A Review of Recent Improvements of Generative Adversarial Networks

Progressive Growing of GANs

Benjamin Wilhelm

February 25, 2019

University of Konstanz

Figure 1: . . .

Figure 1: People imagined by a numpy random number generator[1].

[1] Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

Table of contents

Introduction

Generative Adversarial Networks

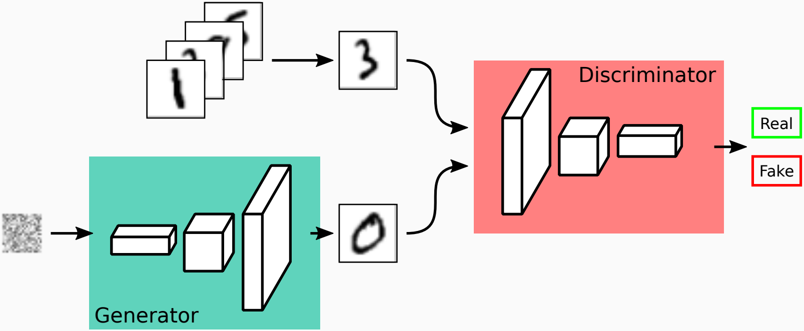

Figure 2: GAN framework[2]

[2] Image reference by Thalles Silva
https://medium.freecodecamp.org/an-intuitive-introduction-to-generative-adversarial-networks-gans-7a2264a81394

Figure 3: Results from the original GAN paper[3].

[3] Goodfellow et al., "Generative adversarial nets".

Problems and Solutions

Variation of Generated Images

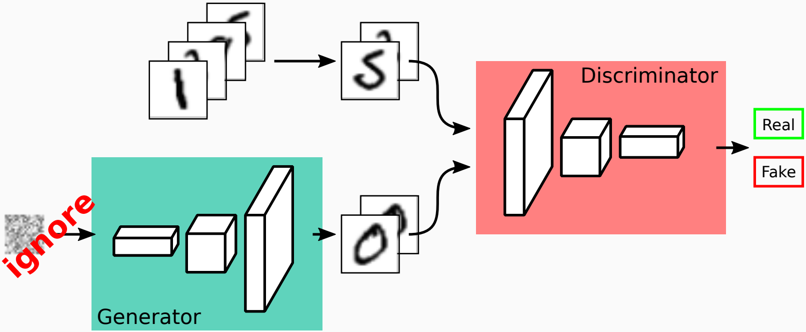

Figure 4: The latent vector can be ignored and only one good image generated.

Variation of Generated Images

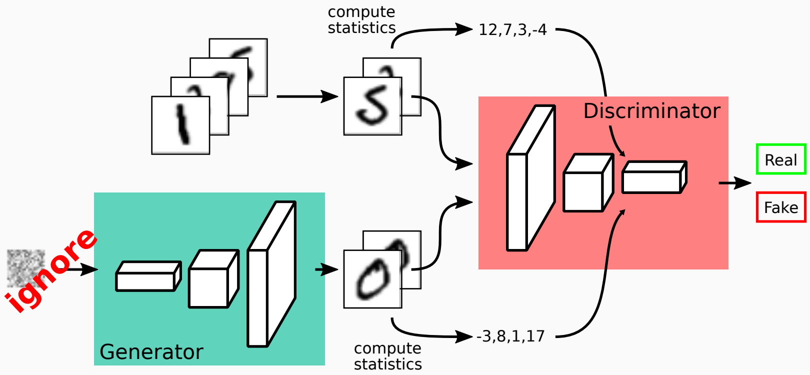

Figure 5: Compute statistics over minibatch and make them available to the discriminator.

High-Resolution Images

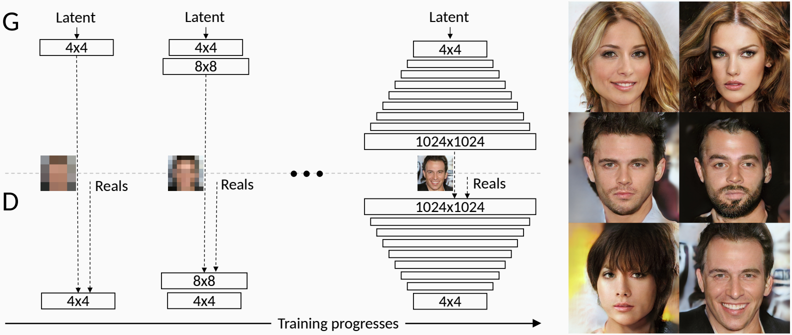

Figure 6: Progressive Growing[4].

[4] Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

High-Resolution Images

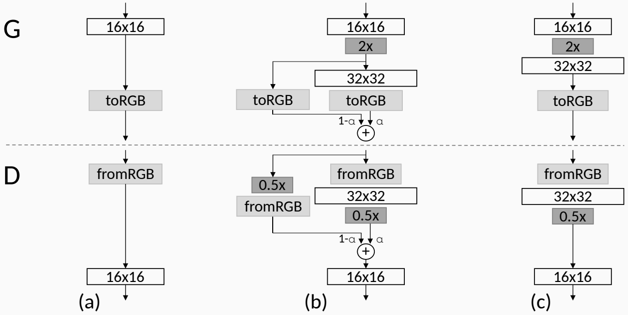

Figure 7: Fading in of new layers[5].

[5] Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

Figure 8: Progressive Growing results[6].

[6] Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

Assessing Results

How good are the generated images?

Option 1: Inception Score (IS)[7]

- If an Inception network gives a good prediction the images are realistic
- If it predicts all classes the variation is good

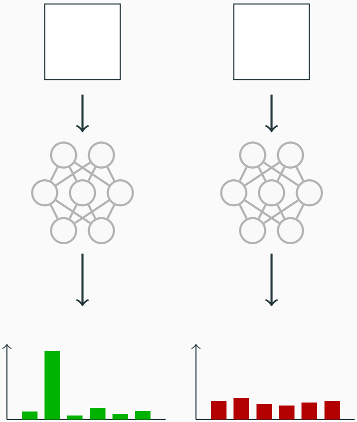

[7] Salimans et al., "Improved techniques for training gans".

Option 1: Inception Score (IS)[7]

- If an Inception network gives a good prediction the images are realistic
- If it predicts all classes the variation is good

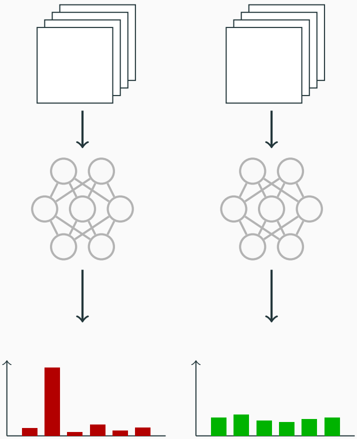

[7] Salimans et al., "Improved techniques for training gans".

Option 2: Fréchet Inception Distance (FID)[8]

- Improvement of the IS
- Use features by an intermediate layer
- Compute μ and Σ for real images and generated images
- Compares statistics

[8] Heusel et al., "Gans trained by a two time-scale update rule converge to a local nash equilibrium".

Option 3: Sliced Wasserstein Distance[9]

- Compute laplacian pyramids of real images and generated images
- Sample image patches on one level
- Compute the sliced Wasserstein distance between the patches

[9] Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

Evaluation

Original Results from the Paper

Figure 9: Progressive Growing results[10].

[10]Karras et al., "Progressive Growing of GANs for Improved Quality, Stability, and Variation".

Generated Images

Figure 10: Images generated with the source code and weights from the paper[12].

[12] https://github.com/tkarras/progressive_growing_of_gans

Generated Images

Figure 11: Images generated with the source code and weights from the paper.

Figure 12: Generated images from the Style-Based GANs paper[13].

[13] Karras, Laine, and Aila, "A Style-Based Generator Architecture for Generative Adversarial Networks".

Style-Based GANs

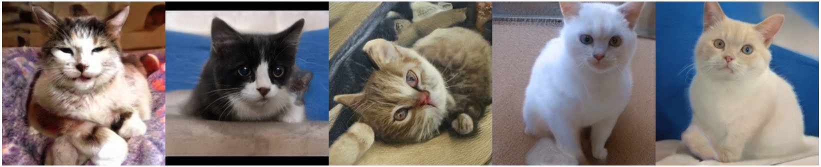

Figure 13: Generated images from the Style-Based GANs paper.

Figure 14: Images generated with the source code and weights from the Style-Based GANs paper[15].

[15] http://stylegan.xyz/code

Style-Based GANs

Figure 15: Images generated with the source code and weights from the Style-Based GANs paper.

Conclusion

Summary

- General idea of GANs
- Problems and Solutions
    - Variation
    - High-resolution
    - Assessing of results
- Generated images

Questions?

References

- Goodfellow, Ian et al. "Generative adversarial nets". In: *Advances in neural information processing systems*. 2014, pp. 2672–2680.
- Heusel, Martin et al. "Gans trained by a two time-scale update rule converge to a local nash equilibrium". In: *Advances in Neural Information Processing Systems*. 2017, pp. 6626–6637.
- Karras, Tero, Samuli Laine, and Timo Aila. "A Style-Based Generator Architecture for Generative Adversarial Networks". In: *arXiv preprint arXiv:1812.04948* (2018).

Karras, Tero et al. "Progressive Growing of GANs for Improved Quality, Stability, and Variation". In: *International Conference on Learning Representations*. 2018. URL: https://openreview.net/forum?id=Hk99zCeAb.

Salimans, Tim et al. "Improved techniques for training gans". In: *Advances in Neural Information Processing Systems*. 2016, pp. 2234–2242.

Videos

Progressive Growing of GANs: `https://youtu.be/G06dEcZ-QTg`

Style-Based GANs: `https://youtu.be/kSLJriaOumA`