

Polarity Sampling: Quality and Diversity Control of Pre-Trained Generative Networks via Singular Values

Ahmed Imtiaz Humayun^{†*}, Randall Balestriero^{‡*}, Richard Baraniuk[†]
Rice University[†], Meta AI Research[‡]

Abstract

We present *Polarity Sampling*, a theoretically justified plug-and-play method for controlling the generation quality and diversity of any pre-trained deep generative network (DGN). Leveraging the fact that DGNs are, or can be approximated by, continuous piecewise affine splines, we derive the analytical DGN output space distribution as a function of the product of the DGN’s Jacobian singular values raised to a power ρ . We dub ρ the **polarity** parameter and prove that ρ focuses the DGN sampling on the modes ($\rho < 0$) or anti-modes ($\rho > 0$) of the DGN output-space probability distribution. We demonstrate that nonzero polarity values achieve a better precision-recall (quality-diversity) Pareto frontier than standard methods, such as truncation, for a number of state-of-the-art DGNs. We also present quantitative and qualitative results on the improvement of overall generation quality (e.g., in terms of the Fréchet Inception Distance) for a number of state-of-the-art DGNs, including StyleGAN3, BigGAN-deep, NVAE, for different conditional and unconditional image generation tasks. In particular, *Polarity Sampling* redefines the state-of-the-art for StyleGAN2 on the FFHQ Dataset to FID 2.57, StyleGAN2 on the LSUN Car Dataset to FID 2.27 and StyleGAN3 on the AFHQv2 Dataset to FID 3.95. *Colab Demo*.

1. Introduction

Deep Generative Networks (DGNs) have emerged as the go-to framework for generative modeling of high-dimensional datasets, such as natural images. Within the realm of DGNs, different frameworks can be used to produce an approximation of the data distribution, e.g., Generative Adversarial Networks (GANs) [18], Variational AutoEncoders (VAEs) [31] or flow-based models [40]. But despite the different training settings and losses that each of these frameworks aim to minimize, the evaluation metric of choice that is used to characterize the overall quality of generation is the *Fréchet Inception Distance* (FID) [22].

^{*}equal contribution

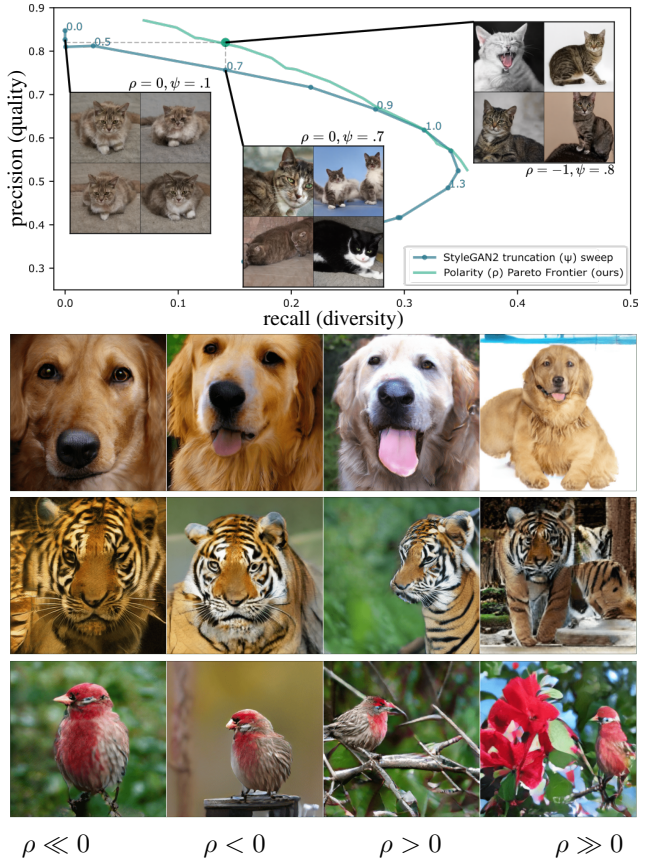


Figure 1. **First row:** Evolution of generation quality and diversity for varying truncation [29] ψ and polarity ρ . Polarity Sampling achieves a better Pareto trade-off than truncation, e.g., polarity can be used to achieve a specified recall at higher precision or a specified precision at higher recall, compared to truncation. For additional Pareto examples, see Fig. 3. **Second, Third, and Fourth row:** Samples obtained from BigGAN-deep on *Golden Retriever*, *Tiger* and *House Finch* classes of Imagenet with samples of greater quality ($\rho < 0$) and greater diversity ($\rho > 0$). For examples with LSUN [54], see Fig. 4.

The FID is obtained by taking the Fréchet Distance in the InceptionV3 [48] embedding space between two distributions; the distributions are usually taken to be the training dataset and samples from a DGN trained on the dataset.

It has been established in prior work [45] that FID nonlinearly combines measures of quality and diversity of the samples, which has inspired further research into disentanglement of these quantities as *precision* and *recall* [32, 45] metrics respectively.

Recent state-of-the-art DGNs such as BigGAN [8], StyleGAN2/3 [28, 30], and NVAE [53], have reached FIDs nearly as low as one could obtain when comparing subsets of real data with themselves. This has led to the deployment of DGNs in a variety of applications, such as real-world high-quality content generation and data-augmentation. However, it is clear that, **depending on the domain of application, generating samples from the best FID model could be suboptimal**. For example, realistic content generation might benefit more from high-quality (precision) samples, while data-augmentation might benefit more from samples of high-diversity (recall), even if in each case, the overall FID slightly diminishes [16, 25]. Therefore, a number of state-of-the-art DGNs have introduced a controllable parameter to trade-off between the precision and recall of the generated samples, e.g., truncated latent space sampling [8], interpolating truncation [29, 30]. However, these methods do not always work “out-of-the-box” [8], e.g., BigGAN requires orthogonal regularization of the DGN’s parameters during training. These methods also lack a clear theoretical understanding which can limit their deployment for sensitive applications.

In this paper, we propose a principled solution to control the quality (precision) and diversity (recall) of DGN samples that does not require retraining nor specific conditioning of model training. Our method, termed *Polarity Sampling*, builds on our previous work on the analytical form of the learned DGN sample distribution [24] and introduces a new hyperparameter, that we dub the *polarity* $\rho \in \mathbb{R}$, that adapts the latent space distribution for post-training control. **The polarity parameter provably forces the latent distribution to concentrate on the modes of the DGN distribution, i.e., regions of high probability ($\rho < 0$), or on the anti-modes, i.e., regions of low-probability ($\rho > 0$); with $\rho = 0$ recovering the original DGN distribution.** The Polarity Sampling process depends only on the top singular values of the DGN’s output Jacobian matrices evaluated at each input sample and can be implemented to perform online sampling. A crucial benefit of Polarity Sampling lies in its theoretical derivation from the analytical DGN data distribution [24] where the product of the DGN Jacobian matrices singular values – raised to the power ρ – provably controls the DGN samples distribution as desired. See Fig. 1 for an initial example of Polarity Sampling in action.

Our main contributions are as follows:

[C1] We first provide the theoretical derivation of Polarity Sampling based on the singular values of the generator Ja-

cobian matrix. We provide pseudocode for Polarity Sampling and an approximation scheme to control its computational complexity as desired (Sec. 3).

[C2] We demonstrate on a range of DGNs and datasets that Polarity Sampling not only enables one to move on the precision-recall Pareto frontier (Sec. 4.1), i.e., it controls the quality and diversity efficiently, but it also reaches improved FID scores for each model (Sec. 4.2).

[C3] We leverage the fact that negative Polarity Sampling provides access to the modes of the learned DGN distribution, which enables us to explore several timely and important questions regarding DGNs. We provide visualization of the modes of trained GANs and VAEs (Sec. 5.1) and assess the perceptual smoothness around the modes (Sec. 5.2).

2. Related Work

Deep Generative Networks as Piecewise-Linear Mappings. In most DGN settings, once training has been completed, sampling new data points is performed by first sampling latent space samples $z_i \in \mathbb{R}^K$ from a latent space distribution $z_i \sim p_z$ and then processing those samples throughout a DGN $G : \mathbb{R}^K \mapsto \mathbb{R}^D$ to obtain the sample $x_i \triangleq G(z_i), \forall i$. One recent line of research that we will rely on through our study consists in formulating DGNs as Continuous Piecewise Affine (CPA) mappings [3, 35], that be expressed as

$$G(z) = \sum_{\omega \in \Omega} (A_\omega z + b_\omega) 1_{\{z \in \omega\}}, \quad (1)$$

where Ω is the input space partition induced by the DGN architecture, ω is a partition-region where z resides, and A_ω, b_ω are the corresponding slope and offset parameters. The CPA formulation of Eq. (1) either represents the exact DGN mapping, when the nonlinearities are CPA e.g. (leaky-)ReLU, max-pooling, or represents a first-order approximation of the DGN mapping. For more background on CPA networks, see [4]. *The key result from [12] that we will leverage is that Eq. (1) is either exact, or can be made close enough to the true mapping G , to be considered exact for practical purposes.*

Post-Training Improvement of a DGN’s Latent Distribution. The idea that the training-time latent distribution p_z might be suboptimal for test-time evaluation has led to multiple research directions to improve the quality of samples post-training. [10, 49] proposed to optimize the samples $z \sim p_z$ based on a Wasserstein discriminator, leading to the *Discriminator Optimal Transport* (DOT) method. That is, after sampling a latent vector z , the latter is repeatedly updated such that the produced datum has greater quality. [50] proposes to simply remove the samples that produce data out of

the true data manifold. This can be viewed as a binary rejection decision of any new sample $z \sim p_z$. [2] were the first to formally introduce rejection sampling based on a discriminator providing a quality estimate used for the rejection sampling of candidate vectors $z \sim p_z$. Replacing rejection sampling with the Metropolis-Hasting algorithm [21] led to the method of [52], coined MH-GAN. An improvement made by [19] was to use the *Sampling-Importance-Resampling* (SIR) algorithm [43]. [26] proposes *latentRS* which consists in training a WGAN-GP [20] on top of any given DGN to learn an improved latent space distribution producing higher-quality samples. [26] also proposes *latentRS+GA*, where the generated samples from that learned distribution are further improved through gradient ascent.

Truncation of the Latent Distribution. Latent space truncation was introduced for high-resolution face image generation by [33] as a method of removing generated artifacts. The authors employed a latent prior of $z \sim \mathcal{U}[-1, 1]$ during training and $z \sim \mathcal{U}[-0.5, 0.5]$ for qualitative improvement during evaluation. The “truncation trick” was formally introduced by [8] where the authors propose re-sampling latents z if they exceed a specified threshold for truncation. The authors also use weight orthogonalization during training to make truncation amenable. Style-based architectures [29, 30] introduce a linear interpolation based truncation in the style-space, which is also designed to converge to the average of the dataset [29]. Ablations for truncation in style-based generators are provided in [32].

3. Introducing The Polarity Parameter From First Principles

In this section, we introduce *Polarity Sampling*, a method that enables us to control the generation quality and diversity of DGNs. We will proceed by first expressing the analytical form of DGNs’ output distribution (Sec. 3.1), and parametrizing the latent space distribution by the singular values of its Jacobian matrix and our *polarity* parameter (Sec. 3.2). We provide pseudo-code and an approximation strategy that enables fast sampling (Sec. 3.3).

3.1. Analytical Output-Space Density Distribution

Given a DGN G , samples are obtained by sampling $G(z)$ with a given latent space distribution, as in $z \sim p_z$. This produces samples that will lie on the *image* of G , the distribution of which is subject to p_z , the DGN latent space partition Ω and per-region affine parameters A_ω, b_ω . We denote the DGN output space distribution as p_G . Under an injective DGN mapping assumption ($g(z) = g(z') \implies z = z'$) (which holds for various architectures, see, e.g., [41]) it is possible to obtain the analytical form of the DGN output distribution by p_G [24]. For a reason that will become clear in the next section, we focus here on the case $z \sim U(\mathcal{D})$

i.e., using a Uniform latent space distribution over the domain \mathcal{D} . Leveraging the Moore-Penrose pseudo inverse [51] $A^\dagger \triangleq (A^T A)^{-1} A^T$, we obtain the following.

Theorem 1. *For $z \sim U(\mathcal{D})$, the probability density $p_G(x)$ is given by*

$$p_G(x) \propto \sum_{\omega \in \Omega} \det(A_\omega^T A_\omega)^{-\frac{1}{2}} \mathbb{1}_{\{A_\omega^\dagger(x - b_\omega) \in \omega \cap \mathcal{D}\}}, \quad (2)$$

where \det is the pseudo-determinant, i.e., the product of the nonzero eigenvalues of $A_\omega^T A_\omega$. (Proof in Appendix B.1.)

Note that one can also view $\det(A_\omega^T A_\omega)^{1/2}$ as the product of the nonzero singular values of A_ω . Theorem 1 is crucial to our development since it demonstrates that the probability of a sample $x = g(z)$ is proportional to the change in volume ($\det(A_\omega^T A_\omega)^{1/2}$) produced by the coordinate system A_ω of the region ω in which z lies in (recall Eq. (1)). If a region $\omega \in \Omega$ has a slope matrix A_ω that contracts the space ($\det(A_\omega^T A_\omega) < 1$) then the output density on that region — mapped to the output space region $\{A_\omega u + b_\omega : u \in \omega\}$ — is increased, as opposed to other regions that either do not contract the space as much, or even expand it ($\det(A_\omega^T A_\omega) > 1$). Hence, the concentration of samples in each output space region depends on how that region’s slope matrix contracts or expands the space, relative to all other regions.

3.2. Controlling the Density Concentration with a Single Parameter

From Theorem 1 we can directly obtain an explicit parametrization of p_z that enables us to control the distribution of samples in the output space, i.e., to control p_G . In fact, note that one can sample from the mode of the DGN distribution by employing $z \sim U(\omega^*), \omega^* = \arg \min_{\omega \in \Omega} \det(A_\omega^T A_\omega)$. Alternatively, one can sample from the region of lowest probability, i.e., the anti-mode, by employing $z \sim U(\omega^*), \omega^* = \arg \max_{\omega \in \Omega} \det(A_\omega^T A_\omega)$. This directly leads to our Polarity Sampling method that adapts the latent space distribution based on the per-region pseudo-determinants.

Corollary 1. *The latent space distribution*

$$p_\rho(z) \propto \sum_{\omega \in \Omega} \det(A_\omega^T A_\omega)^{\frac{\rho}{2}} \mathbb{1}_{\{z \in \omega\}}, \quad (3)$$

where $\rho \in \mathbb{R}$ is the *polarity* parameter, produces the DGN output distribution

$$p_G(x) \propto \sum_{\omega \in \Omega} \det(A_\omega^T A_\omega)^{\frac{\rho-1}{2}} \mathbb{1}_{\{A_\omega^\dagger(x - b_\omega) \in \omega \cap \mathcal{D}\}}, \quad (4)$$

which falls back to the standard DGN distribution for $\rho = 0$, to sampling of the mode(s) for $\rho \rightarrow -\infty$ and to sampling of the anti-mode(s) for $\rho \rightarrow \infty$. (Proof in Appendix B.2.)

Polarity Sampling consists of using the latent space distribution Eq. (3) with a polarity parameter ρ , that is either negative, concentrating the samples toward the mode(s) of the DGN distribution p_G , positive, concentrating the samples towards the anti-modes(s) of the DGN distribution p_G or zero, which removes the effect of polarity. Note that Polarity Sampling changes the output density in a continuous fashion. Its practical effect, as we will see in Sec. 4.1, is to control the quality and diversity of the obtained samples.

3.3. Approximation and Implementation

We now provide the details and pseudocode for the Polarity Sampling procedure that implements Corollary 1.

Computing the \mathbf{A}_ω Matrix. The per-region slope matrix as in Eq. (1), can be obtained given any DGN by first sampling a latent vector $\mathbf{z} \in \omega$, and then obtaining the Jacobian matrix of the DGN $\mathbf{A}_\omega = \mathbf{J}_z G(\mathbf{z}), \forall \mathbf{z} \in \omega$. This has the benefit of directly employing automatic differentiation libraries and thus does not require any exhaustive implementation nor derivation. Computing $\mathbf{J}_z G(\mathbf{z})$ of a generator is not uncommon in practice, e.g., it is employed during path length regularization of StyleGAN2 [30].

Discovering the Regions $\omega \in \Omega$. As per Eq. (3), we need to obtain the singular values of \mathbf{A}_ω (see next paragraph) for each region $\omega \in \Omega$. This is often a complicated task, especially for state-of-the-art DGNs that can have a partition Ω whose number of regions grows with the architecture depth and width [36]. Furthermore, checking if $\mathbf{z} \in \omega$ requires one to solve a linear program [15], which is expensive. As a result, we develop an approximation that consists of sampling many $\mathbf{z} \sim U(\mathcal{D})$ vectors from the latent space (hence our uniform prior assumption in Corollary 1), and computing their corresponding matrices $\mathbf{A}_{\omega(\mathbf{z})}$. This way, we are guaranteed that $\mathbf{A}_{\omega(\mathbf{z})}$ corresponds to the slope of the region ω in which \mathbf{z} falls in, removing the need to check whether $\mathbf{z} \in \omega$. We do so over N samples obtained uniformly from the DGN latent space (based on the original latent space domain). Selection of N can impact performance as this exploration needs to discover as many regions from Ω as possible.

Singular Value Computation. Computing the singular values of \mathbf{A}_ω is an $\mathcal{O}(\min(K, D)^3)$ operation [17]. However, not all singular values might be relevant, e.g., the smallest singular values that are nearly constant across regions ω can be omitted without altering Corollary 1. Hence, we employ only the top- k singular values of \mathbf{A}_ω to speed up singular value computation to $\mathcal{O}(Dk^2)$, details provided in Appendix A.3. (Further approximation could be employed if needed, e.g., power iteration [39]).

While the required number of latent space samples N and the number of top singular values k might seem to be a limitation of Polarity Sampling, we have found in practice that N and k for state-of-the-art DGNs can be set at

Algorithm 1 Polarity Sampling procedure with polarity ρ ; online version and 2D examples in Appendix. Algorithm 2 and Fig. 11. For implementation details, see Sec. 3.3.

Input: $K > 0, S > 0, N \gg S, G, \mathcal{D}, \rho \in \mathbb{R}$
 $\mathcal{Z}, \mathcal{S}, \mathcal{R} \leftarrow [], [], []$
for $n = 1, \dots, N$ **do**
 $\mathbf{z} \sim U(\mathcal{D})$
 $\sigma = \text{SingularValues}(\mathbf{J}_z G(\mathbf{z}), \text{decreasing} = \text{True})$
 $\mathcal{Z}.\text{append}(\mathbf{z})$
 $\mathcal{S}.\text{append}(\rho \sum_{k=1}^K \log(\sigma[k] + \epsilon))$
for $n = 1, \dots, S$ **do**
 $i \sim \text{Categorical}(\text{prob} = \text{softmax}(\mathcal{S}))$
 $\mathcal{R}.\text{append}(\mathcal{Z}[i])$

Output: \mathcal{R}




Figure 2. Effect of Polarity Sampling on FID of a StyleGAN2-F model pretrained on FFHQ for varying number of top- k singular values (left) and varying number of latent space samples N used to obtain per-region slope matrix \mathbf{A}_ω singular values (right) (recall Sec. 3.3 and Algorithm 1). The trend in FIDs to evaluate the impact of ρ stabilizes when using around $k = 40$ singular values and $N \approx 200,000$ latent space samples. For the effect of k and N on precision and recall, see Fig. 9.

$N \approx 200K, k \in [30, 100]$. We conduct a careful ablation study and demonstrate the impact of different choices for N and k in Fig. 2 and Tabs. 3 and 4 in Appendix A.2. Computation times and software/hardware details are provided in Appendix A.3. To reduce round-off errors that can occur for extreme values of ρ , we compute the product of singular values in log-space, as shown in Algorithm 1.

We summarize how to obtain S samples using the above steps in the pseudocode given in Algorithm 1 and provide an efficient solution to reduce the memory requirement incurred when computing the large matrix \mathbf{A}_ω in Appendix A.4. We also provide an implementation that enables online sampling in Algorithm 2 (Appendix A.1). It is also possible to control the DGN prior p_z with respect to a different space than the data-space e.g. inception-space, or with a different input space than the latent-space e.g. style-space in StyleGAN2/3. This incurs no changes in Algorithm 1 except that the DGN is now considered to be either a subset of the original one, or to be composed with a VGG/InceptionV3 network. We provide the implementation details for style-space, VGG-space, and Inception-

space in Appendix A.5. In those cases, the partition Ω and the per-region mapping parameters $\mathbf{A}_\omega, \mathbf{b}_\omega$ are the ones of the corresponding sub-network or composition of networks (recall Eq. (1)). Polarity Sampling adapts the DGN prior distribution to obtain the modes or anti-modes with respect to the considered output spaces.

4. Controlling Precision, Recall, and FID via Polarity

We now provide empirical validation of Polarity Sampling with an extensive array of experiments. Since calculation of distribution metrics such as FID, precision, and recall are sensitive to image processing nuances, we use each model’s original code repository except for BigGAN-deep on ImageNet [13], for which we use the evaluation pipeline specified for ADM [14]. For NVAE (trained on colored-MNIST [1]), we use a modified version of the StyleGAN3 evaluation pipeline. Precision and recall metrics are all based on the implementation of [32]. Metrics in Tab. 2 are calculated for 50K training samples to be able to compare with existing latent reweighing methods. For all other results, the metrics are calculated using $\min\{N_D, 100K\}$ training samples, where N_D is the number of samples in the dataset.

4.1. Polarity Efficiently Parametrizes the Precision-Recall Pareto Frontier

As we have discussed above, Polarity Sampling can explicitly sample from the modes or anti-modes of any learned DGN distribution. Since the DGN is trained to fit the training distribution, sampling from the modes and anti-modes correspond to sampling from regions of the data manifold that are approximated better/worse by the DGN. Therefore, Polarity Sampling is an efficient parameterization of the trade-off between precision and recall of generation [32] since regions with higher precision are regions where the manifold approximation is more accurate.

As experimental proof, we provide in Fig. 3 the precision-recall trade-off when sweeping polarity, and compare it with truncation [29] for pretrained StyleGAN{2,3} architectures. We see that Polarity Sampling offers a competitive alternative to truncation for controlling the precision-recall trade-off of DGNs across datasets and models. For any given precision, the ρ parameter allows us to reach greater recall than what is possible via latent space truncation [29]. And conversely, for any given recall, it is possible to reach a higher precision than what can be attained using latent space truncation. We see that diversity collapses rapidly for latent truncation compared to Polarity Sampling, across all architectures, which is a major limitation. In addition to that, controlling both truncation and polarity allows us to further extend the Pareto frontier for all of our experiments.

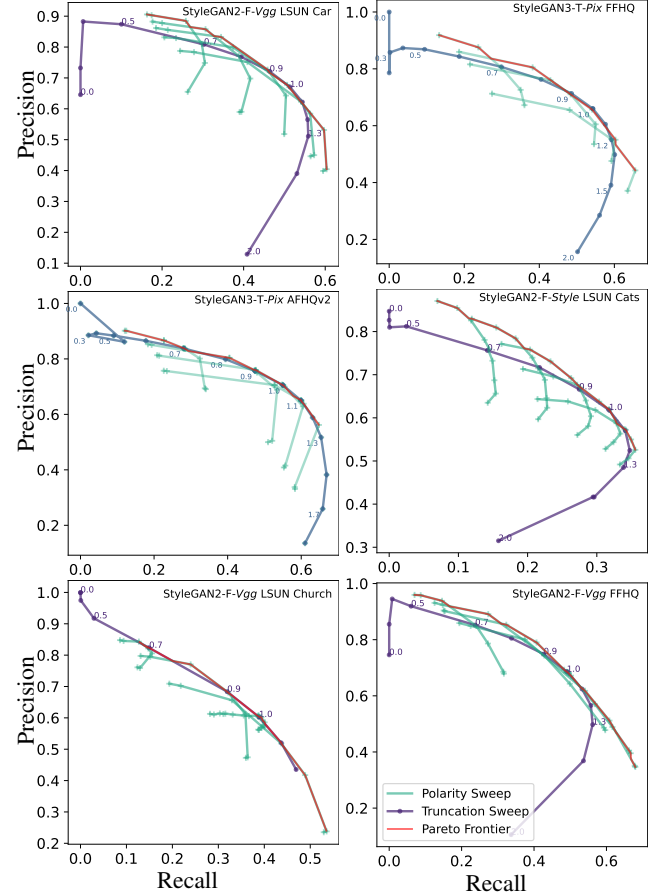


Figure 3. Pareto frontier of the precision-recall metrics can be obtained solely by varying the polarity parameter, for any given truncation level. We depict here six different models and datasets. Results for additional models and datasets are provided in Fig. 1 and Fig. 8.

Apart from the results presented here, we also see that polarity can be used to effectively control the precision-recall trade-off for BigGAN-deep [8] and ProGAN [27]. ProGAN unlike BigGAN and StyleGAN, is not compatible with truncation based methods, i.e., latent space truncation has negligible effect on precision-recall. Hence, polarity offers a great benefit over those existing solutions: Polarity Sampling can be applied regardless of training or controllability factors that are preset in the DGN design. We provide additional results in Appendix C.

4.2. Polarity Improves Any DGN’s FID

We saw in Sec. 4.1 that polarity can be used to control quality versus diversity in a meaningful and controllable manner. In this section, we connect the effect of polarity with FID. Recall that the FID metric nonlinearly combines quality and diversity [45] into a distribution distance measure. Since polarity allows us to control the output distri-

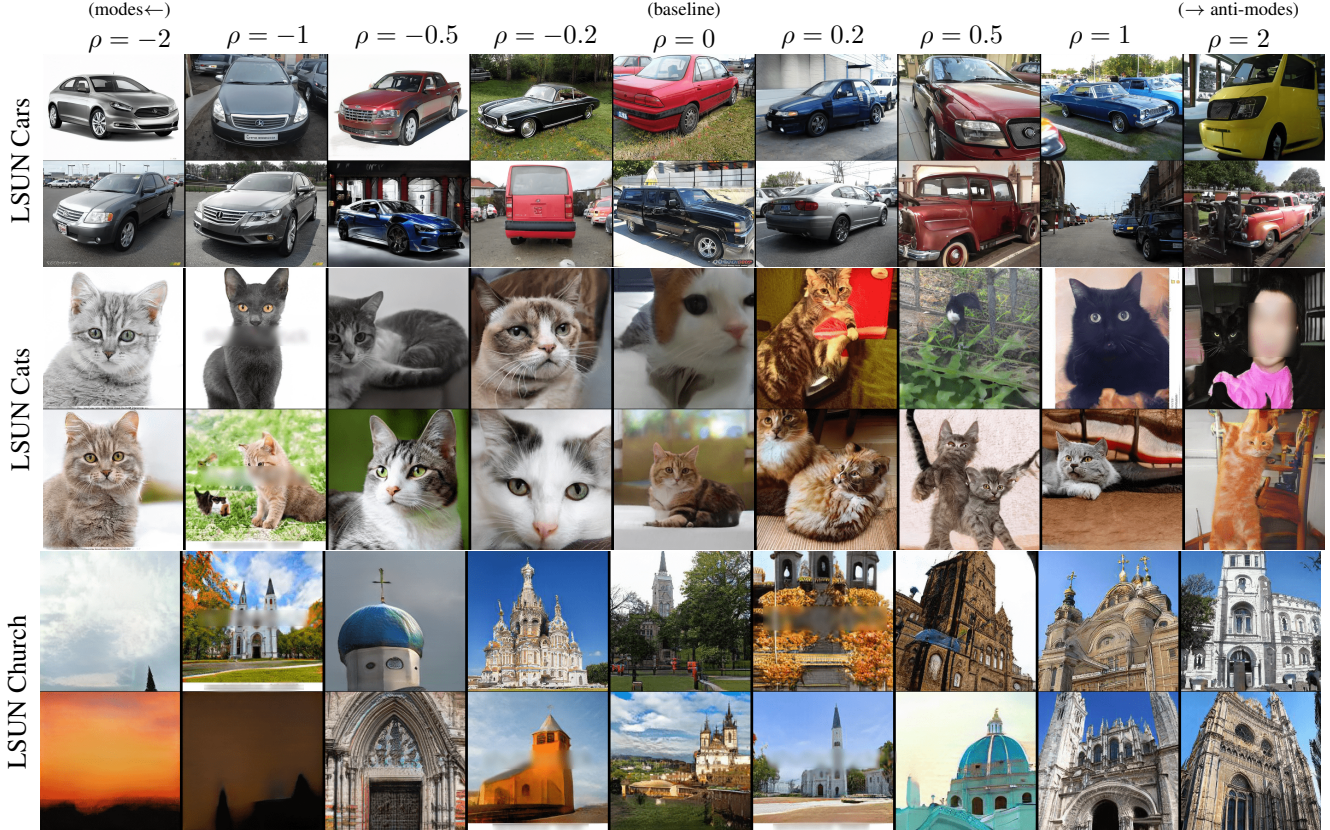


Figure 4. Curated samples of cars and cats for Polarity Sampling in style-space, and church for Polarity Sampling in pixel-space. (Qualitative comparison with truncation sweep in Fig. 10 and nearest training samples in Fig. 12 in the Appendix.) None of the images correspond to training samples, as we discuss in Sec. 5.1.

bution of the DGN, an indirect result of polarity is the reduction of FID by matching the inception embedding distribution of the DGN with that of the training set distribution. Recall that $\rho = 0$ recovers the baseline DGN sampling; for all the state-of-the-art methods in question, we reach lower (better) FID by using a nonzero polarity. In Tab. 1, we compare Polarity Sampling with state-of-the-art solutions that propose to improve FID by learning novel DGN latent space distributions, as were discussed in Sec. 2. We see that for a StyleGAN2 pre-trained on the LSUN church [54] dataset, by increasing the diversity ($\rho = 0.2$) of the VGG embedding distribution, Polarity Sampling surpasses the FID of methods reported in literature that post-hoc improves quality of generation.

In Tab. 2, we present for LSUN {Church, Car, Cat} [54], ImageNet [13], FFHQ [29], and AFHQv2 [11, 28] improved FID obtained *solely by changing the polarity* ρ of a state-of-the-art DGN. This implies that Polarity Sampling provides an efficient solution to adapt the DGN latent space.

We observe that, given any specific setting, $\rho \neq 0$ always improves a model’s FID. We see that in a case specific manner, both positive and negative ρ improves the FID. For

| LSUN Church 256×256 | | | |
|-------------------------------|-------------|------------|------------|
| StyleGAN2 variant | FID ↓ | Prec ↑ | Recall ↑ |
| Standard | 6.29 | .60 | .51 |
| SIR [†] [43] | 7.36 | .61 | .58 |
| DOT [†] [49] | 6.85 | .67 | .48 |
| latentRS [†] [26] | 6.31 | .63 | .58 |
| latentRS+GA [†] [26] | 6.27 | .73 | .43 |
| ρ -sampling 0.2 | 6.02 | .57 | .53 |

Table 1. Comparison of Polarity Sampling with latent reweighting techniques from literature. FID, Precision and Recall is calculated using 50,000 samples. [†]Metrics reported from papers due to unavailability of code. [†]Precision-recall is calculated with 1024 samples only.

StyleGAN2-F trained on FFHQ, *increasing the diversity* of the inception space embedding distribution helps reach a new state-of-the-art FID. By *increasing the precision* of StyleGAN3-T via Polarity Sampling in the Vgg space, we are able to surpass the FID of baseline StyleGAN2-F [28]. We observe that controlling the polarity of the InceptionV3 embedding distribution of StyleGAN2-F gives the most significant gains in terms of FID. This is due to the fact that the

| Model | FID ↓ | Precision ↑ | Recall ↑ | Model | FID ↓ | Precision ↑ | Recall ↑ |
|---------------------------------|-------------|-------------|-------------|---------------------------------|-------------|-------------|-------------|
| LSUN Church 256×256 | | | | LSUN Cat 256×256 | | | |
| DDPM [†] [23] | 7.86 | - | - | ADM (dropout) [†] | 5.57 | 0.63 | 0.52 |
| StyleGAN2 | 3.97 | 0.59 | 0.39 | StyleGAN2 | 6.49 | 0.62 | 0.32 |
| + ρ -sampling Vgg 0.001 | 3.94 | 0.59 | 0.39 | + ρ -sampling Pix 0.01 | 6.44 | 0.62 | 0.32 |
| + ρ -sampling Pix -0.001 | 3.92 | 0.61 | 0.39 | + ρ -sampling Sty -0.1 | 6.39 | 0.64 | 0.32 |
| LSUN Car 512×384 | | | | FFHQ 1024×1024 | | | |
| StyleGAN [†] | 3.27 | 0.70 | 0.44 | StyleGAN2-E | 3.31 | 0.71 | 0.45 |
| StyleGAN2 | 2.34 | 0.67 | 0.51 | Projected GAN [†] [46] | 3.08 | 0.65 | 0.46 |
| + ρ -sampling Vgg -0.001 | 2.33 | 0.68 | 0.51 | StyleGAN3-T | 2.88 | 0.65 | 0.53 |
| + ρ -sampling Sty 0.01 | 2.27 | 0.68 | 0.51 | + ρ -sampling Vgg -0.01 | 2.71 | 0.66 | 0.54 |
| + ρ -sampling Pix 0.01 | 2.31 | 0.68 | 0.50 | | | | |
| ImageNet 256×256 | | | | StyleGAN2-F | 2.74 | 0.68 | 0.49 |
| DCTransformer [†] [37] | 36.51 | 0.36 | 0.67 | + ρ -sampling Ic3 0.01 | 2.57 | 0.67 | 0.5 |
| VQ-VAE-2 [†] [42] | 31.11 | 0.36 | 0.57 | + ρ -sampling Pix 0.01 | 2.66 | 0.67 | 0.5 |
| SR3 [†] [44] | 11.30 | - | - | AFHQv2 512×512 | | | |
| IDDPM [†] [38] | 12.26 | 0.70 | 0.62 | StyleGAN2 [†] | 4.62 | - | - |
| ADM [†] [14] | 10.94 | 0.69 | 0.63 | StyleGAN3-R [†] | 4.40 | - | - |
| ICGAN+DA [†] [9] | 7.50 | - | - | StyleGAN3-T | 4.05 | 0.70 | 0.55 |
| BigGAN-deep | 6.86 | 0.85 | 0.29 | + ρ -sampling Vgg -0.001 | 3.95 | 0.71 | 0.55 |
| + ρ -sampling Pix 0.0065 | 6.82 | 0.86 | 0.29 | | | | |
| ADM+classifier guidance | 4.59 | 0.82 | 0.52 | | | | |

Table 2. [†]Paper reported metrics. We observe that moving away from $\rho = 0$, Polarity Sampling improves FID across models and datasets, empirically validating that the top singular values of a DGN’s Jacobian matrices contain meaningful information to improve the overall quality of generation

Frechet distance between real and generated distributions is directly affected while performing Polarity Sampling in the Inception space. We provide generated samples in Fig. 4 varying the style-space ρ for LSUN cars and LSUN cats, whereas varying the pixel-space ρ for LSUN Church. It is clear that $\rho < 0$ i.e. sampling closer to the DGN distribution modes produce samples of high visual quality, while $\rho > 0$ i.e. sampling closer to the regions of low-probability produce samples of high-diversity, with some samples which are off the data manifold due to the approximation quality of the DGN in that region. Using Polarity Sampling, we are able to advance the state-of-the-art performance on three different settings: for StyleGAN2 on the FFHQ [29] Dataset to FID 2.57, StyleGAN2 on the LSUN [54] Car Dataset to FID 2.27, and StyleGAN3 on the AFHQv2 [28] Dataset to FID 3.95. For additional experiments with ProGAN, and NVAE under controlled training and reference dataset distribution shift, see Appendix C.

5. New Insights into DGN Distributions

In Sec. 4 we demonstrated that Polarity Sampling is a practical method to manipulate DGN output distributions to control their quality and diversity. We now demonstrate that Polarity Sampling has more foundational theoretical applications as well. In particular, we dive into several timely questions regarding DGNs that can be probed using our

framework.

5.1. Are GAN/VAE Modes Training Samples?

Mode collapse [5, 34, 47] has complicated GAN training for many years. It consists of the entire DGN collapsing to generate a few different samples or modes. For VAEs, modes can be expected to be related to the modes of the empirical dataset distribution, as reconstruction is part of the objective. But this might not be the case with GANs e.g., the modes can correspond to parts of the space where the discriminator is the least good at differentiating between true and fake samples. There has been no reported methods in literature that allows us to observe the modes of a trained GAN. Existing visualization techniques focus on finding the role of each DGN unit [6] or finding images that GANs cannot generate [7]. Using Polarity Sampling, we can visualize the modes of DGNs for the first time. In Fig. 5, we present samples from the modes of BigGAN-deep trained on ImageNet, StyleGAN3 trained on AFHQv2, and NVAE trained on colored-MNIST. We observe that BigGAN modes tend to reproduce the unique features of the class, removing the background and focusing more on the object that the class is assigned to. AFHQv2 modes on the other hand, focus on younger animal faces and smoother textures. NVAE mode sampling predominately produce the digit ‘1’ which corresponds to the dataset mode (digit with the least intra-class

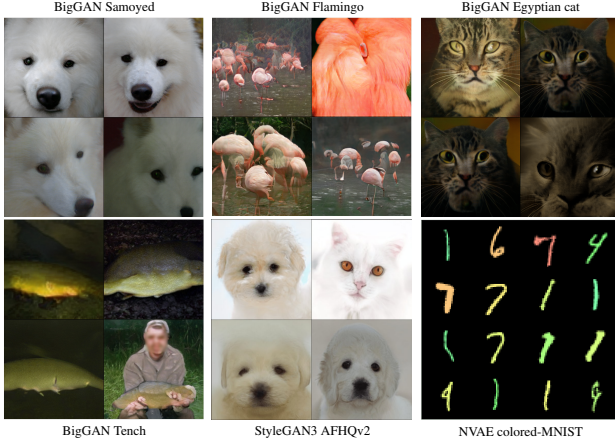


Figure 5. Modes for BigGAN-deep, StyleGAN3-T and NVAE obtained via $\rho \ll 0$ Polarity Sampling. **This is, to the best of our knowledge, the first visualization of the modes of DGNs in pixel space.**




Figure 6. Distribution of l_2 distance to 3 training set nearest neighbors at 32×32 resolution, for 1000 generated samples from LSUN Church StyleGAN2 (left) and colored-MNIST NVAE (right). Samples closer to the modes ($\rho < 0$) have a significant shift in the distribution closer to the training samples for NVAE, while for StyleGAN2 the distribution shift is minimal with significant overlap. This behavior is expected as **VAE models are encouraged to position their modes on the training samples, as opposed to GANs whose modes depend on the discriminator.**

variation). We also provide in Fig. 6 the distribution of the l_2 distances between generated samples and their 3 nearest training samples for modal ($\rho = -5$) and anti-modal ($\rho = 1$) polarity. We see that even after reducing the polarity, StyleGAN2 nearest neighbor distributions have overlap whereas for NVAE the modes move significantly closer to the training samples. In Appendix. Fig. 15 we observe a similar effect for WGAN and NVAE trained on MNIST.

5.2. Perceptual Path Length Around Modes

Perceptual Path Length (PPL) is the distance between the Vgg space image of two latent space points. It has previously been proposed as a measure of perceptual distance [30]. In Fig. 7, we report the PPL of a StyleGAN2-F trained on FFHQ, for an interpolation step of length 10^{-4} between endpoints from the latent/style space. We sample points using Polarity Sampling varying $\rho \in [1, -1]$, es-




Figure 7. Distribution of PPL for StyleGAN2-F trained on FFHQ with varying Polarity Sampling (in Vgg space) setting (ρ given in the legend) for endpoints in the input latent space (left) and endpoints in style-space (right). The means of the distributions (PPL score) are provided as markers on the horizontal axis.

entially measuring the PPL for regions of the data manifold with increasing density as we increase ρ . We see that for negative values of polarity, we have significantly lower PPL compared to positive polarity or even baseline sampling ($\rho = 0$). This result shows that for StyleGAN2, there are smoother perceptual transitions closer to modes. While truncation also reduces the PPL, it essentially does so by sampling points closer to the style space mean [29], see Appendix C.5 for comparisons. Polarity Sampling in the Vgg space, can be used to directly sample from Vgg modes, making it the first method that can be used to explicitly sample regions that are perceptually smoother. It can therefore be used to develop sophisticated interpolation methods where, the interpolation is done along a high-likelihood path on a feature space manifold.

6. Conclusions

We have proposed a new parameterization of the DGN prior p_z in terms of a single parameter – the polarity ρ – to force the DGN samples to be concentrated on the distribution modes or anti-modes (Sec. 3). As a byproduct, for a range of DGNs, we improve the state-of-the-art FID performance. On the theoretical side, Polarity Sampling’s guarantee that it samples from the modes of a DGN enabled us to explore some timely open questions, including the relation between distribution modes and training samples (Sec. 5.1), and the effect of going from mode to anti-mode generation on the perceptual path length (Sec. 5.2). We show that Polarity sampling can also be performed on feature space distributions of classifiers appended with a generator, which can be possibly used for fair attribute generation, out-of-distribution synthetic data generation and much more.

Acknowledgements

Humayun and Baraniuk were supported by NSF grants CCF-1911094, IIS-1838177, and IIS-1730574; ONR grants N00014-18-12571, N00014-20-1-2534, and MURI N00014-20-1-2787; AFOSR grant FA9550-22-1-0060; and a Vannevar Bush Faculty Fellowship, ONR grant N00014-18-1-2047.

References

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 5
- [2] Samaneh Azadi, Catherine Olsson, Trevor Darrell, Ian Goodfellow, and Augustus Odena. Discriminator rejection sampling. *arXiv preprint arXiv:1810.06758*, 2018. 3
- [3] Randall Balestriero and Richard Baraniuk. A spline theory of deep learning. In *ICML*, pages 374–383, 2018. 2
- [4] Randall Balestriero and Richard Baraniuk. Mad max: Affine spline insights into deep learning. *Proceedings of the IEEE*, 109(5):704–727, 2020. 2, 3
- [5] Duhyeon Bang and Hyunjung Shim. Mggan: Solving mode collapse using manifold-guided training. In *ICCV*, pages 2347–2356, 2021. 7
- [6] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B Tenenbaum, William T Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. *arXiv preprint arXiv:1811.10597*, 2018. 7
- [7] David Bau, Jun-Yan Zhu, Jonas Wulff, William Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. In *ICCV*, pages 4502–4511, 2019. 7
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2, 3, 5, 4
- [9] Arantxa Casanova, Marlene Careil, Jakob Verbeek, Michal Drozdal, and Adriana Romero. Instance-conditioned GAN. In *NeurIPS*, 2021. 7
- [10] Tong Che, Ruixiang Zhang, Jascha Sohl-Dickstein, Hugo Larochelle, Liam Paull, Yuan Cao, and Yoshua Bengio. Your gan is secretly an energy-based model and you should use discriminator driven latent sampling. *arXiv preprint arXiv:2003.06060*, 2020. 2
- [11] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, pages 8188–8197, 2020. 6
- [12] Ingrid Daubechies, Ronald A. DeVore, Nadav Dym, Shira Faigenbaum-Golovin, Shahar Z. Kovalsky, Kung-Ching Lin, Josiah Park, Guergana Petrova, and Barak Sober. Neural network approximation of refinable functions. *CoRR*, abs/2107.13191, 2021. 2
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009. 5, 6
- [14] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *arXiv preprint arXiv:2105.05233*, 2021. 5, 7
- [15] Matteo Fischetti and Jason Jo. Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *arXiv preprint arXiv:1712.06174*, 2017. 4
- [16] Edoardo Giacomello, Pier Luca Lanzi, and Daniele Loiacono. Doom level generation using generative adversarial networks. In *Games, Entertainment, Media Conference*, pages 316–323, 2018. 2
- [17] Gene H Golub and Christian Reinsch. Singular value decomposition and least squares solutions. In *Linear algebra*, pages 134–151. Springer, 1971. 4
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 1
- [19] Aditya Grover, Jiaming Song, Alekh Agarwal, Kenneth Tran, Ashish Kapoor, Eric Horvitz, and Stefano Ermon. Bias correction of learned generative models using likelihood-free importance weighting. *arXiv preprint arXiv:1906.09531*, 2019. 3
- [20] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017. 3
- [21] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970. 3
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017. 1
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *arXiv preprint arXiv:2006.11239*, 2020. 7
- [24] Ahmed Imtiaz Humayun, Randall Balestriero, and Richard Baraniuk. Magnet: Uniform sampling from deep generative network manifolds without retraining. In *ICLR*, 2022. 2, 3
- [25] Zubayer Islam, Mohamed Abdel-Aty, Qing Cai, and Jinghui Yuan. Crash data augmentation using variational autoencoder. *Accident Analysis & Prevention*, 151:105950, 2021. 2
- [26] Thibaut Issenhuth, Ugo Tanielian, David Picard, and Jeremie Mary. Latent reweighting, an almost free improvement for gans. *arXiv preprint arXiv:2110.09803*, 2021. 3, 6
- [27] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5, 4, 7
- [28] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021. 2, 6, 7
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410, 2019. 1, 2, 3, 5, 6, 7, 8
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020. 2, 3, 4, 8
- [31] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1
- [32] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *arXiv preprint arXiv:1904.06991*, 2019. 2, 3, 5

- [33] Marco Marchesi. Megapixel size image creation using generative adversarial networks. *arXiv preprint arXiv:1706.00082*, 2017. 3
- [34] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016. 7
- [35] Guido Montúfar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. *arXiv preprint arXiv:1402.1869*, 2014. 2
- [36] Guido Montúfar, Yue Ren, and Leon Zhang. Sharp bounds for the number of regions of maxout networks and vertices of minkowski sums. *arXiv preprint arXiv:2104.08135*, 2021. 4
- [37] Charlie Nash, Jacob Menick, Sander Dieleman, and Peter W Battaglia. Generating images with sparse representations. *arXiv preprint arXiv:2103.03841*, 2021. 7
- [38] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. *arXiv preprint arXiv:2102.09672*, 2021. 7
- [39] Andy Packard, Michael KH Fan, and John C Doyle. A power method for the structured singular value. Technical report, 1988. 4
- [40] Ryan Prenger, Rafael Valle, and Bryan Catanzaro. Waveglow: A flow-based generative network for speech synthesis. In *ICASSP*, pages 3617–3621. IEEE, 2019. 1
- [41] Michael Puthawala, Konik Kothari, Matti Lassas, Ivan Dokmanić, and Maarten de Hoop. Globally injective relu networks. *arXiv preprint arXiv:2006.08464*, 2020. 3
- [42] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *NeurIPS*, pages 14866–14876, 2019. 7
- [43] Donald B Rubin. Using the sir algorithm to simulate posterior distributions. *Bayesian statistics*, 3:395–402, 1988. 3, 6
- [44] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *arXiv preprint arXiv:2104.07636*, 2021. 7
- [45] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *arXiv preprint arXiv:1806.00035*, 2018. 2, 5
- [46] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. *arXiv preprint arXiv:2111.01007*, 2021. 7
- [47] Akash Srivastava, Lazar Valkov, Chris Russell, Michael U Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *NeurIPS*, pages 3310–3320, 2017. 7
- [48] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 1
- [49] Akinori Tanaka. Discriminator optimal transport. *arXiv preprint arXiv:1910.06832*, 2019. 2, 6
- [50] Ugo Tanielian, Thibaut Issenhuth, Elvis Dohmatob, and Jérémie Mary. Learning disconnected manifolds: a no gan’s land. In *ICML*, pages 9418–9427, 2020. 2
- [51] Lloyd N Trefethen and David Bau III. *Numerical linear algebra*, volume 50. Siam, 1997. 3
- [52] Ryan Turner, Jane Hung, Eric Frank, Yunus Saatchi, and Jason Yosinski. Metropolis-hastings generative adversarial networks. In *ICML*, pages 6345–6353, 2019. 3
- [53] Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *arXiv preprint arXiv:2007.03898*, 2020. 2, 4, 6
- [54] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 1, 6, 7

Polarity Sampling: Quality and Diversity Control of Pre-Trained Generative Networks via Singular Values

Supplementary Materials

Codes available at github.com/AhmedImtiazPrio/magnet-polarity

Google Colab demo bit.ly/polarity-samp

We provide the following supplementary materials (SMs) as support of our theoretical and empirical claims. This SM is organized as follows.

Appendix A provides all the implementation details. We first provide the pseudocode for the online sampling algorithm in Appendix A.1 that allows for polarity sampling to be performed rapidly for real-time applications. Along it, we provide further details on the effect of the two hyper-parameters of polarity sampling that are N and k , namely the number of samples used to estimate the per-region singular values of \mathbf{A}_ω for as many ω i.e. N of them, and the number k of top-singular values to utilize (Appendix A.2). We then describe the computation times we observed on our hardware/software (Appendix A.3).

Appendix B provides the proofs for Theorem 1 and Corollary 1. Appendix C supports our claims with additional experiments on various dataset and models. First, Appendix C.1 studies how polarity sampling can help under distribution shift between the training distribution and a target distribution, this is done through colored-MNIST and NVAE. Then Appendix C.2 proposes to study the effect of polarity sampling on ProGAN which is crucial as ProGAN does not allow for truncation based control of its samples.

We conclude with Appendix D and Appendix E that provide descriptions of the datasets and additional qualitative samples from the empirical experiments performed in the main part of the paper, respectively.

A. Implementation Details and Online Sampling Solution

A.1. Online Algorithm

One important aspect of polarity sampling, as summarized in Algorithm 1 is the need to first sample the DGN latent space to obtain the top singular values of as many per-region slope matrices \mathbf{A}_ω as possible. This might seem as a bottleneck if one wants to repeatedly apply polarity sampling on a same DGN. However, this is to provide an estimate of the DGN per-region change of variables, and as the DGN is not retrained nor fine-tuned, it only needs to be done once. Furthermore, this allows for an online sampling algorithm that we provide in Algorithm 2. In short, one first perform this task of estimating as many per-region top singular values as possible, once this is completed, only sampling of latent vectors z and rejection sampling based

on the corresponding \mathbf{A}_ω matrix is done online, \mathbf{A}_ω of the sampled z being obtained easily via $\mathbf{A}_\omega = \mathbf{JG}(z)$.

Algorithm 2 Online Rejection Sampling Algorithm

Input: Latent space domain, \mathcal{D} ; Generator \mathbf{G} ; N change of volume scalars $\{\sigma_1, \sigma_2, \dots, \sigma_N\}$; Number of singular values K ;

while True **do**

$z \sim U(\mathcal{D})$

$\alpha \sim U[0, 1]$

$\mathbf{A} = \mathbf{JG}(z)$

$\sigma_z = \prod_{k=1}^K K\text{-SingularValues}(\mathbf{A}, K)$

if $\frac{\sigma_z^p}{\sigma_z^p + \sum_{i=1}^N \sigma_i^p} \geq \alpha$ **then**

$x \leftarrow \mathbf{G}(z)$

return x

A.2. Effect of N and k

One important aspect of our algorithm comes from the two hyper-parameters N and k . They represent respectively the number of latent space samples to use to estimate as much \mathbf{A}_ω as possible (recall Algorithm 1), and the number of top singular values to compute. Both represent a trade-off between exact polarity sampling, and computation complexity. We argue that in practice, $N \approx 150K$ and $k \approx 100$ is enough to obtain a good estimate of the polarity sampling distribution (Eq. (4)). To demonstrate that, we first provide an ablation study of the number of N and k used for polarity sampling in Tab. 3 and Tab. 4. We also present a visual inspection of the impact of N and k on the precision and recall in Fig. 9.

A.3. Computation Times and Employed Software/Hardware

All the experiments were run on a Quadro RTX 8000 GPU, which has 48 GB of high-speed GDDR6 memory and 576 Tensor cores. For the software details we refer the reader to the provided codebase. In short, we employed TF2 (2.4 at the time of writing), all the usual Python scientific libraries such as NumPy and PyTorch. We employed the official repositories of the various models we employed with official pre-trained weights. As a note, most of the architectures can not be run on GPUs with less or equal to 12 GB of memory.

| N | FFHQ 1024×1024 | | | LSUN Cat 256×256 | | |
|------|-----------------|--------------------|-----------------|------------------|--------------------|-----------------|
| | FID (lowest) | Precision (max) | Recall (max) | FID (lowest) | Precision (max) | Recall (max) |
| 100K | 2.63 | 0.80 | 0.59 | 6.38 | 0.69 | 0.31 |
| 200K | 2.62 | 0.82 | 0.63 | 6.38 | 0.71 | 0.32 |
| 250K | 2.59 | 0.84 | 0.64 | 6.39 | 0.74 | 0.31 |
| 300K | 2.61 | 0.87 | 0.65 | 6.37 | 0.75 | 0.33 |
| 500K | 2.58 | 0.90 | 0.67 | 6.34 | 0.77 | 0.33 |

Table 3. Ablation of N and its effect on best FID, Precision and Recall values that can be obtained by a StyleGAN2 ($\psi = 1$) on the FFHQ and LSUN Cat dataset. We vary the polarity in the VGG space for FFHQ dataset, and style space for LSUN Cats, for number of singular values $k = 30$.

| K | FFHQ 1024×1024 | | |
|-----|-----------------|--------------------|-----------------|
| | FID (lowest) | Precision (max) | Recall (max) |
| 10 | 2.71 | 0.89 | 0.65 |
| 20 | 2.67 | 0.90 | 0.66 |
| 40 | 2.57 | 0.90 | 0.67 |
| 60 | 2.62 | 0.90 | 0.66 |
| 80 | 2.67 | 0.90 | 0.66 |
| 100 | 2.70 | 0.90 | 0.67 |

Table 4. Ablation of K and its effect on the best FID, Precision and Recall values that can be obtained by a StyleGAN2 ($\psi = 1$) on the FFHQ dataset. We vary the polarity in the inception space for FFHQ dataset

We report here the Jacobian computation times for Tensorflow 2.5 with CUDA 11 and Cudnn 8 on an NVIDIA Titan RTX GPU. For StyleGAN2 pixel space, 5.03s/it; StyleGAN2 style-space, 1.12s/it; BigGAN 5.95s/it; ProgGAN 3.02s/it. For NVAE on Torch 1.6 it takes 20.3s/it. Singular value calculation for StyleGAN2 pixel space takes .005s/it, StyleGAN2 style space .008s/it, BigGAN .001s/it, ProgGAN .004s/it and NVAE .02s/it on NumPy. According to this, for StyleGAN2-e, $N=250,000$ requires 14 days to obtain. This only needs to be done once, and it is also possible to perform online sampling once it is calculated. The time required for this is relatively small compared to the training time required for only one set of hyperparameters, which is 35 days and 11 hours¹. We have added pseudocode for MaGNET sampling and online sampling in Appendix G.

Computational Complexity. We are computing the top- k singular values of the $D \times K$ Jacobian matrix. This can be performed in $\mathcal{O}(DKk + Dk^2)$. In fact, one has to project k K -dimensional vectors onto the $D \times K$ Jacobian’s matrix: $\mathcal{O}(DKk)$ and then perform QR-decomposition of the $D \times k$ matrix: $\mathcal{O}(Dk^2)$. Then, k D -dimensional vectors are pro-

jected onto the transpose of the Jacobian matrix: $\mathcal{O}(DKk)$ followed by their QR-decomposition: $\mathcal{O}(Kk^2)$, dominated by $\mathcal{O}(Dk^2)$ (full SVD runs in $\mathcal{O}(DK^2)$).

A.4. Reducing Memory Requirements

The core of polarity sampling relies on computing the top-singular values of the possibly large matrix A_ω for a variety of regions $\omega \in \Omega$, discovered through latent space sampling (recall Algorithm 1). One challenge for state-of-the-art DGNs lies in the size of the matrices A_ω . Multiple solutions exist, such as computing the top singular values through block power iterations. Doing so, the matrices A_ω do not need to be computed entirely, only the matrix-matrix product $A_\omega W$ and $A_\omega^T V$ needs to be performed repeatedly (interleaved with QR decompositions). After many iterations, W estimate the top right-singular vectors of A_ω , and V the corresponding top left-singular vectors from which the singular values can be obtained. However, we found this solution to remain computationally extensive, and found that in practice, a simpler approximation that we now describe provided sufficiently accurate estimates.

Instead of the above iterative estimation, one can instead compute the top-singular values of $W A_\omega$ with W a semi-orthogonal matrix of shape $D' \times D$ with $D' < D$ (recall that A_ω is of shape $D \times K$). Doing so, we are now focusing on the singular values of A_ω whose left-singular vectors are not orthogonal with the right singular vectors of W . While this possibly incurs an approximation error, we found that the above was sufficient to provide polarity sampling and adequate precision-recall control.

A.5. Applying Polarity Sampling in Style, VGG and Inception Space

We call the ambient space of the images the pixel-space, because each dimension in this space corresponds to individual pixels of the images. Apart from controlling the density of the pixel-space manifold, polarity can also be used to control the density of the style-space manifold for style based architectures such as StyleGAN{1,2,3} [28–30]. We also extend the idea of intermediate manifolds to feature space manifolds such as VGG or InceptionV3 space, which can be assumed continuous mappings of the pixel space to the corresponding models’ bottleneck embedding space. In Fig. 8-left we present comparisons between Style, Pixel, VGG and Inception space precision-recall curves for StyleGAN2-F FFHQ with $\psi = 1$, top- $k = 30$ and $\rho = [-2, 2]$. We see that the VGG and InceptionV3 curves trace almost identically. This is expected behavior since both these feature spaces correspond to perceptual features, therefore the transform they induce on the pixel space distribution is almost identical. On the other hand, the pixel space distribution saturates at high polarity at almost equal values. The point of equal precision and

¹<https://github.com/NVlabs/stylegan2>

recall for both the Inception and VGG spaces, occurs at a polarity of 0.1. It's clear from the figures that feature space polarity changes have a larger effect on precision and recall compared to pixel-space and style-space has the least effect on precision and recall. This could be due to the number of density transforms the style-space distribution undergoes until the VGG space, where precision and recall is calculated. In Fig. 8-right we present the polarity characteristics for StyleGAN2-E, StyleGAN2-F and StyleGAN3. For each model, we choose the best space w.r.t the pareto frontier, VGG and Inception space for StyleGAN2-E and StyleGAN2-F, and pixel-space for StyleGAN3. Notice that StyleGAN3 exceeds the recall of the other two models for negative polarity, while matching the precision for StyleGAN2-E.

B. Proofs

The proofs of the two main claims of the paper heavily rely on the spline form of the DGN input-output mapping from Eq. (1). For more background on the form of the latent space partition Ω , the per-region affine mappings $\mathbf{A}_\omega, \mathbf{b}_\omega, \forall \omega \in \Omega$ and further discussion on how to deal with DGN including smooth activation functions, we refer the reader to [4], and in particular to [24] for DGN specific results.

B.1. Proof of Theorem 1

Proof. We will be doing the change of variables $\mathbf{z} = (\mathbf{A}_\omega^T \mathbf{A}_\omega)^{-1} \mathbf{A}_\omega^T (\mathbf{x} - \mathbf{b}_\omega) \triangleq \mathbf{A}_\omega^\dagger (\mathbf{x} - \mathbf{b}_\omega)$, also notice that $J_{G^{-1}}(\mathbf{x}) = \mathbf{A}_\omega^\dagger$. First, we know that $P_{G(\mathbf{z})}(\mathbf{x} \in w) = P_{\mathbf{z}}(\mathbf{z} \in G^{-1}(w)) = \int_{G^{-1}(w)} p_{\mathbf{z}}(\mathbf{z}) d\mathbf{z}$ which is well defined based on our full rank assumptions. We then proceed by

$$\begin{aligned} P_G(\mathbf{x} \in w) &= \sum_{\omega \in \Omega} \int_{\omega \cap w} p_{\mathbf{z}}(G^{-1}(\mathbf{x})) \\ &\quad \times \sqrt{\det(J_{G^{-1}}(\mathbf{x})^T J_{G^{-1}}(\mathbf{x}))} d\mathbf{x} \\ &= \sum_{\omega \in \Omega} \int_{\omega \cap w} p_{\mathbf{z}}(G^{-1}(\mathbf{x})) \\ &\quad \times \sqrt{\det((\mathbf{A}_\omega^\dagger)^T \mathbf{A}_\omega^\dagger)} d\mathbf{x} \\ &= \sum_{\omega \in \Omega} \int_{\omega \cap w} p_{\mathbf{z}}(G^{-1}(\mathbf{x})) \frac{1}{\sqrt{\det(\mathbf{A}_\omega^T \mathbf{A}_\omega)}} d\mathbf{x}, \end{aligned}$$

where the second to third equality follows by noticing that $\sigma_i(\mathbf{A}^\dagger) = (\sigma_i(\mathbf{A}))^{-1}$ which can be showed easily by replacing \mathbf{A}_ω with its SVD and unrolling the product of matrices. Now considering a uniform latent distribution case on a bounded domain U in the DGN latent space we obtain

by substitution in the above result

$$p_G(\mathbf{x}) = \frac{\sum_{\omega \in \Omega} 1_{\mathbf{x} \in \omega} \det(\mathbf{A}_\omega^T \mathbf{A}_\omega)^{-\frac{1}{2}}}{\text{Vol}(U)}, \quad (5)$$

leading to the desired result. \square

B.2. Proof of Corollary 1

Proof. The proof of this result largely relies on Theorem 1. Taking back our previous result, we know that

$$p_G(\mathbf{x}) = \sum_{\omega \in \Omega} p_{\mathbf{z}}(G^{-1}(\mathbf{x})) 1_{\{G^{-1}(\mathbf{x}) \in \omega\}} \frac{1}{\sqrt{\det(\mathbf{A}_\omega^T \mathbf{A}_\omega)}} d\mathbf{x}. \quad (6)$$

However, recall that polarity sampling leverages the prior probability given by

$$p_\rho(\mathbf{z}) = \frac{1}{\kappa} \sum_{\omega \in \Omega} \det(\mathbf{A}_\omega^T \mathbf{A}_\omega)^{\frac{\rho}{2}} \mathbb{1}_{\{\mathbf{z} \in \omega\}}, \quad (7)$$

which, after replacing $G^{-1}(\mathbf{x})$ with its corresponding \mathbf{z} becomes

$$p_G(\mathbf{x}) = \sum_{\omega \in \Omega} \frac{1}{\kappa} \frac{\det(\mathbf{A}_\omega^T \mathbf{A}_\omega)^{\frac{\rho}{2}}}{\sqrt{\det(\mathbf{A}_\omega^T \mathbf{A}_\omega)}} \mathbb{1}_{\{\mathbf{z} \in \omega'\}} d\mathbf{x}, \quad (8)$$

and simplifies to

$$p_G(\mathbf{x}) = \sum_{\omega \in \Omega} \frac{1}{\kappa} \det(\mathbf{A}_\omega^T \mathbf{A}_\omega)^{\frac{\rho-1}{2}} \mathbb{1}_{\{\mathbf{z} \in \omega'\}} d\mathbf{x}, \quad (9)$$

leading to the desired result. Note that when $\rho = 1$ then the density is uniform onto the DGN manifold, when $\rho = 0$, one recovers the original DGN density onto the manifold, and in the extreme cases, only the region with highest or lowest probability would be sampled i.e. the modes or anti-modes. \square

C. Extra experiments

C.1. Polarity Helps Under Distribution Shift

A last benefit of polarity sampling is to adapt a sampling distribution to a reference distribution that suffered a distribution shift. For example, this could occur when training a DGN on a training set, and using it for content generation with a slightly different type target samples. In fact, as long as the distribution shift remains reachable by the model i.e. in the support of p_G , altering the value of ρ will help to shift the sampling distribution, possible to better match the target one. In all generality, there is no guarantee that any benefit would happen for $\rho \neq 0$, however, for the particular case where the distribution shift only changes the way the samples are distributed (on the same domain), we observe in Fig. 14 that $\rho \neq 0$ can provide benefit. To control the




Figure 8. Top: Precision Recall tradeoff for polarity sweep on VGG, Inception and Pixel space distributions. Bottom: BigGAN-deep Imagenet pareto curves obtained for a few classes, in red is the baseline while each scatter point can be reached by varying truncation and ρ . Calculated with 1300 real and generated samples.




Figure 9. Effect of Polarity Sampling on Precision (top) and Recall (bottom) of a StyleGAN2-F model pretrained on FFHQ for varying number of top- k singular values (right) and varying number of latent space samples N (left) used to obtain per-region slope matrix \mathbf{A}_ω singular values (recall Sec. 3.3 and Algorithm 1). The trend in metrics stabilizes when using around $N \approx 300,000$ latent space samples. Increasing the number of top- k singular values to use, amplifies the effect of polarity, saturating at around $k = 50$.

experimental setting, we took the colored-MNIST dataset and the NVAE DGN model [53] and produce a training set with a Gaussian hue distribution favoring blue and two test set, one with same hue distribution and one with uniform hue distribution. We observe that ρ can provide a beneficial distribution shift to go from the biased-hue samples to the uniform-hue one.

C.2. ProGAN Polarity Sweep

Previously in Sec. 4.1, we have drawn note to the fact that ProGAN [27], an architecture which is widely used, but does not incorporate truncation, can also be controlled via polarity sampling. In Tab. 5 we present precision-recall characteristics for polarity sweep on ProGAN. As control, we also perform latent space truncation as in [8] by sampling a truncated gaussian distribution, parameterized by its support $[-\beta, \beta]$. We change β between $[10^{-10}, 10]$ and notice that for β smaller than 10^{-4} , the generator collapses to 0 precision and recall. Other than that, it maintains a precision of 0.72 and recall of 0.34. Using polarity sweep, we also exceed the baseline FID on CelebAHQ 1024x1024 attained by ProGAN; polarity of $-.01$ in pixel-space reduces the FID from 7.37 to 7.28.

C.3. FID for truncated models

While the FID improvement for some of the methods we present are not significantly large, we see that for truncated models, i.e., models with $\psi < 1$, $\rho > 0$ provides significant FID improvements by increasing diversity, e.g., for StyleGAN2-FFHQ with $\psi = \{.9, .7, .5\}$, increasing $\rho > 0$ improves FID by $\{.69, 8.11, 11.1\}$ points. Such truncation is commonly used in practice for qualitative experiments, making polarity sampling particularly relevant in such settings. Since truncation reduces the range of the generator and polarity increases the diversity of sampling within the range, both can be combined to achieve greater FID improvements.

C.4. NVAE Negative Log-Likelihood for Varying ρ

To validate the effect of ρ on the likelihood of generated samples, we estimate the negative log-likelihood of samples generated via an NVAE trained on colored-MNIST while varying ρ . We generate 5000 samples each for $\rho = \{-5, -1, -.5, -.1, 0, .1, .5, 1, 5\}$ which yields negative log-likelihood values of $\{3.0, 3.2, 3.5, 3.8, 3.9, 4.1, 4.3, 4.6, 5.1\} \times 10^{-2}$ bits/dim. This shows that decreasing $\rho < 0$ samples high-likelihood points while increasing $\rho > 0$ samples lower-likelihood points compared to standard sampling ($\rho = 0$).

C.5. Perceptual Path Length for Constant Latent Shifts

In Sec. 5.2 we present the PPL variation for a 10^{-4} interpolation step from a latent space point towards another random latent space point. To evaluate the perceptual smoothness around regions of the latent space, we also calculate the PPL for paths of length .3 starting from individual latent space points towards random directions. We see that for both StyleGAN2-FFHQ and BigGAN-Imagenet, PPL decreases monotonically with $\rho < 0$ whereas first de-




Figure 10. Modes for BigGAN-deep trained on Imagenet, and conditioned on a specific class: “pug” (top left), “lion” (top right), “cheeseburger” (bottom left) and “pomerian” (bottom right). We observe that the modes correspond to nearly aligned faces with little to no background. Variation of colors and sizes can be seen across the modes. The same observation can be made for the cheeseburger, nearly no background is present, and the shape is consistent to a typical cheeseburger “template”. See Fig. 19 for additional classes.

creases and then increases for $\psi < 1$. For StyleGAN2-FFHQ we acquire $\text{PPL} = \{281, 316\}$ for $\psi = \{0.1, 0\}$ and $\text{PPL} = \{274, 271\}$ for $\rho = \{-2, -10\}$. For BigGAN-GoldenRetriever we get PPL 35.9 for $\psi = 0$ and 0.20 for

$\rho = -2$. This possibly indicates that decreasing truncation might not always lead to perceptually smoother regions whereas decreasing polarity does.




Figure 11. Polarity sweep for a WGAN trained on 2D toy datasets with 4 gaussians (top row), 4 gaussians with triangular domain (middle row) and two circles (bottom row).




Figure 12. Modes for a VAE trained on colored-MNIST with 8 nearest neighbors. The leftmost column for each figure contains generated samples. Notice the higher prevalence of digit 1. Due to low pixel variations, digit 1 samples have high density on the manifold.




Figure 13. Singular value distribution and Marchenko Pastur distribution fit for StyleGAN2-FFHQ (left) and BigGAN-Imagenet (middle). Log-sigma distribution for StyleGAN2-FFHQ with varying ρ (right).




Figure 14. FID, precision and recall for an NVAE trained on colored-MNIST with hue bias. Metrics are calculated for a test dataset with hue distribution **Left**: identical to training and **Right**: uniformly distributed across digit classes. **Polarity allows adapting the DGN output distribution to balance possible distribution shifts.**




Figure 15. Distribution of l_2 distance to 3 MNIST training set nearest neighbors, for 1000 generated MNIST samples from WGAN (left) and NVAE (right). For $\rho = -5$ we see that both distributions have a peak around 5. For WGAN the distribution has a significantly longer tail compared to NVAE, indicating that the WGAN modes don't necessarily coincide with training points.

D. Dataset Description

D.1. colored-MNIST

We perform controlled experiments on NVAE [53] by training on datasets with and without controllable distribution shifts. To control the shift, we colorize MNIST with hue ranging $[0, \pi]$ by 1) uniformly sampling the hue and 2)

| CelebAHQ 1024×1024 | | | | | | |
|--------------------|-------------|---------------|--------|-------|------------|------------|
| $ \rho $ | FID | $\rho \leq 0$ | | FID | $\rho > 0$ | |
| | | Precision | Recall | | Precision | Recall |
| 0 | 7.37 | .73 | .34 | - | - | - |
| 0.01 | 7.28 | .73 | .34 | 7.45 | .73 | .35 |
| 0.1 | 7.41 | .76 | .31 | 8.95 | .68 | .38 |
| 1 | 12.65 | .85 | .19 | 17.96 | .58 | .48 |
| 2 | 13.09 | .86 | .19 | 18.54 | .58 | .48 |

Table 5. FID, Precision and Recall metrics of ProGAN [27] with polarity sweep in the pixel space.

sampling the hue for each image from a truncated normal distribution, with a truncation scale of 2.

D.2. LSUN Dataset

We use the LSUN dataset [54] available at the official website². We preprocess the dataset using the StyleGAN2 repository³.

D.3. AFHQv2 and FFHQ

We use the version 2 of AFHQ that was released prepackaged with StyleGAN3 for our experiments. For FFHQ we use also use TFRecords provided with StyleGAN2.

D.4. License

The majority of Polarity-Sampling is licensed under CC-BY-NC, however portions of the project are available under separate license terms: NVAE, StyleGAN2 and StyleGAN3 are licensed under the NVIDIA license; guided-diffusion is licensed under the MIT license.⁴

E. Qualitative results

We provide in the following pages, Figs. 16 to 19 that correspond to LSUN Cats, LSUN Cars, AFHQv2 samples with varying ρ values and 800 Imagenet modes. For LSUN Cars and LSUN Cats we draw comparisons between varying truncation and varying polarity independently.

²<https://www.yf.io/p/lsun>

³<https://github.com/NVlabs/stylegan2>

⁴<https://github.com/openai/guided-diffusion/blob/main/LICENSE>



Figure 16. Uncurated samples of LSUN Cats using (top) $\rho = \{-1, -.5, -.2, -.1, 0\}$, $\psi = .8$ and (bottom) $\psi = \{.7, .73, .75, .77, .8\}$; both representing regions with roughly an equal span of recall score on Fig. 3. Notice the significant precision of the left-most columns of top compared to the left-most of bottom, where at equal diversity, top has significantly higher precision score.



Figure 17. Uncurated samples of LSUN Cars using (top) $\rho = \{-.1, -.075, -.05, -.025, 0\}$, $\psi = .7$ and (bottom) $\psi = \{.5, .55, .6, .65, .7\}$; both representing regions with roughly an equal span of precision score on Fig. 3 Notice the significant diversity in (top) especially in the leftmost columns, where the recall score is significantly higher than that of the leftmost column of bottom left, with equal precision.



Figure 18. Uncurated samples of AFHQv2 using $\rho = \{-2, -1, -.5, -.2, 0, .01, .1, .2, .5\}$ and $\psi = .9$ in pixel-space. As we move right from baseline (middle column) we see an increase in texture diversity of images, whereas, moving left, we see images with smoother textures.




Figure 19. Depiction of a single mode (large negative ρ) for each class of the first 800 Imagenet classes.