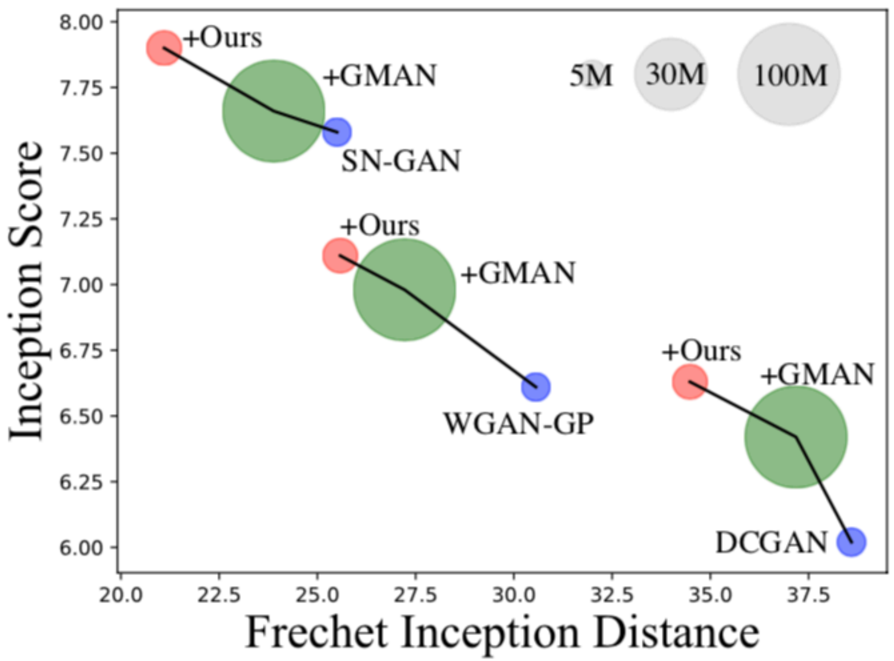

Results – Toy Dataset

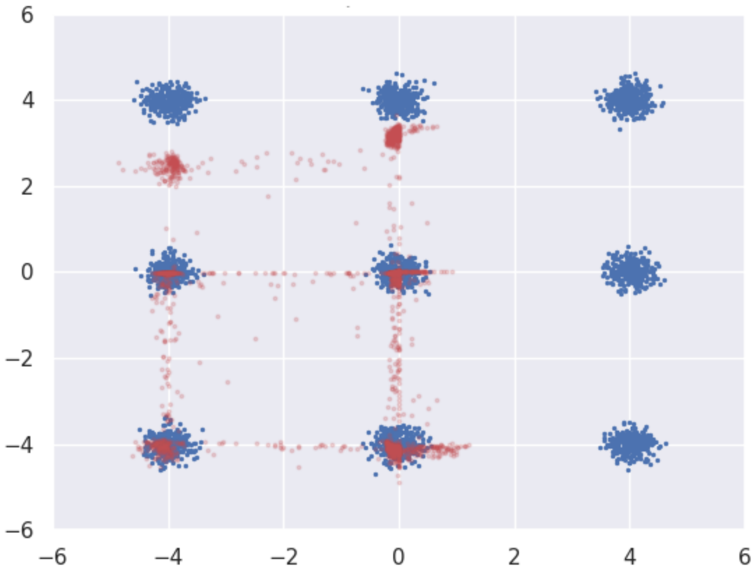

GMAN

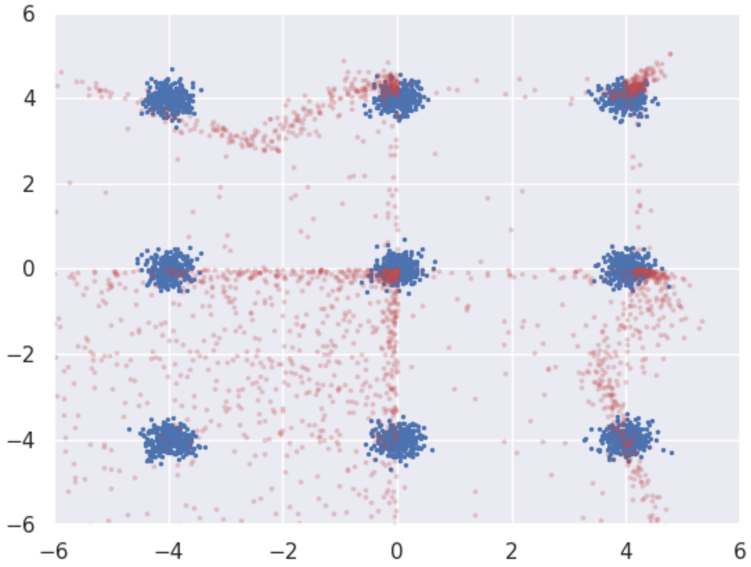

Maximize DDL

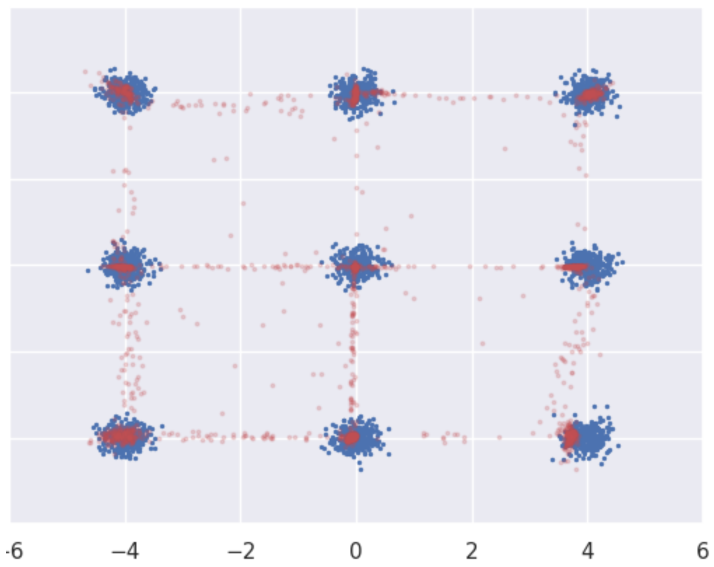

Minimax DDL

Results – Cifar10/STL10

Model	DCGAN	WGAN-GP	SN-GAN
Vanilla	6.02 / 38.59	6.61 / 30.56	7.58 / 25.50
+ GMAN	6.42 / 37.18	6.98 / 27.22	7.66 / 23.89
+ DDL	**6.63 / 34.48**	**7.11 / 25.58**	**7.90 / 21.01**
+ DDL*	6.37 / 35.16	7.04 / 26.14	7.71 / 23.64

Table 1. IS/FID results on CIFAR10. DDL* is a variant of our method without shared layers between discriminators.

Model	WGAN	WGAN-GP	SN-GAN
Vanilla	7.57 / 64.20	8.42 / 55.10	8.79 / 43.20
+ GMAN	7.82 / 54.93	8.72 / 47.26	8.86 / 41.67
+ DDL	**7.92 / 48.05**	**8.94 / 44.80**	**9.21 / 39.68**

Table 2. IS/FID results on the STL-10 dataset.

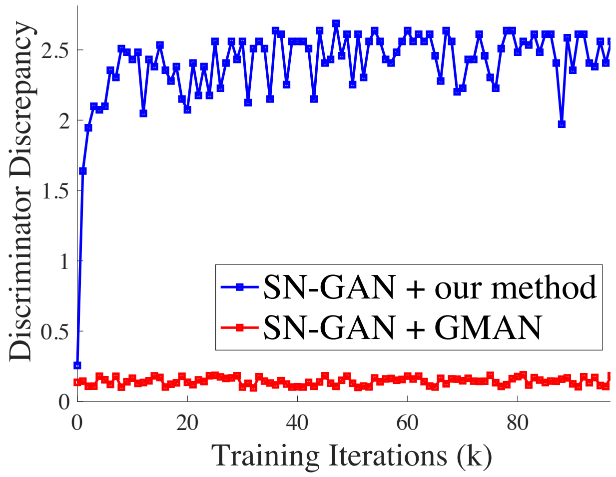

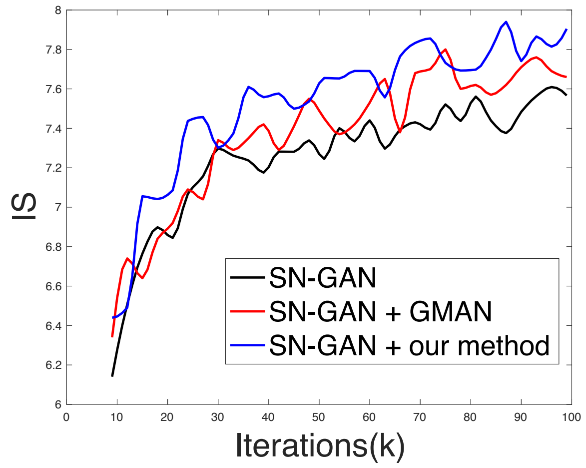

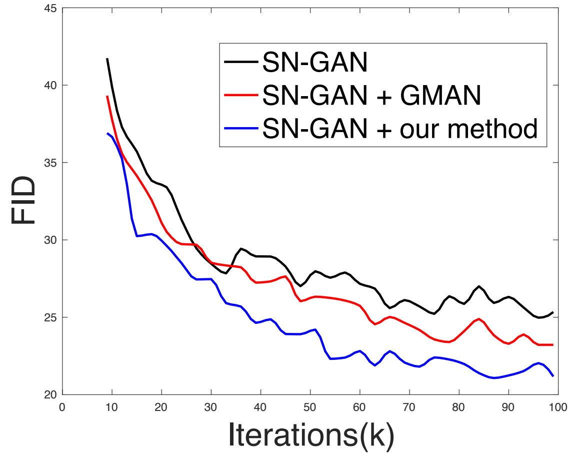

Results – CelebA/ImageNet/LSUN

CelebA	IS / FID	ImageNet	IS / Intra FID
WGAN	1.67 / 45.17	SN-GAN-Proj	36.8 / 92.4
+ GMAN	1.66 / 41.09	+ GMAN	37.6 / 89.5
+ DDL	**1.75 / 39.15**	+ DDL	**39.7 / 83.7**

Table 3. Results of our method on CelebA and ImageNet.

Model	FID	Perceptual Path Length	
		Full	End
StyleGAN	3.324	2419.78	1349.88
+ GMAN	2.862	2378.29	1302.09
+ DDL	**2.606**	**2314.87**	**1282.97**

Table 4. Results of our method on LSUN-Bedroom.

Vanilla

Ours

Results – Ablation Study

K	SN-GAN + GMAN	SN-GAN + DDL
1	7.58 / 25.50	7.58 / 25.50
4	7.60 / 24.17	7.63 / 22.87
8	7.59 / 23.88	7.70 / 22.82
12	7.60 / 23.33	7.62 / 22.66
16	7.66 / 23.89	**7.90 / 21.01**
20	7.63 / 22.57	7.70 / 22.33
32	7.58 / 23.22	7.59 / 22.96

Table 5. The IS/FID results of GMAN and our method with respect to the number of discriminators on CIFAR10 (backbone: SN-GAN). $K = 1$ is equivalent to the vanilla SN-GAN. The best result is achieved at $K = 16$. For all candidates of K, our method consistently outperforms GMAN.

λ	0.0	0.001	0.1	0.3	1.0	2.0
IS	7.58	7.48	7.45	7.63	**7.90**	7.64
FID	25.50	23.04	24.01	23.92	**21.01**	23.53

Table 6. The influence of λ when applying DDL to SN-GAN on CIFAR10. $\lambda = 0$ is equivalent to the vanilla SN-GAN.

Conclusions

1. Discriminator Discrepancy Loss (DDL) to diversify multi-discriminators of GANs.
2. A multi-player minimax game for GANs, where Ds maximize DDL and G minimizes DDL.
3. Layer-sharing architecture for hyperparameter efficiency and collaboration.
4. Orthogonal to existing GANs and consistently outperforms GMAN.

Questions?