Progressive Growing of GANs for Improved Quality, Stability, and Variation

Paper: T.Karras, T.Aila, S.Laine, J. Lehtinen
Nvidia and Aalto Univ.

ICLR Oral 2018

Stefano Blumberg, UCL

Why I chose Paper

- Excellent Results
- Introduced New Pipeline
- Idea for a current project
- ‘Traning paper’ many applications
- Made me feel poor

Everybody Knows about GANs?

Learning Across Different Scales

- Training at high-resolution too difficult

Fading in Layers

Increasing Variation

- Prevent Mode Collapse
- Minibatch Discrimination (Salimans et al.), feature statistics across minibatch – encourage statistics across different training, generated images.

Minibatch Standard Deviation

- Compute std for each feature in each spatial location
- Average estimates over all features
- Replicate value and concatenate to all spatial locations
- Add value as a feature map towards end discriminator
- (More complicated methods didn’t improve results)

Equalised Learning Rate

- Ignore complex weight initialisation, scale weights at runtime
- Parameter dynamic range and learning speed is same for all weights

Pixelwise Feature Vector Normalisation

- Prevent Magnitudes of generator and Discriminator from spiralling out of control
- Normalise feature vector in each pixel to unit length
- Replaces Batch,Layer et.c-Norm

Multi Scale Statistical Similarity for Assessing GAN Results

- Current methods fail to react to variation in colour or textures, or assess image quality
- Intuition: Samples produced have local image structure similar to training set in all scales
- Statistical similarity from Laplacian Pyramid (specific spatial frequency band)
- Then Wasserstein distance

CelebA-HQ

- High quality version of CelebA
- 30000 images 1024**2

Contributions

- Learning across Different Scales
- Fading in Layers
- Increasing Variation
- Equalised Learning Rate
- Multi Scale Statistical Similarity for Asessing GANs
- CelebA-HQ

Loss and Evaluation

- Design choices orthogonal to loss function chosen (LSGAN,WGAN-GP)
- Sliced Wasserstein Distance and Multi-Scale Structural Similarity

Training

- 8 Tesla V100 (10-11K$ each), 4 days
- Reduce MB size to preserve memory

Results

- Progressive Growing 2-5.4x speedup and better minima
- LSUN bedroom
- Record inception score of 8.80 in unsupervised CIFAR10

Takeaway

- Better hardware is nice
- Progressive growing is very good
- Being ‘very hacky’ does produce advantages