Summary of "Freeze the Discriminator: a Simple Baseline for Fine-Tuning GANs"

Abstract

-GANs have the problem of needing lots of time/data to train. Transfer learning tackles this problem, but problem of overfitting still exists

-Paper shows that fine-tuning + freezing lower layers of discriminator "performs surprisingly well"

Introduction

-GANs slow, transfer learning good

-This paper proposes a baseline, FreezeD, which involves fine-tuning + freezing lower layers of discriminator

-"Intuitively, the lower layers of the discriminator learn generic features of images while the upper layers learn to classify whether the image is real or fake based on the extracted features."

Methods

-This paper reviews "previous methods for transfer learning of GANs"

-Fine-tuning, scale/shift, generative latent optimization (GLO), FreezeD (this paper's baseline), L2-SP, Feature distillation

Experiments

-Conducted on both unconditional GANs and conditional GANs

Experiments on unconditional GANs

-Used StyleGAN model pretrained on FFHQ dataset, fine-tuned it on Animal Face (20 classes, about 100 samples each) and Anime Face datasets (first 10 classes, about 100 samples each). Image resolution 256x256. Fine-tuned for 50,000 iterations.

-performed fine-tuning, and freezeD separately. FreezeD had better (lower) FID scores for both Animal Face and Anime Face datasets.

-continue comparing FreezeD against all the other transfer learning methods (scale/shift, GLO, MineGAN, L2-SP, and feature distillation (FD))

-provided a table of FID scores for each of those methods. FreezeD had best FID scores out of all methods. Paper mentions some downsides of scale/shift, L2-SP, GLO, and MineGAN in particular. (downsides = blurry images and poor diversity)

Experiments on conditional GANs

Pre-trained model was fine-tuned on three different datasets that each contained 102, 200, 256 classes, each class having 50-100 samples. Image resolution 128x128. Fine-tuned for 20,000 iterations.

-compared finetuning vs freezeD

--"FreezeD generates more class-consistent samples than fine-tuning"

"FreezeD improves both the performance and stability for most cases, but harms the stability for Oxford Flower." They leave that investigation for future work.

Conclusion

"FreezeD splits the discriminator into a feature extractor and a classifier and then fine-tune the classifier only." And that performs better than the other mentioned methods.

Reference

Mo, S., Cho, M., & Shin, J. (2020). Freeze the discriminator: a simple baseline for fine-tuning gans. *arXiv preprint arXiv:2002.10964*