

ETH zürich

High-Fidelity Image Generation With Fewer Labels

Michael Tschannen*



Mario Lucic*



Marvin Ritter*



Xiaohua Zhai

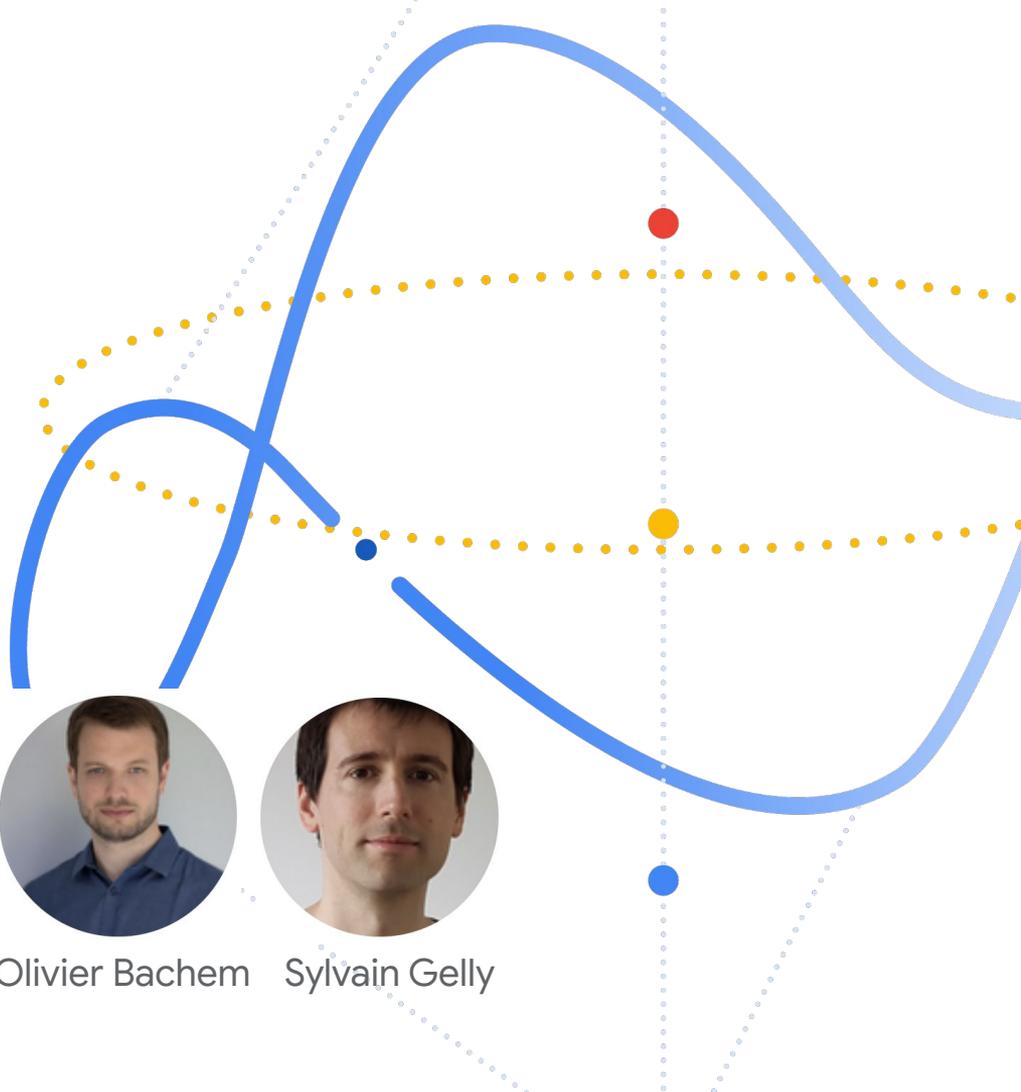


Olivier Bachem

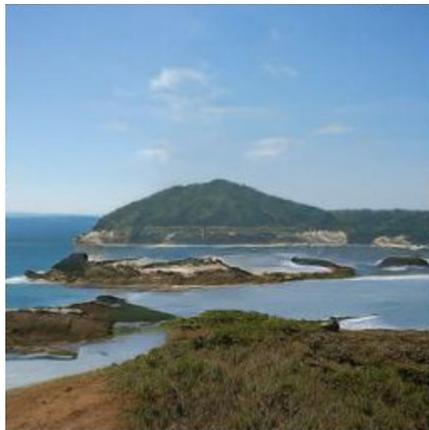


Sylvain Gelly

*equal contribution



Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)

Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)

class-conditional

Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)

class-conditional

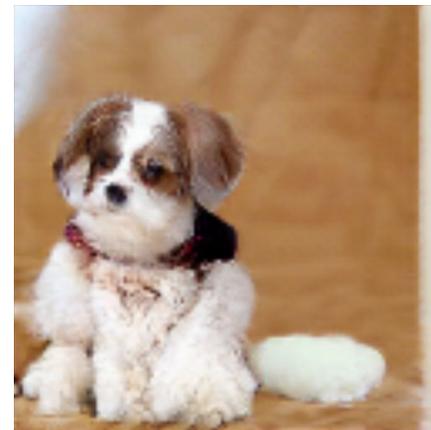
Conditioning reduces the diverse generation problem to a per-class problem

Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)

class-conditional



SS-GAN (Chen et al. 2019)

unsupervised

Conditioning reduces the diverse generation problem to a per-class problem

Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)



class-conditional



SS-GAN (Chen et al. 2019)

unsupervised

Conditioning reduces the diverse generation problem to a per-class problem

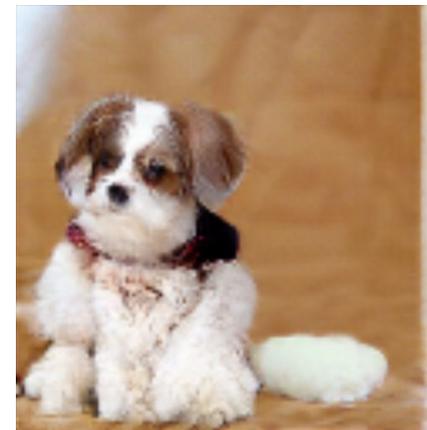
Unsupervised models are considerably less powerful

Generative Adversarial Networks (GANs): Recent Progress



BigGAN (Brock, Donahue, Simonyan 2019)

class-conditional



SS-GAN (Chen et al. 2019)

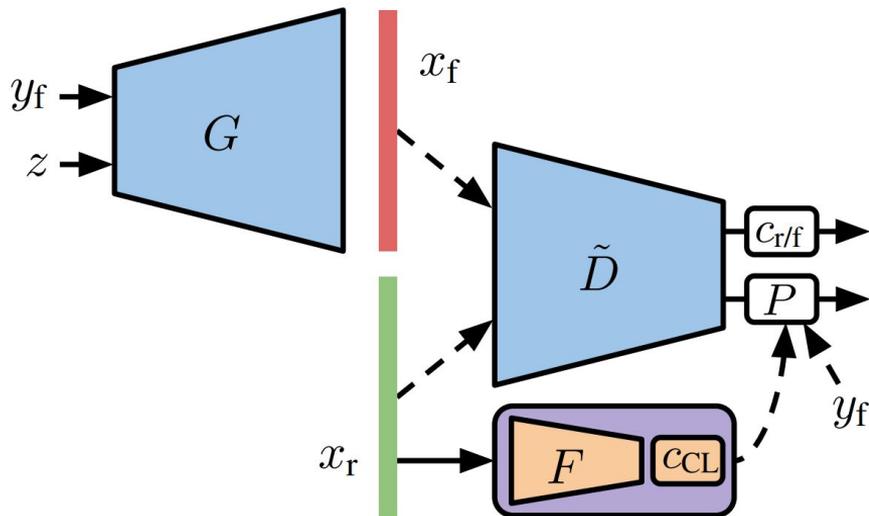
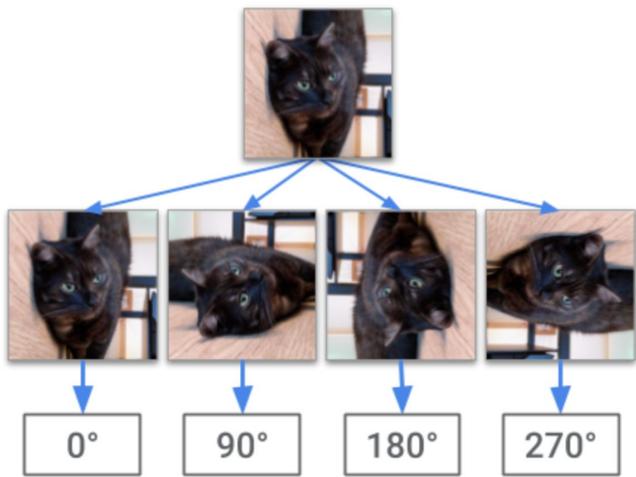
unsupervised

This work: How to close the gap between conditional and unsupervised GANs?

Proposed methods: Overview

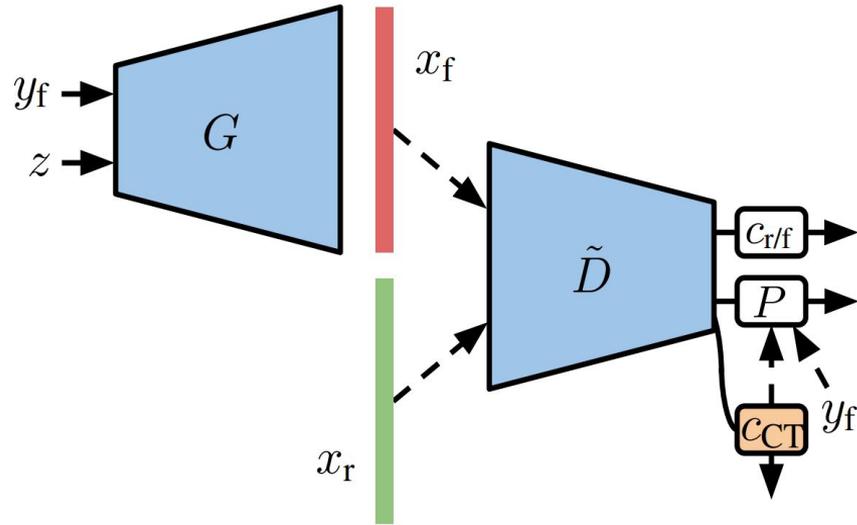
- Replace ground-truth labels with synthetic/inferred labels
→ *No changes in the GAN architecture required*
- Infer labels for the real data using *self-supervised and semi-supervised learning techniques*

Proposed methods: Pre-training



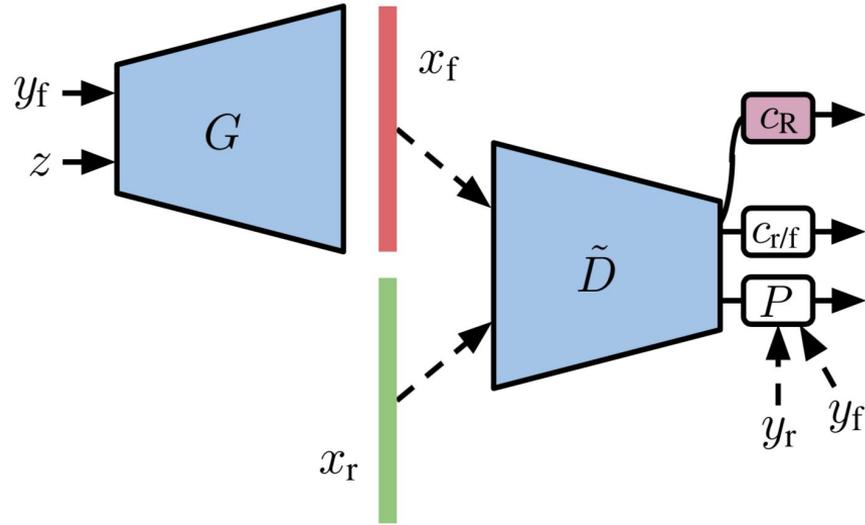
1. Learn a semantic representation F of the data using self-supervision by rotation prediction (Gidaris et al. 2018)
2. Clustering or semi-supervised learning on the representation F
3. Train GAN with inferred labels

Proposed methods: Co-training



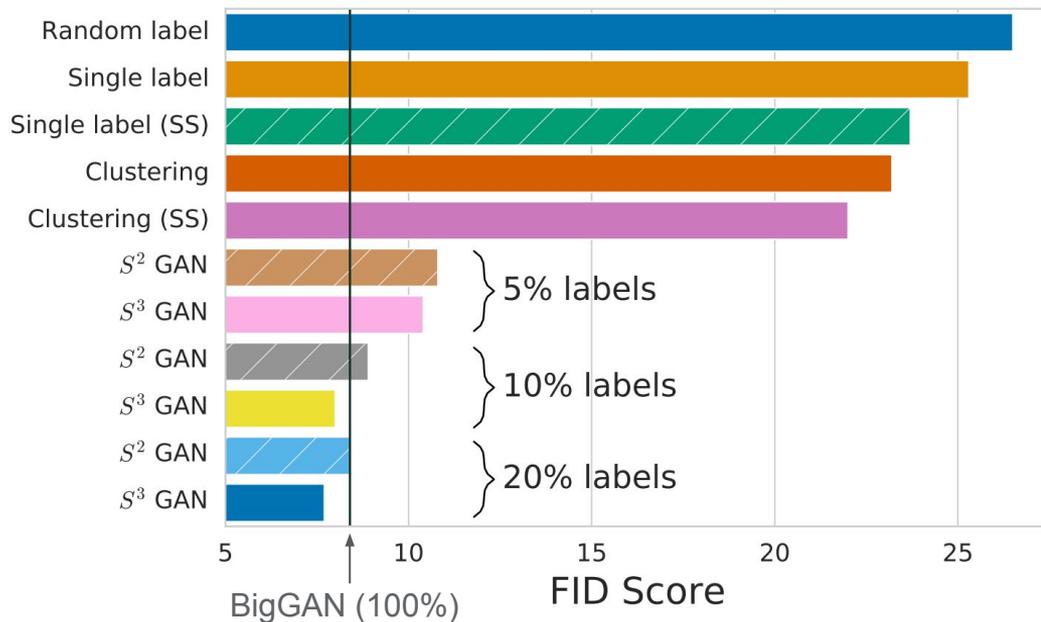
- Semi-supervised classification head on discriminator

Improve pre- and co-training methods



- Rotation-self supervision *during GAN training* (Chen et al. 2019)

Results



- Clustering (SS) is unsupervised **SOTA (FID 22.0)**
- S^2 GAN (20%) and S^3 GAN (10%) match BigGAN (100%)
- S^3 GAN (20%) outperforms BigGAN (100%) (**SOTA**)

Samples: BigGAN (our implementation) vs proposed

BigGAN (100%)



S³GAN (10%)



256 x 256 px

Results



S³GAN (10%)

256 x 256 px

Code, pretrained models and Colabs:



github.com/google/compare_gan

Check out our **poster #13** tonight 6:30-9:00 pm!