NEURAL INFORMATION PROCESSING SYSTEMS

ILLINOIS

DigGAN: Discriminator gradIent Gap Regularization for GAN Training with Limited Data

Tiantian Fang, Ruoyu Sun, Alex Schwing

GAN training with limited data

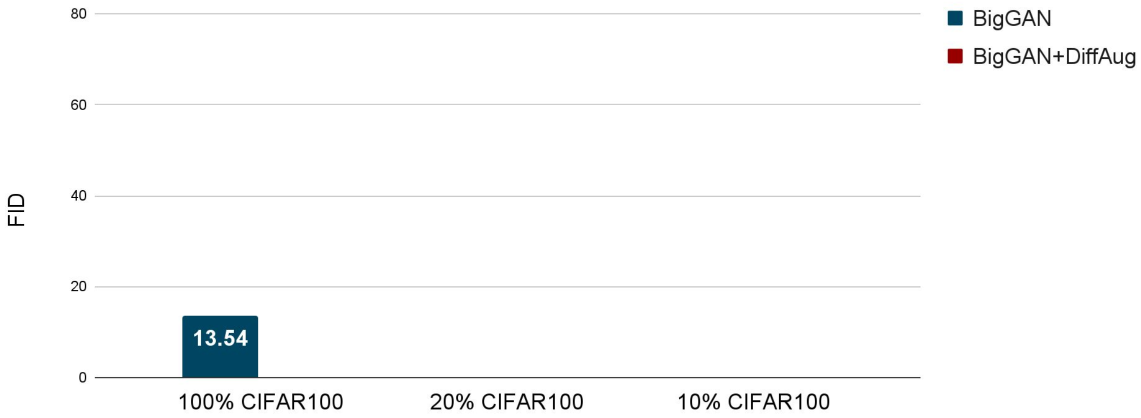

GAN training with limited data

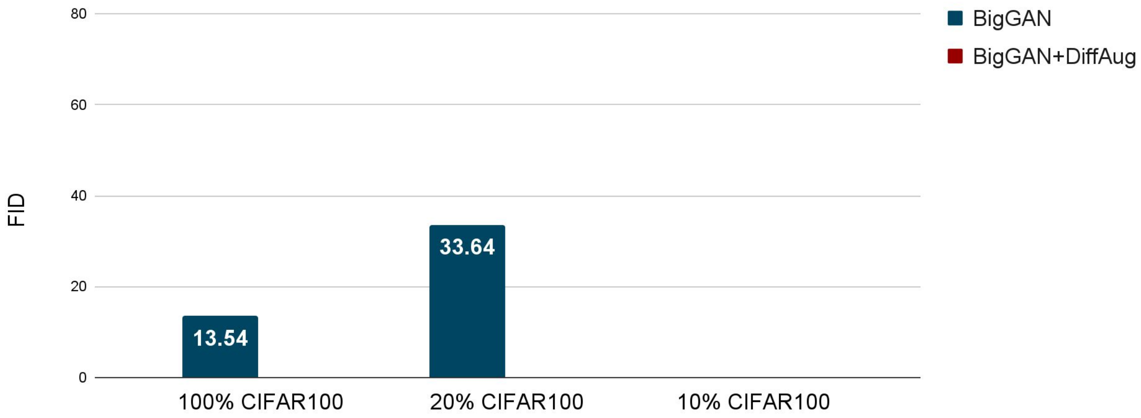

GAN training with limited data

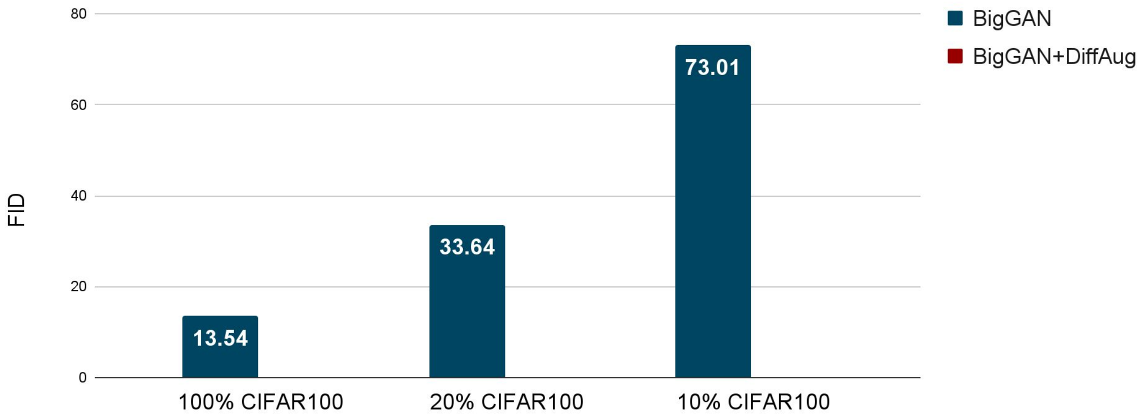

Data augmentation benefit is limited

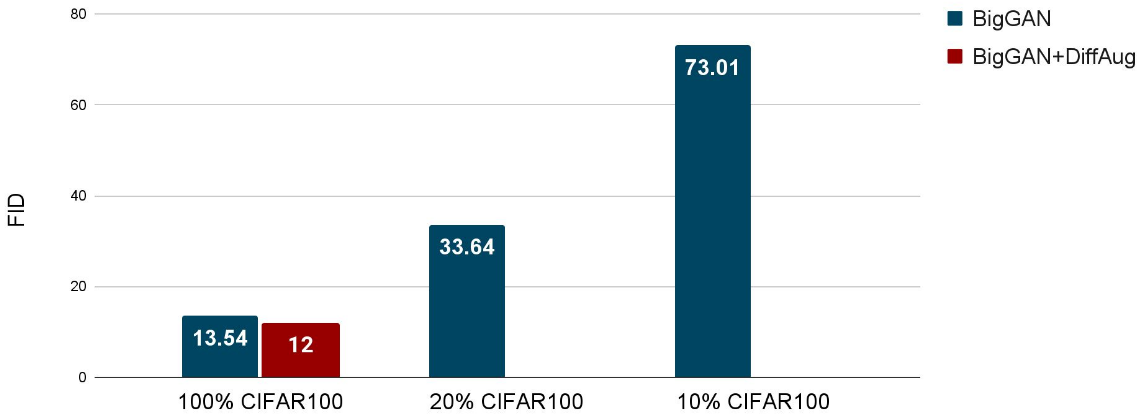

Data augmentation benefit is limited

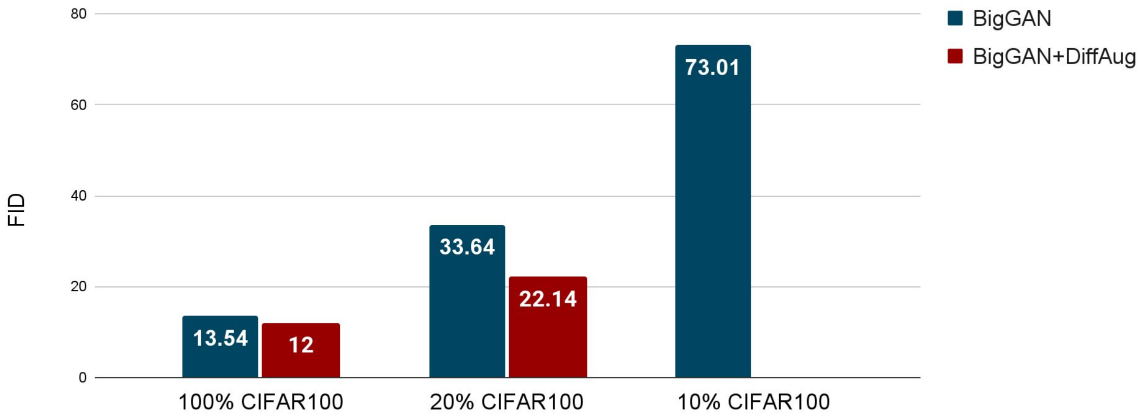

Data augmentation benefit is limited

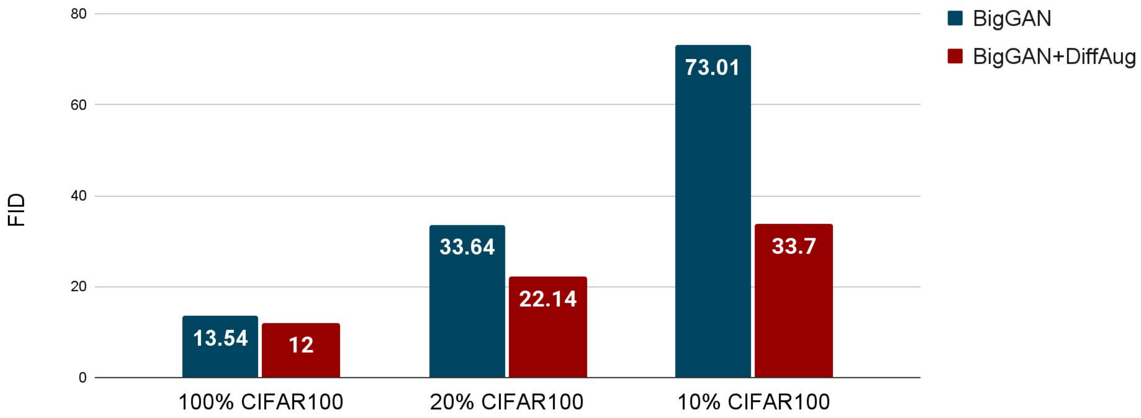

Observation: large gap of gradient norms

Large gap between

- The norm of the gradient of a discriminator's prediction w.r.t. real images
- The norm of the gradient of a discriminator's prediction w.r.t. generated images

$$R(D, x_R, x_F) = \left(\left\| \frac{\partial D}{\partial x_R} \right\|_2 - \left\| \frac{\partial D}{\partial x_F} \right\|_2 \right)^2$$

Observation: large gap of gradient norms

Large gap between

- The norm of the gradient of a discriminator's prediction w.r.t. real images
- The norm of the gradient of a discriminator's prediction w.r.t. generated images

$$R(D, x_R, x_F) = \left(\left\| \frac{\partial D}{\partial x_R} \right\|_2 - \left\| \frac{\partial D}{\partial x_F} \right\|_2 \right)^2$$

BigGAN:

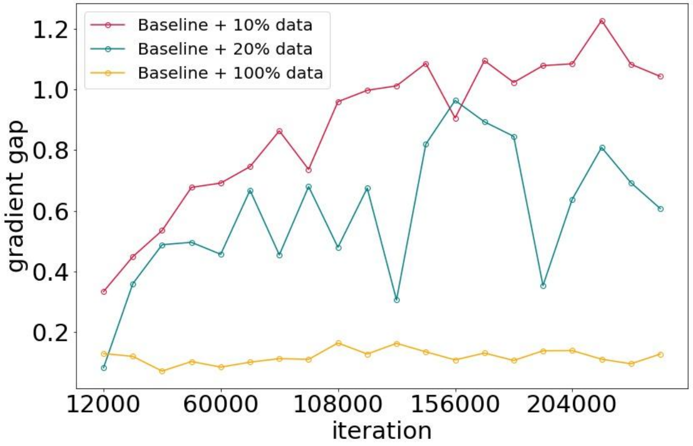

Observation: large gap of gradient norms

Large gap between

- The norm of the gradient of a discriminator's prediction w.r.t. real images
- The norm of the gradient of a discriminator's prediction w.r.t. generated images

$$R(D, x_R, x_F) = \left(\left\| \frac{\partial D}{\partial x_R} \right\|_2 - \left\| \frac{\partial D}{\partial x_F} \right\|_2 \right)^2$$

BigGAN:

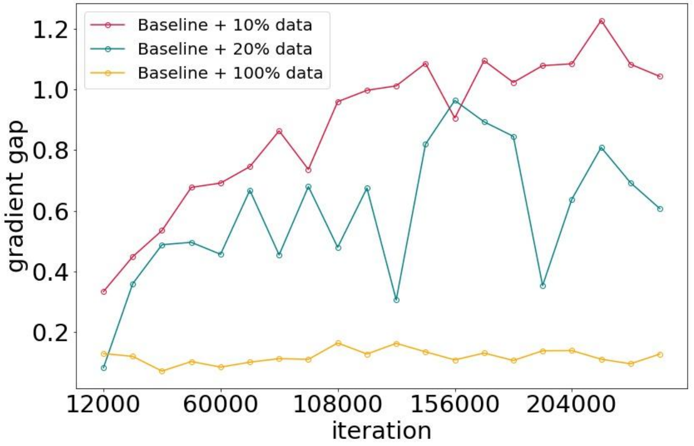

BigGAN + DiffAug:

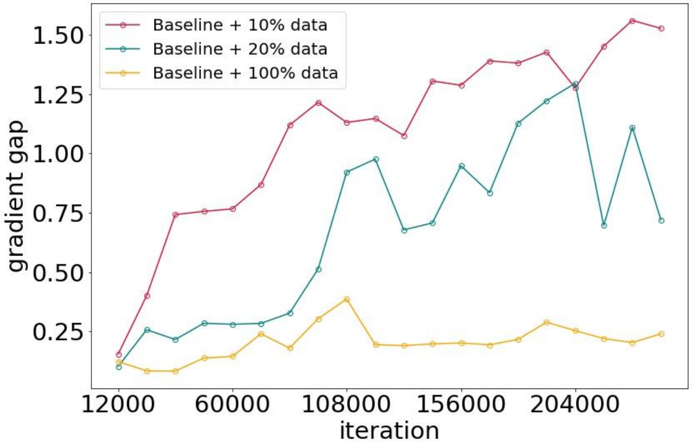

Observation: large gap of gradient norms

Large gap between

- The norm of the gradient of a discriminator's prediction w.r.t. real images
- The norm of the gradient of a discriminator's prediction w.r.t. generated images

$$R(D, x_R, x_F) = \left(\left\| \frac{\partial D}{\partial x_R} \right\|_2 - \left\| \frac{\partial D}{\partial x_F} \right\|_2 \right)^2$$

BigGAN:

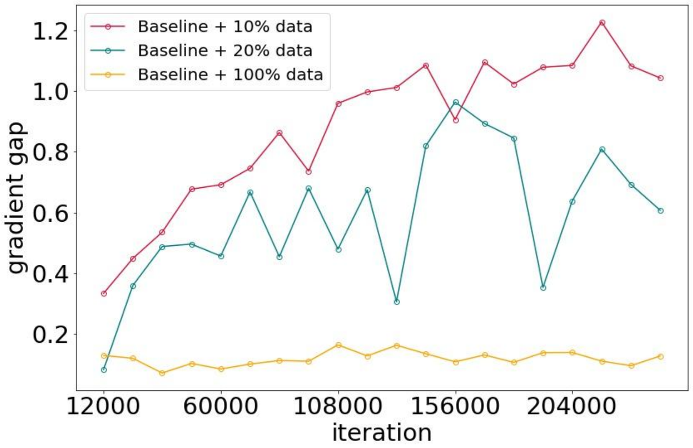

BigGAN + DiffAug:

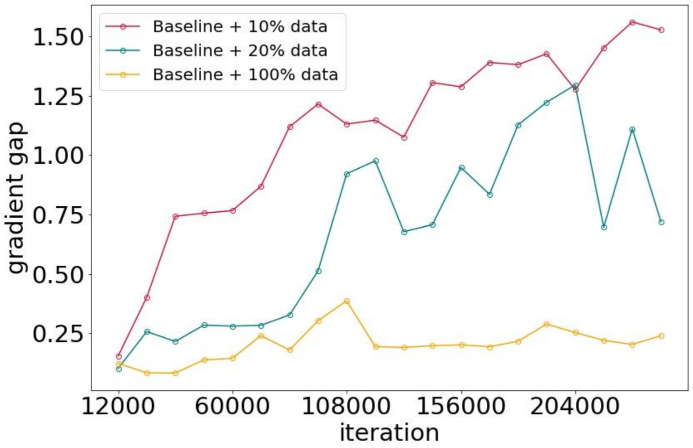

DigGAN: Discriminator gradIent Gap Regularization

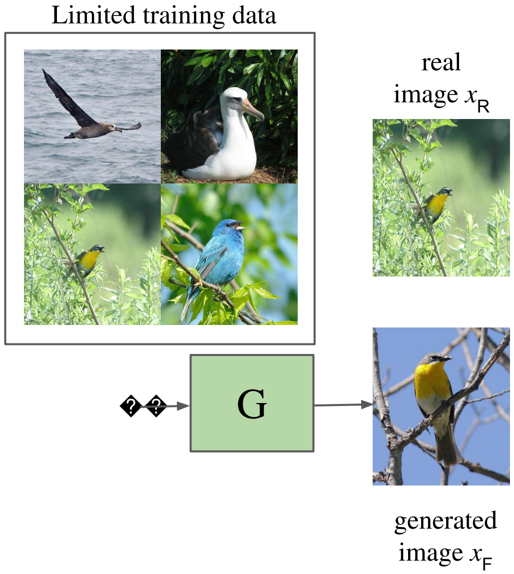

DigGAN: Discriminator gradIent Gap Regularization

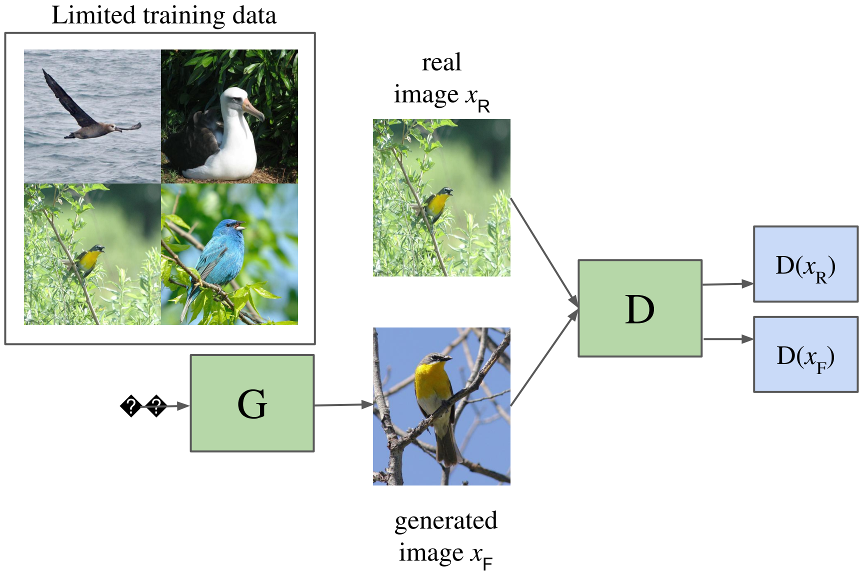

DigGAN: Discriminator gradIent Gap Regularization

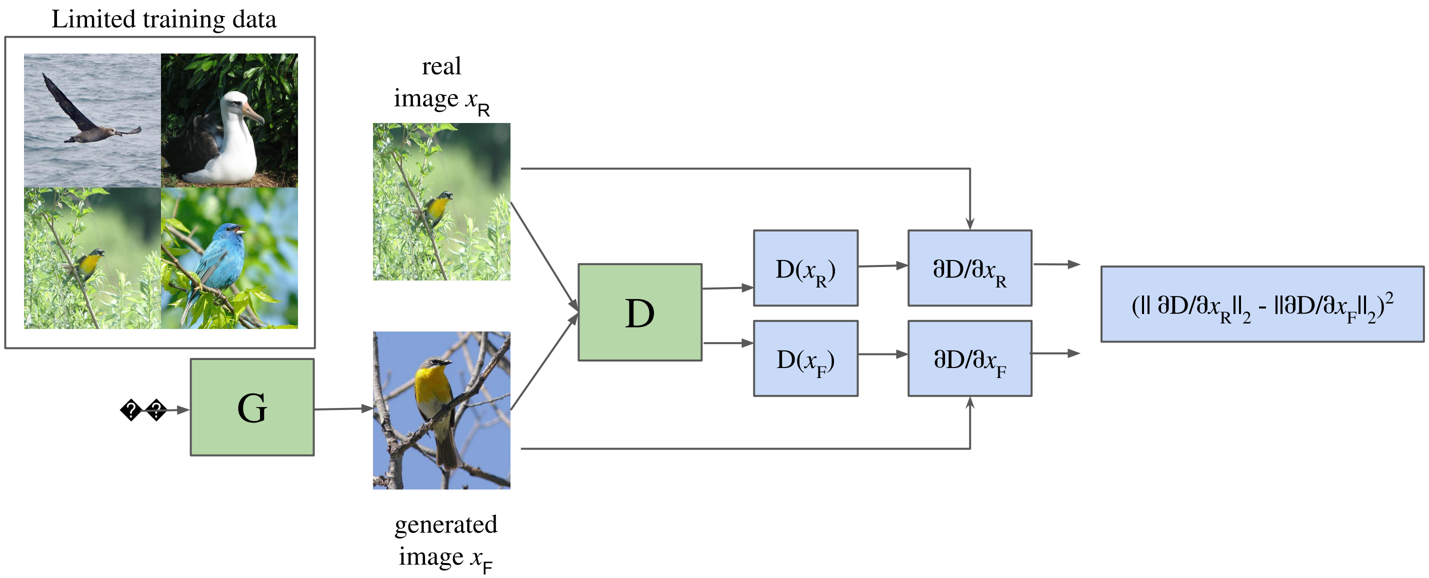

DigGAN: Discriminator gradIent Gap Regularization

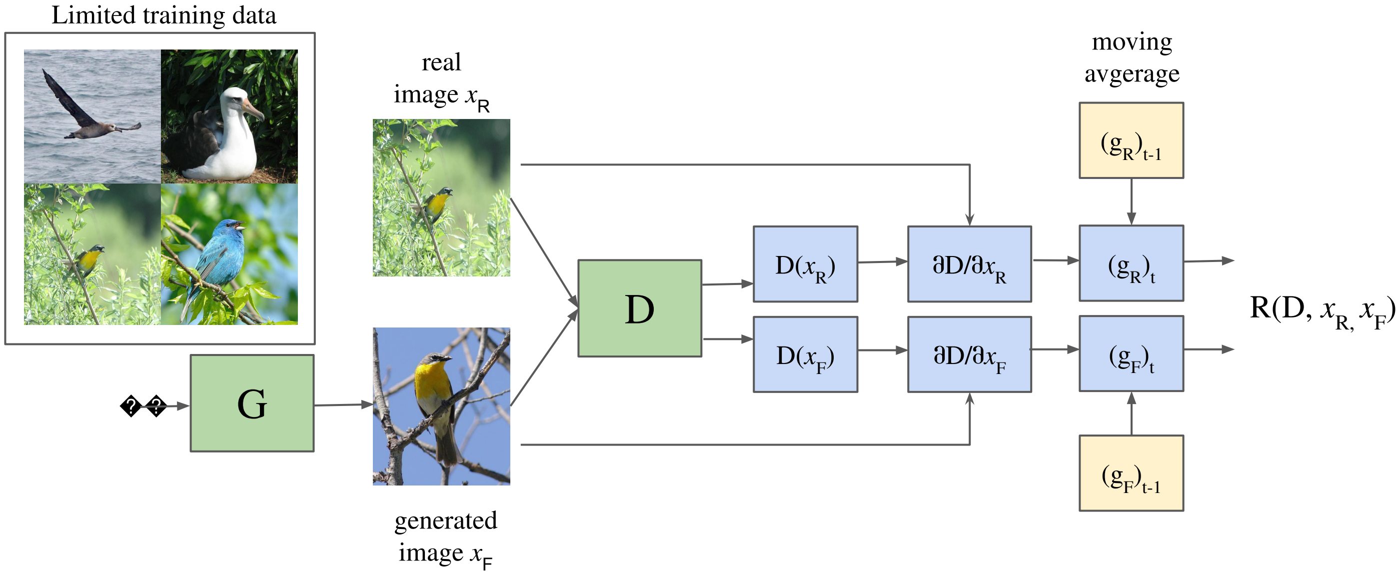

Intuition for DigGAN Regularizer and Attractors

Compared to Vanilla GANs, DigGAN shows empirical advantages:

- Avoids getting trapped in bad local attractors
- Escapes from bad local attractors even if starting at local attractors

Results: CIFAR-100

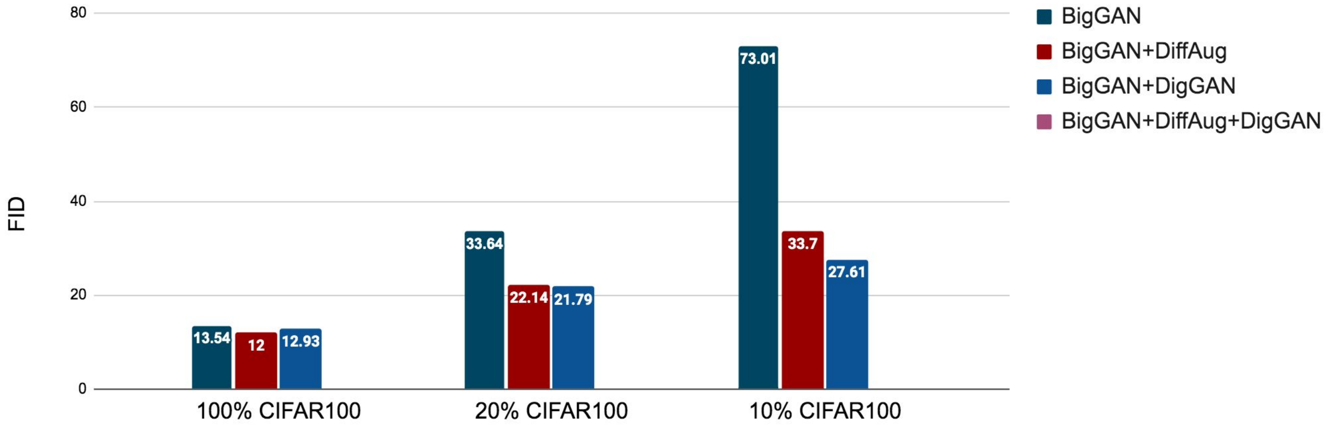

Results: CIFAR-100

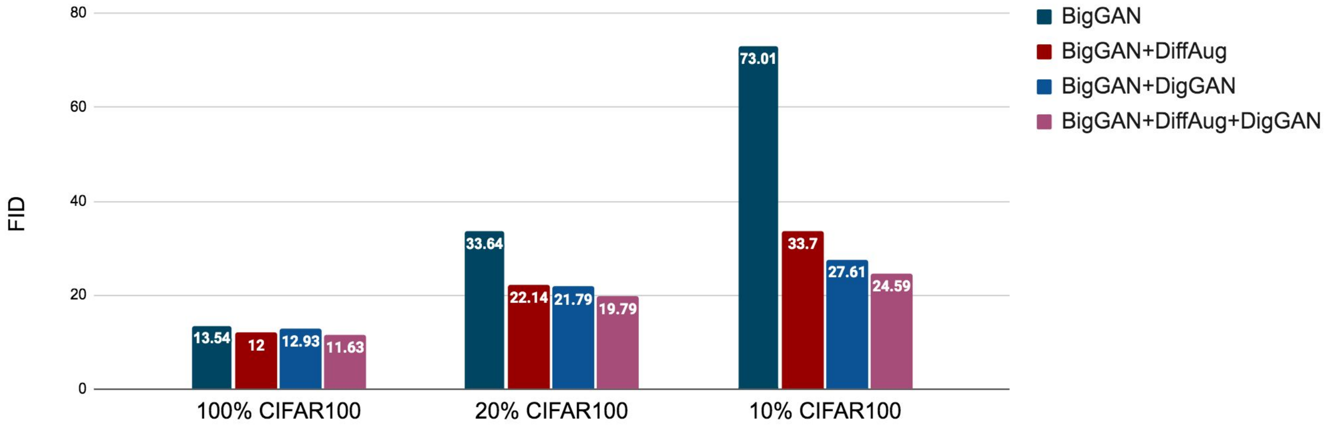

Results: CIFAR-100

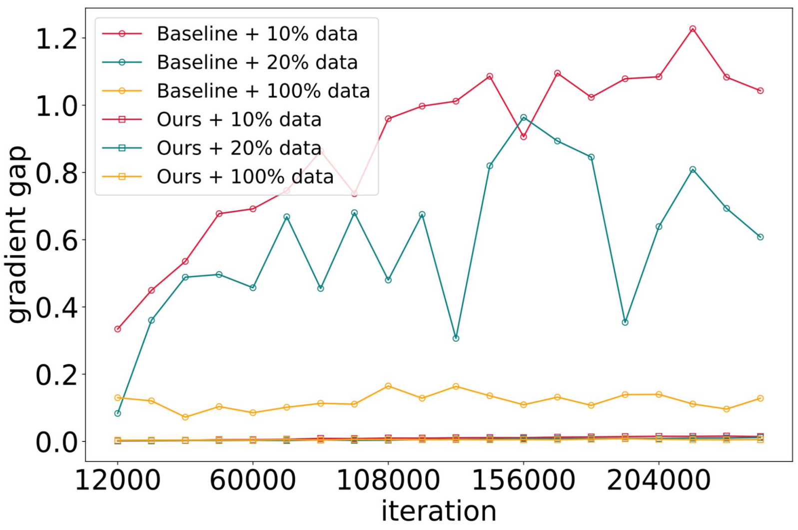

BigGAN

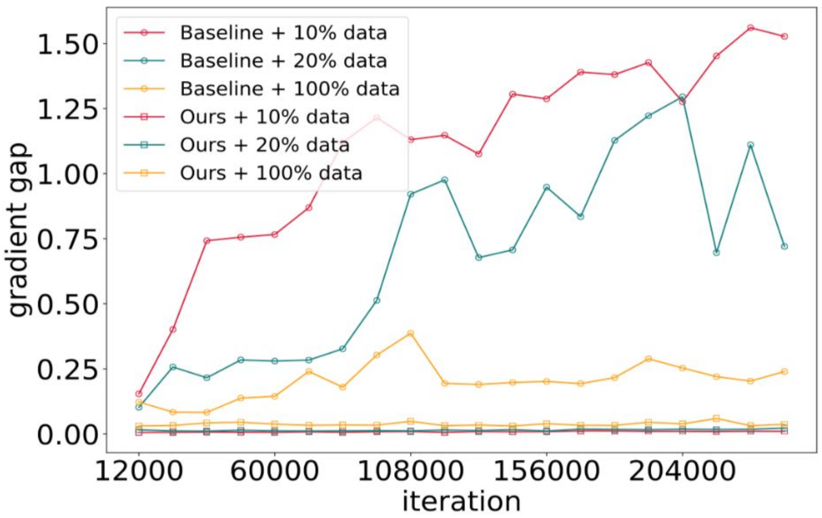

BigGAN + DiffAug

Results

	100% Tiny ImageNet	50% Tiny ImageNet	10% Tiny ImageNet	100% CUB-200	50% CUB-200
BigGAN	31.92	43.45	130.77	20.15	48.67
BigGAN+R_{LC} [51]	28.11	36.11	121.16	40.37	98.38
BigGAN + DigGAN (ours)	**17.76**	**24.63**	**84.27**	**14.45**	**23.20**
BigGAN + DiffAug	16.33	24.50	95.40	13.49	24.35
BigGAN+R_{LC} [51]+DiffAug	16.30	23.67	83.76	12.81	23.49
BigGAN + DiffAug + DigGAN (ours)	**14.84**	**22.66**	**51.18**	**11.58**	**21.12**

Table 3: Fréchet Inception distance (FID) for BigGAN with Tiny-ImageNet and CUB-200.


100% Tiny-ImageNet


50% Tiny-ImageNet


100% CUB200


50% CUB200

Results

	100-shot Obama	100-shot grumpy cat	AnimalFace Dog	AnimalFace Cat
StyleGAN+ADA	49.78	27.34	66.25	41.40
StyleGAN+ADA+DigGAN	**41.34**	**26.75**	**59.00**	**37.61**

Table 4: Fréchet Inception distance (FID) for StyleGAN2 with ADA on low-shot datasets.