
TransGAN: Two Pure Transformers Can Make One Strong GAN, and That Can Scale Up

Code used for TransGAN: Two Pure Transformers Can Make One Strong GAN, and That Can Scale Up.

Implementation

- ☐ checkpoint gradient using torch.utils.checkpoint
- ☐ 16bit precision training
- ☒ Distributed Training (Faster!)
- ☒ IS/FID Evaluation
- ☒ Gradient Accumulation
- ☒ Stronger Data Augmentation
- ☒ Self-Modulation

Guidance

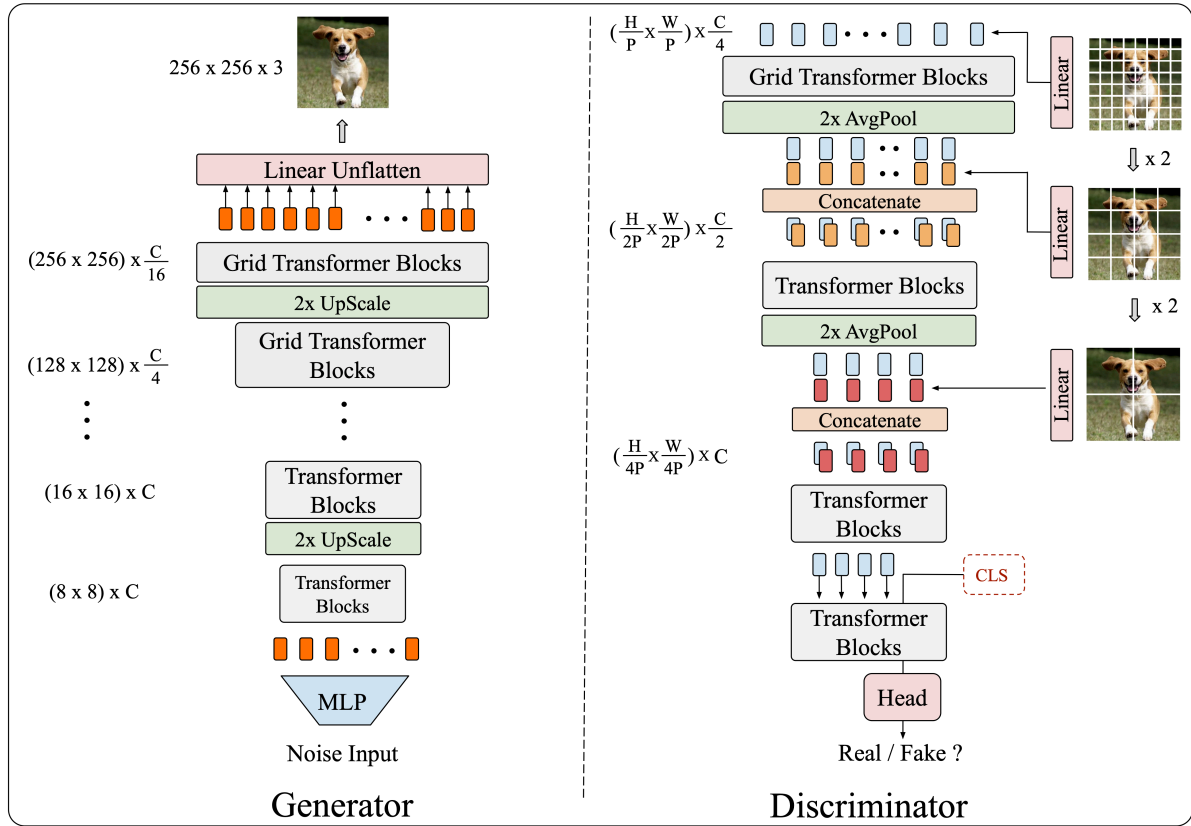
Cifar training script

```
1 python exp/cifar_train.py
```

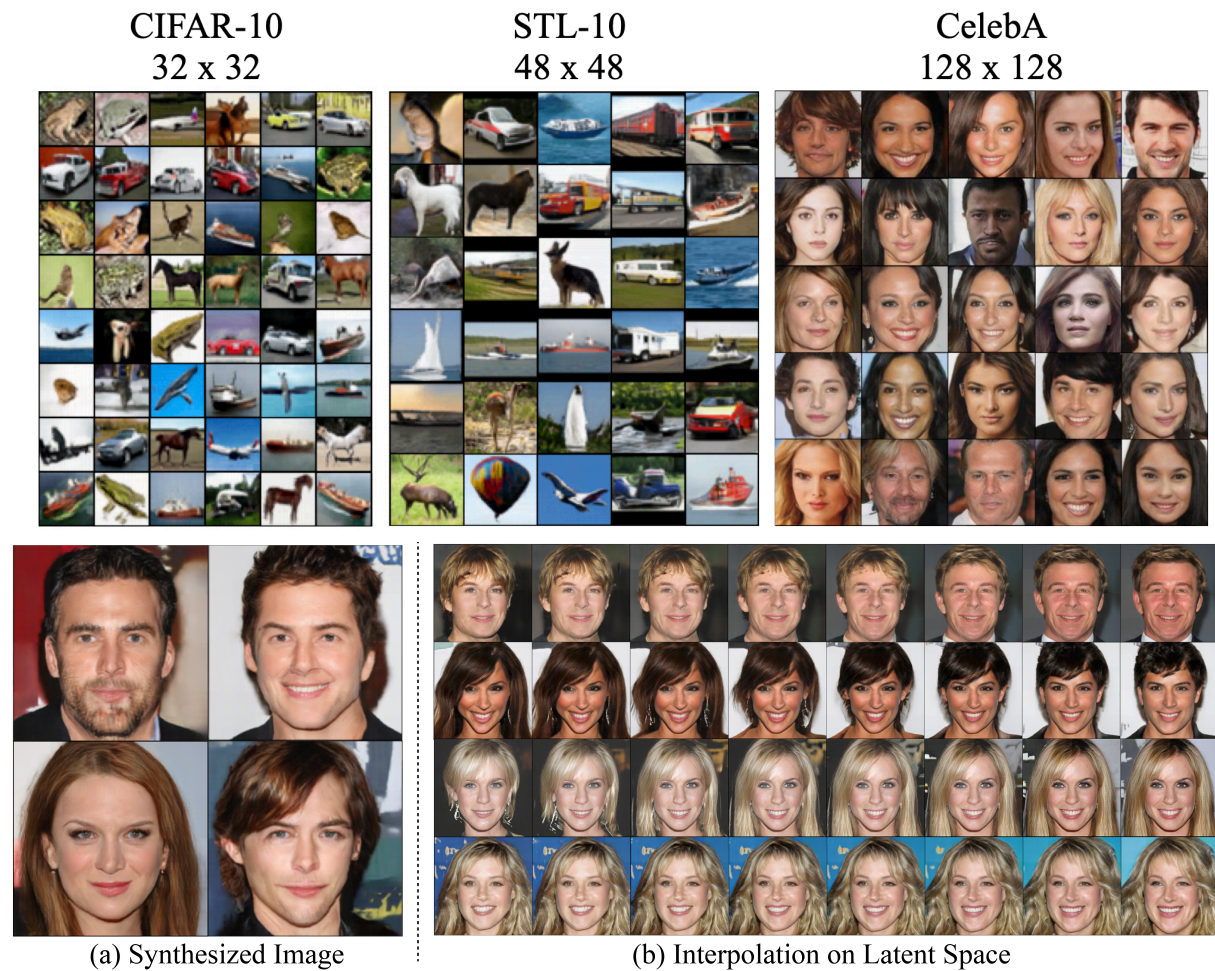
I disabled the evaluation during training job as it causes strange bug. Please launch another evaluation job simultaneously by copying the [path](#) to test script. ##### Cifar test First download the cifar checkpoint and put it on `./cifar_checkpoint`. Then run the following script.

```
1 python exp/cifar_test.py
```

Main Pipeline



Representative Visual Results



README waits for updated ## Acknowledgement Codebase from AutoGAN, pytorch-image-models

Citation

if you find this repo is helpful, please cite

```
1 @article{jiang2021transgan,  
2   title={Transgan: Two pure transformers can make one strong gan, and  
   that can scale up},  
3   author={Jiang, Yifan and Chang, Shiyu and Wang, Zhangyang},  
4   journal={Advances in Neural Information Processing Systems},  
5   volume={34},  
6   year={2021}  
7 }
```