Improved Precision and Recall Metric for Assessing Generative Models

Tuomas Kynkäänniemi* Aalto University NVIDIA tuomas.kynkaanniemi@aalto.fi **Tero Karras** NVIDIA tkarras@nvidia.com Samuli Laine NVIDIA slaine@nvidia.com

Jaakko Lehtinen Aalto University NVIDIA jlehtinen@nvidia.com Timo Aila NVIDIA taila@nvidia.com

Abstract

The ability to automatically estimate the quality and coverage of the samples produced by a generative model is a vital requirement for driving algorithm research. We present an evaluation metric that can separately and reliably measure both of these aspects in image generation tasks by forming explicit, non-parametric representations of the manifolds of real and generated data. We demonstrate the effectiveness of our metric in StyleGAN and BigGAN by providing several illustrative examples where existing metrics yield uninformative or contradictory results. Furthermore, we analyze multiple design variants of StyleGAN to better understand the relationships between the model architecture, training methods, and the properties of the resulting sample distribution. In the process, we identify new variants that improve the state-of-the-art. We also perform the first principled analysis of truncation methods and identify an improved method. Finally, we extend our metric to estimate the perceptual quality of individual samples, and use this to study latent space interpolations.

1 Introduction

The goal of generative methods is to learn the manifold of the training data so that we can subsequently generate novel samples that are indistinguishable from the training set. While the quality of results from generative adversarial networks (GAN) [7], variational autoencoders (VAE) [14], autoregressive models [29, 30], and likelihood-based models [6, 13] have seen rapid improvement recently [11, 8, 28, 20, 4, 12], the automatic evaluation of these results continues to be challenging.

When modeling a complex manifold for sampling purposes, two separate goals emerge: individual samples drawn from the model should be faithful to the examples (they should be of "high quality"), and their variation should match that observed in the training set. The most widely used metrics, such as Fréchet Inception Distance (FID) [9], Inception Score (IS) [25], and Kernel Inception Distance (KID) [2], group these two aspects to a single value without a clear tradeoff. We illustrate by examples that this makes diagnosis of model performance difficult. For instance, it is interesting that while recent state-of-the-art generative methods [4, 13, 12] claim to optimize FID, in the end the (uncurated) results are almost always produced using another model that explicitly sacrifices variation, and often FID, in favor of higher quality samples from a truncated subset of the domain [17].

^{*}This work was done during an internship at NVIDIA.



Figure 1: Definition of precision and recall for distributions [24]. (a) Denote the distribution of real images with P_r (blue) and the distribution of generated images with P_g (red). (b) Precision is the probability that a random image from P_g falls within the support of P_r . (c) Recall is the probability that a random image from P_r falls within the support of P_q .

Meanwhile, insufficient coverage of the underlying manifold continues to be a challenge for GANs. Various improvements to network architectures and training procedures tackle this issue directly [25, 19, 11, 15]. While metrics have been proposed to estimate the degree of variation, these have not seen widespread use as they are subjective [1], domain specific [19], or not reliable enough [23].

Recently, Sajjadi et al. [24] proposed a novel metric that expresses the quality of the generated samples using two separate components: *precision* and *recall*. Informally, these correspond to the average sample quality and the coverage of the sample distribution, respectively. We discuss their metric (Section 1.1) and characterize its weaknesses that we later demonstrate experimentally. Our primary contribution is an improved precision and recall metric (Section 2) which provides explicit visibility of the tradeoff between sample quality and variety. Source code of our metric is available at https://github.com/kynkaat/improved-precision-and-recall-metric.

We demonstrate the effectiveness of our metric using two recent generative models (Section 3), StyleGAN [12] and BigGAN [4]. We then use our metric to analyze several variants of StyleGAN (Section 4) to better understand the design decisions that determine result quality, and identify new variants that improve the state-of-the-art. We also perform the first principled analysis of truncation methods [17, 13, 4, 12]. Finally, we extend our metric to estimate the quality of individual generated samples (Section 5), offering a way to measure the quality of latent space interpolations.

1.1 Background

Sajjadi et al. [24] introduce the classic concepts of precision and recall to the study of generative models, motivated by the observation that FID and related density metrics cannot be used for making conclusions about precision and recall: a low FID may indicate high precision (realistic images), high recall (large amount of variation), or anything in between. We share this motivation.

From the classic viewpoint, precision denotes the fraction of generated images that are realistic, and *recall* measures the fraction of the training data manifold covered by the generator (Figure 1). Both are computed as expectations of binary set membership over a distribution, i.e., by measuring how likely is it that an image drawn from one distribution is classified as falling under the support of the other distribution. In contrast, Sajjadi et al. [24] formulate precision and recall through the relative probability densities of the two distributions. The choice of modeling the relative densities comes from an ambiguity, i.e., should the differences between the two distributions be attributed to the generator covering the real distribution inadequately or is the generator producing samples that are unrealistic. The authors resolve this ambiguity by modeling a continuum of precision/recall values where the extrema correspond to the classic definitions. In addition to raising the question of which value to use, their practical algorithm cannot reliably estimate the extrema due to its reliance on relative densities: it cannot, for instance, correctly interpret situations where large numbers of samples are packed together, e.g., as a result of mode collapse or truncation. The k-nearest neighbors based two-sample test by Lopez-Paz et. al. [16] suffers from the same problem. Parallel with our work, Simon et al. [26] extend Sajjadi's formulation to arbitrary probability distributions and provide a practical algorithm that estimates precision and recall by training a post hoc classifier.



Figure 2: (a) An example manifold in a feature space. (b) Estimate of the manifold obtained by sampling a set of points and surrounding each with a hypersphere that reaches its kth nearest neighbor.

We argue that the classic definition of precision and recall is sufficient for disentangling the effects of sample quality and manifold coverage. This can be partially justified by observing that precision and recall correspond to the vertical and horizontal extremal cases in Lin et al.'s [15] theoretically founded analysis of mode collapse regions. In order to approximate these quantities directly, we construct adaptive-resolution finite approximations to the real and generated manifolds that are able to answer binary membership queries: "does sample x lie in the support of distribution P?". Together with existing density-based metrics, such as FID, our precision and recall scores paint a highly informative picture of the distributions produced by generative image models. In particular, they make effects in the "null space" of FID clearly visible.

2 Improved precision and recall metric using k-nearest neighbors

We will now describe our improved precision and recall metric that does not suffer from the weaknesses listed in Section 1.1. The key idea is to form explicit non-parametric representations of the manifolds of real and generated data, from which precision and recall can be estimated.

Similar to Sajjadi et al. [24], we draw real and generated samples from $X_r \sim P_r$ and $X_g \sim P_g$, respectively, and embed them into a high-dimensional feature space using a pre-trained classifier network. We denote feature vectors of the real and generated images by ϕ_r and ϕ_g , respectively, and the corresponding sets of feature vectors by Φ_r and Φ_g . We take an equal number of samples from each distribution, i.e., $|\Phi_r| = |\Phi_g|$.

For each set of feature vectors $\Phi \in {\Phi_r, \Phi_g}$, we estimate the corresponding manifold in the feature space as illustrated in Figure 2. We obtain the estimate by calculating pairwise Euclidean distances between all feature vectors in the set and, for each feature vector, forming a hypersphere with radius equal to the distance to its *k*th nearest neighbor. Together, these hyperspheres define a volume in the feature space that serves as an estimate of the true manifold. To determine whether a given sample ϕ is located within this volume, we define a binary function

$$f(\boldsymbol{\phi}, \boldsymbol{\Phi}) = \begin{cases} 1, \text{ if } \|\boldsymbol{\phi} - \boldsymbol{\phi}'\|_2 \le \|\boldsymbol{\phi}' - \mathrm{NN}_k(\boldsymbol{\phi}', \boldsymbol{\Phi})\|_2 \text{ for at least one } \boldsymbol{\phi}' \in \boldsymbol{\Phi} \\ 0, \text{ otherwise,} \end{cases}$$
(1)

where NN_k (ϕ', Φ) returns kth nearest feature vector of ϕ' from set Φ . In essence, $f(\phi, \Phi_r)$ provides a way to determine whether a given image looks realistic, whereas $f(\phi, \Phi_g)$ provides a way to determine whether it could be reproduced by the generator. We can now define our metric as

$$\operatorname{precision}(\boldsymbol{\Phi}_r, \boldsymbol{\Phi}_g) = \frac{1}{|\boldsymbol{\Phi}_g|} \sum_{\boldsymbol{\phi}_g \in \boldsymbol{\Phi}_g} f(\boldsymbol{\phi}_g, \boldsymbol{\Phi}_r) \quad \operatorname{recall}(\boldsymbol{\Phi}_r, \boldsymbol{\Phi}_g) = \frac{1}{|\boldsymbol{\Phi}_r|} \sum_{\boldsymbol{\phi}_r \in \boldsymbol{\Phi}_r} f(\boldsymbol{\phi}_r, \boldsymbol{\Phi}_g) \quad (2)$$

In Equation (2), precision is quantified by querying for each generated image whether the image is within the estimated manifold of real images. Symmetrically, recall is calculated by querying for each real image whether the image is within estimated manifold of generated images. See Appendix A in the supplement for pseudocode.

In practice, we compute the feature vector ϕ for a given image by feeding it to a pre-trained VGG-16 classifier [27] and extracting the corresponding activation vector after the second fully connected layer. Brock et al. [4] show that the nearest neighbors in this feature space are meaningful in the sense that they correspond to semantically similar images. Meanwhile, Zhang et al. [32] use the intermediate activations of multiple convolutional layers of VGG-16 to define a perceptual metric, which they



Figure 3: (a) Our metric behaves similarly to FID in terms of varying sample count. (b) Precision (blue) and recall (orange) for several neighborhood sizes k. Larger k increases both numbers. Here a trained model ($\psi = 1$) was expanded to a family of models by artificially limiting the variation in the results. We would expect the precision and recall to reach 1.0 and 0.0, respectively, when $\psi \rightarrow 0$. (c) Using Inception-v3 features instead of VGG-16 yields a substantially similar result.

show to correlates well with human judgment for image corruptions. We have tested both approaches and found that feature space, used by Brock at al., works considerably better for the purposes of our metric, presumably because it places less emphasis on the exact spatial arrangement — sparsely sampled manifolds rarely include near-exact matches in terms of spatial structure.

Like FID, our metric is weakly affected by the number of samples taken (Figure 3a). Since it is standard practice to quote FIDs with 50k samples, we adopt the same design point for our metric as well. The size of the neighborhood, k, is a compromise between covering the entire manifold (large values) and overestimating its volume as little as possible (small values). In practice, we have found that higher values of k increase the precision and recall estimates in a fairly consistent fashion, and lower values of k decrease them, until they start saturating at 1.0 or 0.0 (Figure 3b). Tests with various datasets and GANs showed that k = 3 is a robust choice that avoids saturating the values most of the time. Thus we use k = 3 and $|\Phi| = 50000$ in all our experiments unless stated otherwise. Figure 3c further shows that the qualitative behavior of our metric is not limited to VGG-16 – which we use in all tests – as Inception-v3 features lead to very similar results.

3 Precision and recall of state-of-the-art generative models

In this section, we demonstrate that precision and recall computed using our method correlate well with the perceived quality and variation of generated distributions, and compare our metric with Sajjadi et al.'s method [24] as well as the widely used FID metric [9]. For Sajjadi et al.'s method, we use 20 clusters and report $F_{1/8}$ and F_8 as proxies for precision and recall, respectively, as recommended by the authors. We examine two state-of-the-art generative models, StyleGAN [12] trained with the FFHQ dataset, and BigGAN [4] trained on ImageNet [5].

StyleGAN Figure 4 shows the results of various metrics in four StyleGAN setups. These setups exhibit different amounts of truncation and training time, and have been selected to illustrate how the metrics behave with varying output image distributions. Setup A is heavily truncated, and the generated images are of high quality but very similar to each other in terms of color, pose, background, etc. This leads to high precision and low recall, as one would expect. Moving to setup B increases variation, which improves recall, while the image quality and thus precision is somewhat compromised. Setup C is the FID-optimized configuration in [12]. It has even more variation in terms of color schemes and accessories such as hats and sunglasses, further improving recall. However, some of the faces start to become distorted which reduces precision. Finally, setup D preserves variation and recall, but nearly all of the generated images have low quality, indicated by much lower precision as expected.

In contrast, the method of Sajjadi et al. [24] indicates that setups B, C and D are all essentially perfect, and incorrectly assigns setup A the lowest precision. Looking at FID, setups B and D appear almost equally good, illustrating how much weight FID places on variation compared to image quality, also evidenced by the high FID of setup A. Setup C is ranked as clearly the best by FID despite the obvious



Figure 4: Comparison of our method (black dots), Sajjadi et al.'s method [24] (red triangles), and FID for 4 StyleGAN setups. We recommend zooming in to better assess the quality of images.



Figure 5: (a) Example images produced by StyleGAN [12] trained using the FFHQ dataset. It is generally agreed [17, 4, 13, 12] that truncation provides a tradeoff between perceptual quality and variation. (b) With our method, the maximally truncated setup ($\psi = 0$) has zero recall but high precision. As truncation is gradually removed, precision drops and recall increases as expected. The final recall value approximates the fraction of training set the generator can reproduce (generally well below 100%). (c) The method of Sajjadi et al. reports both precision and recall increasing as truncation is removed, contrary to the expected behavior, and the final numerical values of both precision and recall seem excessively high.

image artifacts. The ideal tradeoff between quality and variation depends on the intended application, but it is unclear which application might favor setup D where practically all images are broken over setup B that produces high-quality samples at a lower variation. Our metric provides explicit visibility on this tradeoff and allows quantifying the suitability of a given model for a particular application.

Figure 5 applies gradually stronger truncation [17, 4, 13, 12] on precision and recall using a single StyleGAN generator. Our method again works as expected, while the method of Sajjadi et al. does not. We hypothesize that their difficulties are a result of truncation packing a large number of generated images into a small region in the embedding space. This may result in clusters that contain no real images in that region, and ultimately causes the metric to incorrectly report low precision. The tendency to underestimate precision can be alleviated by using fewer clusters, but doing so leads to overestimation of recall. Our metric does not suffer from this problem because the manifolds of real and generated images are estimated separately, and the distributions are never mixed together.

BigGAN Brock et al. recently presented BigGAN [4], a high-quality generative network able to synthesize images for ImageNet [5]. ImageNet is a diverse dataset containing 1000 classes with \sim 1300 training images for each class. Due to the large amount of variation within and between classes, generative modeling of ImageNet has proven to be a challenging problem [21, 31, 4]. Brock et al. [4] list several ImageNet classes that are particularly easy or difficult for their method. The difficult classes often contain precise global structure or unaligned human faces, or they are underrepresented



Figure 6: Our precision and recall for four easy (a) and four difficult (b) ImageNet classes using BigGAN. For each class we sweep the truncation parameter ψ linearly from 0.3 to 1.0, left-to-right. The FIDs refer to a non-truncated model, i.e., $\psi = 1.0$. The per-class metrics were computed using all available training images of the class and an equal number of generated images, while the curve for the entire dataset was computed using 50k real and generated images.

in the dataset. The easy classes are largely textural, lack exact global structure, and are common in the dataset. Dogs are a noteworthy special case in ImageNet: with almost a hundred different dog breeds listed as separate classes, there is much more training data for dogs than for any other class, making them artificially easy. To a lesser extent, the same applies to cats that occupy ~ 10 classes.

Figure 6 illustrates the precision and recall for some of these classes over a range of truncation values. We notice that precision is invariably high for the suspected easy classes, including cats and dogs, and clearly lower for the difficult ones. Brock et al. state that the quality of generated samples increases as more truncation is applied, and the precision as reported by our method is in line with this observation. Recall paints a more detailed picture. It is very low for classes such as "Lemon" or "Broccoli", implying much of the variation has been missed, but FID is nevertheless quite good for both. Since FID corresponds to a Wasserstein-2 distance in the feature space, low intrinsic variation implies low FID even when much of that variation is missed. Correspondingly, recall is clearly higher for the difficult classes. Based on visual inspection, these classes have a lot of intra-class variation that BigGAN training has successfully modeled. Dogs and cats show recall similar to the difficult classes, and their image quality and thus precision is likely boosted by the additional training data.

4 Using precision and recall to analyze and improve StyleGAN

Generative models have seen rapid improvements recently, and FID has risen as the de facto standard for determining whether a proposed technique is considered beneficial or not. However, as we have shown in Section 3, relying on FID alone may hide important qualitative differences in the results and it may inadvertently favor a particular tradeoff between precision and recall that is not necessarily aligned with the actual goals. In this section, we use our metric to shed light onto some of the design decisions associated with the model itself. Appendix C in the supplement performs a similar, principled analysis for truncation methods. We use StyleGAN [12] in all experiments, trained with FFHQ at 1024×1024 .

4.1 Network architectures and training configurations

To avoid drawing false conclusions when comparing different training runs, we must properly account for the stochastic nature of the training process. For example, we have observed that FID can often vary by up to $\pm 14\%$ between consecutive training iterations with StyleGAN. The common approach is to amortize this variation by taking multiple snapshots of the model at regular intervals and selecting the best one for further analysis [12]. With our metric, however, we are faced with the problem of multiobjective optimization [3]: the snapshots represent a wide range of different tradeoffs between precision and recall, as illustrated in Figure 7a. To avoid making assumptions about the desired tradeoff, we identify the Pareto frontier, i.e., the minimal subset of snapshots that is guaranteed to contain the optimal choice for any given tradeoff.



Figure 7: (a) Precision and recall for different snapshots of StyleGAN taken during the training, along with their corresponding Pareto frontier. We use the standard training configuration by Karras et al. [12] with FFHQ and $\psi = 1$. (b) Different training configurations lead to vastly different tradeoffs between precision and recall. (c) Best FID obtained for each configuration (lower is better).

Figure 7b shows the Pareto frontiers for several variants of StyleGAN. The baseline configuration (A) has a dedicated *minibatch standard deviation* layer that aims to increase variation in the generated images [11, 15]. Using our metric, we can confirm that this is indeed the case: removing the layer shifts the tradeoff considerably in favor of precision over recall (B). We observe that R_1 regularization [18] has a similar effect: reducing the γ parameter by $100 \times$ shifts the balance even further (C). Karras et al. [11] argue that their progressive growing technique improves both quality and variation, and indeed, disabling it reduces both aspects (D). Moreover, we see that randomly translating the inputs of the discriminator by $-16 \dots 16$ pixels improves precision (E), whereas disabling instance normalization in the AdaIN operation [10], unexpectedly, improves recall (F).

Figure 7c shows the best FID obtained for each configuration; the corresponding snapshots are highlighted in Figure 7a,b. We see that FID favors configurations with high recall (A, F) over the ones with high precision (B, C), and the same is also true for the individual snapshots. The best configuration in terms of recall (F) yields a new state-of-the-art FID for this dataset. Random translation (E) is an exceptional case: it improves precision at the cost of recall, similar to (B), but also manages to slightly improve FID at the same time. We leave an in-depth study of these effects for future work.

5 Estimating the quality of individual samples

While our precision metric provides a way to assess the overall quality of a population of generated images, it yields only a binary result for an individual sample and therefore is not suitable for ranking images by their quality. Here, we present an extension of the classification function f (Equation 1) that provides a continuous estimate of how close a given sample is to the manifold of real images.

We define a *realism score* R that increases the closer an image is to the manifold and decreases the further an image is from the manifold. Let ϕ_g be a feature vector of a generated image and ϕ_r a feature vector of a real image from set Φ_r . Realism score of ϕ_q is calculated as

$$R(\boldsymbol{\phi}_{g}, \boldsymbol{\Phi}_{r}) = \max_{\boldsymbol{\phi}_{r}} \left\{ \frac{\left\| \boldsymbol{\phi}_{r} - \mathrm{NN}_{k} \left(\boldsymbol{\phi}_{r}, \boldsymbol{\Phi}_{r} \right) \right\|_{2}}{\left\| \boldsymbol{\phi}_{g} - \boldsymbol{\phi}_{r} \right\|_{2}} \right\}.$$
(3)

This is a continuous extension of $f(\phi_g, \Phi_r)$ with the simple relation that $f(\phi_g, \Phi_r) = 1$ iff $R(\phi_g, \Phi_r) \ge 1$. In other words, when $R \ge 1$, the feature vector ϕ_g is inside the (k-NN induced) hypersphere of at least one ϕ_r .

With any finite training set, the *k*-NN hyperspheres become larger in regions where the training samples are sparse, i.e., regions with low representation. When measuring the quality of a large population of generated images, these underrepresented regions have little impact as it is unlikely that too many generated samples land there — even though the hyperspheres may be large, they are sparsely located and cover a small volume of space in total. However, when computing the realism



Figure 8: Quality of individual samples of BigGAN from eight classes. Top: Images with high realism. Bottom: Images with low realism. We show two images with the highest and lowest realism score selected from 1000 non-truncated images.

score for a single image, a sample that happens to land in such a fringe hypersphere may obtain a wildly inaccurate score. Large errors, even if they are rare, would undermine the usefulness of the metric. We tackle this problem by discarding half of the hyperspheres with the largest radii. In other words, the maximum in Equation 3 is not taken over all $\phi_r \in \Phi_r$ but only over those ϕ_r whose associated hypersphere is smaller than the median. This pruning yields an overconservative estimate of the real manifold, but it leads to more consistent realism scores. Note that we use this approach only with R, not with f.

Figure 8 shows example images from BigGAN with high and low realism. In general, the samples with high realism display a clear object from the given class, whereas the object is often distorted to unrecognizable for the low realism images. Appendix D in the supplement provides more examples.

5.1 Quality of interpolations

An interesting application for the realism score is to evaluate the quality of interpolations. We do this with StyleGAN using linear interpolation in the intermediate latent space W as suggested by Karras et al. [12]. Figure 9 shows four example interpolation paths with randomly sampled latent vectors as endpoints. Paths A appears to be located completely inside the real manifold, path D completely outside it, and paths B and C have one endpoint inside the real manifold and one outside it. The realism scores assigned to paths A–D correlate well with the perceived image quality: Images with low scores contain multiple artifacts and can be judged to be outside the real manifold, and vice versa for high-scoring images. See Appendix D in the supplement for additional examples.

We can use interpolations to investigate the shape of the subset of W that produces realistic-looking images. In this experiment, we sampled without truncation 1M latent vectors in W for which $R \ge 1$, giving rise to 500k interpolation paths with both endpoints on the real manifold. It would be unrealistic to expect all intermediate images on these paths to also have $R \ge 1$, so we chose to consider an interpolation path where more than 25% of the intermediate images have R < 0.9 as straying too far from the real manifold. Somewhat surprisingly, we found that only 2.4% of the paths crossed unrealistic parts of W under this definition, suggesting that the subset of W on the real manifold is highly convex. We see potential in using the realism score for measuring the shape of this region in W with greater accuracy, possibly allowing the exclusion of unrealistic images in a more refined manner than with truncation-like methods.



Figure 9: Realism score for four interpolation paths as function of linear interpolation parameter t and corresponding images from paths A–D. We did not use truncation when generating the images.

6 Conclusion

We have demonstrated through several experiments that the separate assessment of precision and recall can reveal interesting insights about generative models and can help to improve them further. We believe that the separate quantification of precision can also be useful in the context of image-to-image translation [33], where the quality of individual images is of great interest.

Using our metric, we have identified previously unknown training configuration-related effects in Section 4.1, raising the question whether truncation is really necessary if similar tradeoffs can be achieved by modifying the training configuration appropriately. We leave the in-depth study of these effects for future work.

Finally, it has recently emerged that density models can be incapable of assessing whether a given example belongs to the training distribution [22]. By explicitly modeling the real manifold, our metrics may provide an alternative way for estimating this.

7 Acknowledgements

We thank David Luebke for helpful comments; Janne Hellsten, and Tero Kuosmanen for compute infrastructure.

References

- [1] S. Arora and Y. Zhang. Do GANs actually learn the distribution? An empirical study. *CoRR*, abs/1706.08224, 2017.
- [2] M. Bińkowski, D. J. Sutherland, M. Arbel, and A. Gretton. Demystifying MMD GANs. *CoRR*, abs/1801.01401, 2018.
- [3] J. Branke, J. Branke, K. Deb, K. Miettinen, and R. Slowiński. *Multiobjective optimization: Interactive and evolutionary approaches*, volume 5252. Springer Science & Business Media, 2008.
- [4] A. Brock, J. Donahue, and K. Simonyan. Large scale GAN training for high fidelity natural image synthesis. In Proc. ICLR, 2019.
- [5] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In Proc. CVPR, 2009.
- [6] L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using Real NVP. CoRR, abs/1605.08803, 2016.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. In NIPS, 2014.
- [8] A. Grover, M. Dhar, and S. Ermon. Flow-GAN: Combining maximum likelihood and adversarial learning in generative models. In *Proc. AAAI*, 2018.
- [9] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *NIPS*, pages 6626–6637, 2017.
- [10] X. Huang and S. J. Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. *CoRR*, abs/1703.06868, 2017.

- [11] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. *CoRR*, abs/1710.10196, 2017.
- [12] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*, 2019.
- [13] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *CoRR*, abs/1807.03039, 2018.
- [14] D. P. Kingma, D. J. Rezende, S. Mohamed, and M. Welling. Semi-supervised learning with deep generative models. In *Proc. NIPS*, 2014.
- [15] Z. Lin, A. Khetan, G. Fanti, and S. Oh. PacGAN: The power of two samples in generative adversarial networks. *CoRR*, abs/1712.04086, 2017.
- [16] D. Lopez-Paz and M. Oquab. Revisiting classifier two-sample tests. In Proc. ICLR, 2017.
- [17] M. Marchesi. Megapixel size image creation using generative adversarial networks. *CoRR*, abs/1706.00082, 2017.
- [18] L. Mescheder, A. Geiger, and S. Nowozin. Which training methods for GANs do actually converge? CoRR, abs/1801.04406, 2018.
- [19] L. Metz, B. Poole, D. Pfau, and J. Sohl-Dickstein. Unrolled generative adversarial networks. *CoRR*, abs/1611.02163, 2016.
- [20] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. *CoRR*, abs/1802.05957, 2018.
- [21] T. Miyato and M. Koyama. cGANs with projection discriminator. CoRR, abs/1802.05637, 2018.
- [22] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan. Do deep generative models know what they don't know? In *Proc. ICLR*, 2019.
- [23] A. Odena, C. Olah, and J. Shlens. Conditional image synthesis with auxiliary classifier GANs. In *ICML*, 2017.
- [24] M. S. M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, and S. Gelly. Assessing generative models via precision and recall. *CoRR*, abs/1806.00035, 2018.
- [25] T. Salimans, I. J. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved techniques for training GANs. In NIPS, 2016.
- [26] L. Simon, R. Webster, and J. Rabin. Revisiting precision and recall definition for generative model evaluation. *CoRR*, abs/1905.05441, 2019.
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [28] I. Tolstikhin, O. Bousquet, S. Gelly, and B. Schoelkopf. Wasserstein auto-encoders. In Proc. ICLR, 2018.
- [29] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, pages 1747–1756, 2016.
- [30] A. van den Oord, N. Kalchbrenner, O. Vinyals, L. Espeholt, A. Graves, and K. Kavukcuoglu. Conditional image generation with PixelCNN decoders. *CoRR*, abs/1606.05328, 2016.
- [31] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena. Self-attention generative adversarial networks. CoRR, abs/1805.08318, 2018.
- [32] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proc. CVPR*, 2018.
- [33] J. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *CoRR*, abs/1703.10593, 2017.