Projected GANs Converge Faster

Axel Sauer, Kashyap Chitta, Jens Müller, Andreas Geiger

TL;DR: Training GANs in pretrained feature spaces improves image quality, training speed, and sample efficiency.

The Discriminator's Task is two-fold

In GAN training, the discriminator aims to distinguish real from fake samples. On closer inspection, the discriminator's task is two-fold

- 1. Learn a representation of the input space (projects the real and fake samples into a meaningful space)
- 2. Discriminate based on this representation

We explore the utility of pretrained representations to facilitate the discriminator's task and improve and stabilize GAN training.

Exploiting the Full Potential of **Pretrained Features**

Key components:

- Pretrained Feature Network F
- Feature Pyramids via random projections:
 - Cross-channel mixing (CCM)
 - Cross-scale mixing (CSM)
- Multi-scale discriminators



eberhard karls UNIVERSITÄT TÜBINGEN

Training Properties

Motivation



Results





MAX PLANCK INSTITUTE











STYLEGAN2-AI FASTGAN [35] PROJECTED GAN PROJECTED GAN





										8
	FID	Imgs	FID	Imgs	FID	Imgs	FID	Imgs	FID	Imgs
	Large Datasets (256 ²)									
	CLEVR		FFHQ		Cityscapes		Bedroom		Church	
	26.04	10 M	16.21	10 M	12.81	10 M	14.06	10 M	6.15	10 M
DA [25]	10.17	10 M	7.32	10 M	8.35	10 M	11.53	10 M	5.85	10 M
[20]	9.24	10 M	7.42	10 M	5.23	10 M	6.15	10 M	5.47	10 M
Sold-	3.24	10 M	12.69	10 M	8.78	1.8 M	8.24	4.8 M	8.43	8.9 M
Ν	0.89	4.5 M	3.39	7.1 M	3.41	1.7 M	1.52	5.2 M	1.59	9.2 M
۷*	3.39	0.5 M	3.56	7.0 M	4.60	1.1 M	2.58	1.5 M	3.18	1.1 M
25, 26, 68]	5.05	25 M	3.62	25 M	-	-	2.65	70 M	3.39	88 M
	Small Datasets (256 ²)									
	Art Painting		Landscape		AnimalFace		Flowers		Pokemon	
DA [25]	43.07	3.2 M	15.99	6.3 M	60.90	2.2 M	21.66	3.8 M	40.38	3.4 M
	44.02	0.7 M	16.44	1.8 M	62.11	0.2 M	26.23	0.8 M	81.86	2.5 M
Ν	27.96	0.8 M	6.92	3.5 M	17.88	10 M	13.86	1.8 M	26.36	0.8 M
N*	40.22	0.2 M	14.99	0.6 M	58.07	0.02 M	21.60	0.2 M	36.57	0.3 M