MaGNET: UNIFORM SAMPLING FROM DEEP GENERATIVE NETWORK MANIFOLDS WITHOUT RETRAINING

Anonymous authors
Paper under double-blind review

ABSTRACT

Deep Generative Networks (DGNs) are extensively employed in Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and their variants to approximate the data manifold, and data distribution on that manifold. However, training samples are often obtained based on preferences, costs, or convenience producing artifacts in the empirical data distribution e.g. the large fraction of smiling faces in the CelebA dataset or the large fraction of dark-haired individuals in FFHQ). *These inconsistencies will be reproduced when sampling from the trained DGN, which has far-reaching potential implications for fairness, data augmentation, anomaly detection, domain adaptation, and beyond.* In response, we develop a differential geometry based sampler -coined MaGNET- that, given any trained DGN, produces samples that are uniformly distributed on the learned manifold. We prove theoretically and empirically that our technique produces a uniform distribution on the manifold regardless of the training set distribution. We perform a range of experiments on various datasets and DGNs. One of them considers the state-of-the-art StyleGAN2 trained on FFHQ dataset, where uniform sampling via MaGNET increases distribution precision & recall by 4.12% & 3.01% and decreases gender bias by 41.2%, without requiring labels or retraining.

Figure 1: Random batches of StyleGAN2 samples with 1024×1024 resolution, generated using standard sampling (**left**), uniform sampling via MaGNET on the learned pixel-space manifold (**middle**), and uniform sampling on the style-space manifold (**right**) of the same model. MaGNET sampling yields a higher number of young faces, better gender balance, and greater background/accessory variation, without the need for labels or retraining. Images are sorted and color coded according to age & gender (red/green) as predicted by the Microsoft Cognitive API. Larger batches of images and attribute distributions are furnished in Appendix E.

1 INTRODUCTION

Deep Generative Networks (DGNs) are Deep Networks (DNs) trained to learn latent representations of datasets with frameworks such as Generative Adversarial Networks (GANs) (Goodfellow et al., 2014), Variational Autoencoders (VAEs) (Kingma & Welling, 2013), flow-based models such as NICE (Dinh et al., 2014), and their variants (Dziugaite et al., 2015; Zhao et al., 2016; Durugkar et al., 2016; Arjovsky et al., 2017; Mao et al., 2017; Yang et al., 2019; Fabius & van Amersfoort, 2014; van den Oord et al., 2017; Higgins et al., 2017; Tomczak & Welling, 2017; Davidson et al., 2018; Dinh et al., 2016; Grathwohl et al., 2018; Kingma & Dhariwal, 2018). A common assumption that we will carry though our study is that the datasets of interest are not uniformly distributed in their ambient space, but rather are concentrated on, or around, manifolds of lower intrinsic dimension e.g. the manifold of natural images (Peyré, 2009). All those DGNs training methods have been developed and refined to obtain DGNs that approximate as closely as possible the training set distribution. This becomes an Achilles heel when the training set, regardless of its size, is not representative of the true data distribution, i.e., when the training samples have been curated based on cost or availability that result in implicit/explicit biases. In such a scenario, while the training samples will lie on the true data manifold, the density distribution of the training set will be different from the natural distribution of the data.

Deploying a DGN trained with an inaccurate data distribution can be catastrophic, in particular, when employed for tasks such as data augmentation (Sandfort et al., 2019), controlled data generation for exploration/interpretation (Thirumuruganathan et al., 2020), or estimation of statistical quantities of the data geometry, such as the Lipschitz constant of the data manifold (Gulrajani et al., 2017; Scaman & Virmaux, 2018). Biased data generation from DGNs due to skewed training distributions also raises serious concerns in terms of fair machine learning (Hwang et al., 2020; Tan et al., 2020). To ameliorate the impact of biased training set sampling on the performance of DGNs, *we propose* **MaGNET** *(for MAximum entropy Generative NETwork), a simple and efficient modification to any DGN that adapts its latent space distribution such that its generated samples lie uniformly on the underlying data manifold, regardless of the training set sampling*. Even in the presence of an incorrect or modified training set distribution, as long as the DGN is able to learn the correct underlying manifold, our method provably samples uniformly on the true data manifold. Importantly, *MaGNET can be employed on any pre-trained DGN regardless of its training setting*, reducing the requirement of fine-tuning or retraining of the DGN. This is crucial, as many models such as BigGAN (Brock et al., 2019) and StyleGAN (Karras et al., 2020), have significant computational and energy requirements for training. A plug-and-play method is thus greatly preferred to ease deployment in any already built/trained deep learning pipeline.

While there has been rigorous work on DGNs aiming at improving the training stability of those models, deriving theoretical approximation results, understanding the role of the DGN architectures, or even numerical approximations to speed-up training and deployment of trained models (Mao et al., 2017; Chen et al., 2018; Arjovsky & Bottou; Miyato et al., 2018; Xu & Durrett, 2018; Liu et al., 2017; Zhang et al., 2017; Biau et al., 2018; Li et al., 2017; Kodali et al., 2017; Roy et al., 2018; Andrés-Terré & Lió, 2019; Chen et al., 2018; Balestriero et al., 2020a; Tomczak & Welling, 2016; Berg et al., 2018). Existing methods (Metz et al., 2016; Tanaka, 2019; Che et al., 2020) also try to improve the data distribution approximation and tackle mode dropping, but this in turn can increase the bias learned implicitly by the DGNs. We are the first to consider the task of providing uniform sampling on the DGN underlying manifold, which has far reaching consequences ranging from producing DGNs stable to data curation and able to handle inconsistencies such as repeated samples in the training set. We provide a first-of-its-kind provable uniform sampling on the data manifold that can be used to speed up estimation of various geometric quantities such as estimation of the Lipschitz constant.

MaGNET applies to any (pretrained) DGN architecture (GAN, VAE, NF, etc.) using continuous piece-wise affine (CPA) non-linearities, such as (leaky) ReLU; smooth nonlinearities are dealt with via a first-order Taylor approximation argument. Our main contributions are as follows:
[C1] We characterize the transformation incurred by a density distribution when composed with a CPA mapping (Sec. 3.1) and derive the analytical sampling strategy that allows to obtain a uniform distribution on a manifold that is continuous and piecewise affine (Sec 3.2).
[C2] We observe that current DGNs produce continuous piecewise affine manifolds, and we demonstrate how to leverage [C1] to produce uniform sampling on the manifold of any DGN (Sec. 3.2).
[C3] We provide several carefully controlled experiments validating the importance of uniform sam-

pling as well as showcasing the performance of MaGNET on pretrained models such as BigGAN (Brock et al., 2019), StyleGAN2 (Karras et al., 2020), progGAN (Karras et al., 2017), and NVAE (Vahdat & Kautz, 2020) (Sec. 4.2 and Sec. 4.3).

Reproducible code for the various experiments and figures will be provided upon completion of the review process. Computation and software details are provided in Appendix F, with the proofs of our results in Appendix G. Discussion on the settings in which MaGNET is desirable and on possible limitations is provided in (Sec. 5).

2 Background

Deep Generative Networks. A deep generative network (DGN) is a (nonlinear) operator $\boldsymbol{G}_\Theta$ with parameters Θ mapping a *latent* input $\boldsymbol{z} \in \mathbb{R}^S$ to an *observation* $\boldsymbol{x} \in \mathbb{R}^D$ by composing L intermediate *layer* mappings. The only assumption we require for our study is that the nonlinearities present in the DGN are continuous piecewise affine as is the case with (leaky-)ReLU, absolute value, max-pooling. If employing smooth nonlinearities, then all our results hold from a first-order Taylor approximation argument. Precise definitions of DGN operators can be found in Goodfellow et al. (2016). We will omit Θ from the $\boldsymbol{G}_\Theta$ operator for conciseness unless needed. It is also common to refer to $\boldsymbol{z}$ as the *latent representation*, and $\boldsymbol{x}$ as the *generated/observed data*, e.g, a time-series or image. They can be viewed equivalently as signals, flattened column vectors, or tensors depending on the context. One property of DGNs that emerge whenever employing nonlinearities such as (leaky-)ReLU, max-pooling and the likes is that the entire input-output mapping becomes a continuous piecewise affine spline; we introduce those operators below.

Continuous Piecewise Affine Mappings. One rich class of functions emerges from piecewise polynomials: spline operators. In short, given a partition Ω of a domain $\mathbb{R}^S$, a spline of order k is a mapping defined by a polynomial of order k on each region $\omega \in \Omega$ with continuity constraints on the entire domain for the derivatives of order $0,\ldots,k-1$. As we will focus entirely on affine splines ($k = 1$) we only define this case for concreteness. An affine spline S produces its output via

$$\boldsymbol{S}(\boldsymbol{z}) = \sum_{\omega \in \Omega} (\boldsymbol{A}_\omega \boldsymbol{x} + \boldsymbol{b}_\omega) \mathbb{1}_{\{\boldsymbol{z} \in \omega\}}, \tag{1}$$

with $\boldsymbol{A}_\omega, \boldsymbol{b}_\omega$ the per-region *slope* and *offset* parameters respectively with constraints that the entire mapping is continuous on the domain i.e. $S \in \mathcal{C}^0(\mathbb{R}^S)$. Spline operators and especially affine spline operators have been extensively used in function approximation theory (Cheney & Light, 2009), optimal control (Egerstedt & Martin, 2009), statistics (Fantuzzi et al., 2002) and all related fields.

3 Continuous Piecewise Affine Mapping of a Probability Density

In this section, we study the properties of a probability density that is transformed by a Continuous Piecewise Affine (CPA) mapping. Our goal is to analytical derive the produced density and characterize its properties such as how the per-region affine mappings impact the density concentration. This section presents some key results that will serve as the backbone of our core result in the next section: how to sample uniformly from the manifold generated by Deep Generative Networks.

3.1 Density on the Generated Manifold

Consider an affine spline operator $\boldsymbol{S}$ going from a space of dimension S to a space of dimension D with $D \geq S$. The image of this mapping is a CPA manifold of dimension at most S, the exact dimension is determined by the rank of the per-region slope matrices. Formally the image of $\boldsymbol{S}$ is given by

$$Im(\boldsymbol{S}) \triangleq \{\boldsymbol{S}(\boldsymbol{z}) : \boldsymbol{z} \in \mathbb{R}^S\} = \bigcup_{\omega \in \Omega} \text{Aff}(\omega; \boldsymbol{A}_\omega, \boldsymbol{b}_\omega) \tag{2}$$

with $\text{Aff}(\omega; \boldsymbol{A}_\omega, \boldsymbol{b}_\omega) = \{\boldsymbol{A}_\omega \boldsymbol{z} + \boldsymbol{b}_\omega : \boldsymbol{z} \in \omega\}$ the affine transformation of region ω by the per-region parameters $\boldsymbol{A}_\omega \boldsymbol{b}_\omega$.

From (Eq. 2) we observe that the generated manifold i.e. surface is made of regions that are the affine transformations of the latent space partition regions $\omega \in \Omega$ based on the coordinate change induced by $\boldsymbol{A}_\omega$ and the shift induced by $\boldsymbol{b}_\omega$. We visualize this in Fig. 2 for a toy spline operator with 2-dimensional latent space and 3-dimensional ambient/output space. In the remaining of our study we will denote for conciseness $\boldsymbol{S}(\omega) \triangleq \text{Aff}(\omega; \boldsymbol{A}_\omega, \boldsymbol{b}_\omega)$.

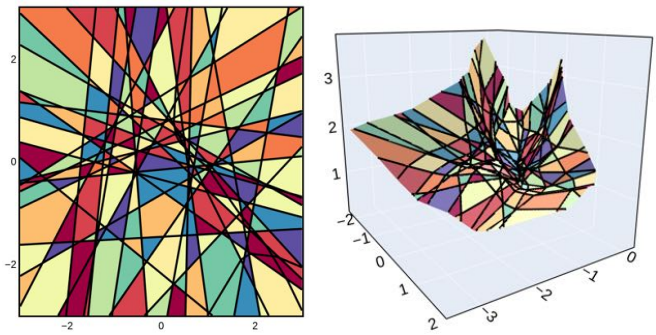

Figure 2: Visual depiction of Eq. 2 with an affine spline mapping $\boldsymbol{S} : \mathbb{R}^2 \mapsto \mathbb{R}^3$. **Left:** latent space partition Ω made of different regions shown with different colors and with boundaries shown in black. **Right:** affine spline image $Im(\boldsymbol{S})$ which is a continuous piecewise affine surface composed of the latent space regions affinely transformed by the per-region affine mappings. The region colors between left and right are the same.

Whenever the input space is equipped with a density distribution, then this density is transformed by the mapping $\boldsymbol{S}$ and will "live" on the surface of the CPA manifold generated by $\boldsymbol{S}$. In fact, given a distribution $\boldsymbol{p_z}$ over the latent space, we can explicitly compute the output distribution after the application of $\boldsymbol{S}$, which leads to an intuitive result exploiting the piecewise affine property of the generator. For this result, we require that operator $\boldsymbol{S}$ to be bijective between its domain and its image. That is, each slope matrix $\boldsymbol{A}_\omega, \forall \omega \in \Omega$ should be full-rank, and there should not be any folding of the generated CPA surface that intersects with itself i.e. $\boldsymbol{S}(\omega) \cap \boldsymbol{S}(\omega') \neq \{\} \iff \omega = \omega'$. We can now derive the key result of this section that characterize the density distribution lying on the manifold.

Lemma 1. *The volume of a region $\omega \in \Omega$ denoted by $\mu(\omega)$ is related to the volume the affinely transformed region $\boldsymbol{S}(\omega)$ by*

$$\frac{\mu(\boldsymbol{S}(\omega))}{\mu(\omega)} = \sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}, \tag{3}$$

where $\mu(\boldsymbol{S}(\omega))$ is using the measure on the S-dimensional affine subspace spanned by the CPA mapping. (Proof in Appendix G.1.)

Theorem 1. *The generated probability density $p_{\boldsymbol{S}}(\boldsymbol{x})$ given $\boldsymbol{p_z}$ and $\boldsymbol{S}$ is given by*

$$p_{\boldsymbol{S}}(\boldsymbol{x}) = \sum_{\omega \in \Omega} \frac{\boldsymbol{p_z}\left(\left(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega\right)^{-1} \boldsymbol{A}_\omega^T \left(\boldsymbol{x} - \boldsymbol{b}_\omega\right)\right)}{\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} \mathbb{1}_{\{\boldsymbol{x} \in \boldsymbol{S}(\omega)\}}. \tag{4}$$

(Proof in Appendix G.2.)

That is, the distribution obtained in the output space naturally corresponds to a piecewise affine transformation of the original latent space distribution, weighted by the change in volume of the per-region mappings from (Eq. 3). The use of the above results to obtain the analytical form of the density covering the generative manifold in the Gaussian and Uniform prior cases is done in Sec. B in the appendix along with its differential Entropy.

3.2 Making the Density on the Manifold Uniform

The goal of this section is to build on Thm. 1 to provide a novel latent space distribution such that the density distribution lying on the generated manifold is uniform.

One important remark that we shall highlight is that having a Uniform density distribution in the latent space of the affine spline is not enough to have a uniform density lying on the manifold; it would be if $\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega) = \det(\boldsymbol{A}_{\omega'}^T \boldsymbol{A}_{\omega'}), \forall \omega \neq \omega'$. This is especially clear from (Eq. 8). So we propose here to design a novel latent space sampler with the purpose that once transformed by the affine spline i.e. the DGN, that distribution becomes uniform on the DGN manifold. We focus here on the technical aspect and will propose precise motivations behind such construction in the next section that deals with practical applications of this following result. To obtain K samples uniformly distributed:

1. sample $N \gg K$ (as large as possible) iid latent vectors with U being the considered latent space domain $(\boldsymbol{z}_1, \ldots, \boldsymbol{z}_N)$, with $\boldsymbol{z}_i \sim \mathcal{U}(U)$

2. compute the per-region slope matrices $\boldsymbol{A}_i = \boldsymbol{J}_{\boldsymbol{S}}(\boldsymbol{z}_i)$ and the change of volume scalar $(\sigma_1, \ldots, \sigma_N) \triangleq \left(\sqrt{\det(\boldsymbol{A}_1^T \boldsymbol{A}_1)}, \ldots, \sqrt{\det(\boldsymbol{A}_N^T \boldsymbol{A}_N)}\right)$,
3. sample (with replacement) K latent vectors $(\boldsymbol{z}_1, \ldots, \boldsymbol{z}_K)$ with probability $\propto (\sigma_1, \ldots, \sigma_N)$

We discuss the possible choices of N and K in Sec. D in the Appendix where we observe that even for state-of-the-art models like StyleGAN2, N =200,000 is sufficient to provide a stable approximation of the true latent space target distribution.

The above provides a Monte-Carlo approximation that does not require knowledge of the partition Ω nor the per-region slope matrices. Those are computed on-demand as $\boldsymbol{z}_i$ are sampled The above procedure produces uniform samples on the manifold learned by a DGN regardless of how it has been trained.

4 Uniform Sampling from Deep Generative Networks

The goal of this section is to first bridge current DGNs with affine splines to then leverage Thm. 1 and Sec. 3.2 in order to effectively produce uniform samples on the manifold of current Deep Generative Networks such as BigGAN, StyleGAN. We build this affine splines-DGN bridge and propose carefully motivation for uniform sampling in Sec. 4.1 and finally present various experiments across architectures in Sec. 4.2, 4.3 and 4.4.

4.1 Uniform Sampling on the Deep Generative Network Manifold

We proposed in the previous section a thorough study of affine splines and how those mappings transform a given input distribution. This now takes high relevance as per the following remark.

Remark 1. *Any DGN that employs CPA nonlinearities (as in Sec. 2) is itself a CPA when considering the entire input-output mapping, or any sub-part of the architecture; that is, the input-output mapping can be expressed as in (Eq. 1).*

We shall remind that such observation in the context of classifier DNs goes back to Montufar et al. (2014) and has been further studied in Unser (2018); Balestriero & Baraniuk (2018b). We also shall emphasize that operators such as Batch-Normalization (Ioffe & Szegedy, 2015) are not continuous piecewise affine during training but become affine operators during evaluation time. For completeness we also provide that analytical form of the per-region affine mappings $\boldsymbol{A}_\omega, \boldsymbol{b}_\omega$ of Eq. 1 in the case of DGNs in Sec. C in the appendix. The key for our method is thus to combine the above with the results from Sec. 3.2 to obtain the following statement.

Theorem 2. *Consider a training set sampled from a true manifold $\mathcal{M}$ with distribution $p_{\boldsymbol{x}}$ (with support being $\mathcal{M}$), as long as the DGN approximation of $\mathcal{M}$ is valid, sampling from a trained DGN as per Sec. 3.2 produces uniform samples that do not depend on $p_{\boldsymbol{x}}$. (Proof in G.4.)*

By using the above *one can take any (possibly pre-trained) DGN regardless of the training setting, and produce uniform samples on the learned underlying manifold*. Hence, as long as the underlying manifold learned by the DGN is correct, our solution allows to produce a generative process that becomes invariant to the training set distribution. While this provides a theoretical guarantee for uniform sampling, it also highlight the main limitation of MaGNET: *the uniform samples will be lying on a continuous piecewise affine manifold*. That is, unless the true manifold $\mathcal{M}$ is also continuous, MaGNET will occasionally introduce abnormal samples that correspond to sampling from the regions of discontinuity of $\mathcal{M}$. We will see in the following sections how even on high-quality image datasets, MaGNET produces very few abnormal samples, one reason being that for complicated data manifold, state-of-the-art DGNs are often built with a (class) conditioning. That is, the above continuity assumption on $\mathcal{M}$ lessens only to a within-class continuity assumption which is a much more realistic assumption. Sampling uniformly on the DGN manifold has many important applications that are deferred to the following sections.

4.2 Quantitative Validation: ϵ-Ball Concentration and GMM Likelihood

We now propose two quantitative controlled experiments to validate the applicability of the theoretical results from Sec. 3.2 for the MaGNET sampling procedure.

First, we consider MNIST so that we can assume that the entire data manifold is approximately covered by the $50,000$ training samples. We have the following observation. Regardless of the

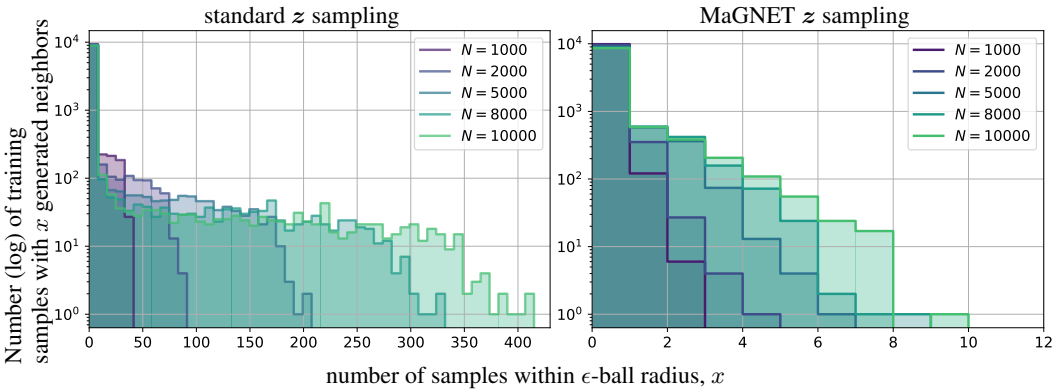

Figure 3: Distribution of the number of MNIST training samples with x neighbors generated within an ϵ-ball radius. N samples are generated using standard sampling and MaGNET sampling using an NVAE model. Here ϵ is taken to be the average nearest neighbor distance for the training samples. For Vanilla NVAE heavier tails are indicative of larger density variations on the manifold as N is increased, whereas for MaGNET the shorter tails are indicative of less variations in neighborhood density, i.e. uniform generation on the data manifold. $10k$ MNIST samples are used for comparison, for additional ϵ see Fig. 13 and Fig. 14 in the Appendix.

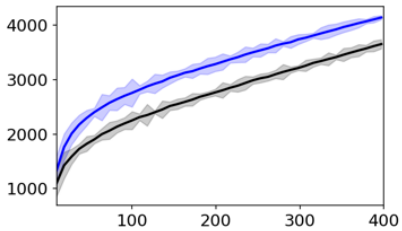

Figure 4: Depiction of the average over 5 runs (and 10σ-bandwidth) log-likelihood of a GMM with varying number of clusters (x-axis) trained on $1,000$ standard MNIST NVAE samples (black) and $1,000$ MaGNET MNIST NVAE samples (blue). Since the manifold and domain of those samples is the same for both cases the ability to reach a higher log-likelihood (given the same number of clusters) in the standard sampling case demonstrates an increased concentration around a few (or many) modes as opposed to the lower log-likelihood obtained with MaGNET which is due to the samples being more uniformly distributed.

training data distribution on the manifold (uniform or not), we can pick a datum at random, count how many generated samples are within this datum ϵ-ball and repeat this process for all training data. If this count does not vary, then it strongly indicates that the generated samples are uniformly distributed on the manifold covered by the training data. We do exactly this experiment using a pre-trained state of the art variational autoencoder NVAE (Vahdat & Kautz, 2020) to perform standard and MaGNET sampling with number of generated samples N ranging from $1,000$ to $10,000$. We report the distribution of counts in Fig. 3. Again, uniform sampling is equivalent to having the same count for all training samples i.e. a Dirac distribution in the reported histograms. We can see that MaGNET sampling approaches that distribution while standard sampling has a heavy-tail count distribution i.e. the generated digits have different ϵ-ball concentration at different places on the data manifold. Another quantitative measure consists in fitting a Gaussian Mixture Model (GMM) with varying number of clusters, and for each case, compare the likelihood obtained on the standard and on MaGNET sampling. As we know that in both cases the samples lie on the same manifold and domain, the sampling with lower likelihood will correspond to the one for which samples are spread more uniformly on the manifold. We report this in Fig. 4, further confirming the ability of MaGNET to produce uniformly spread samples. We report, for both sampling, the generated samples in Appendix E.

4.3 Qualitative Validation: High-Dimensional State-of-the-art Image Generation

We now turn into the qualitative evaluation of MaGNET sampling, to do so we propose extensive experiments on various state-of-the-art image DGNs. We also remind the reader that in all cases, standard and MaGNET sampling are performed on the same DGN (same weights) as discussed in Sec. 3.2.

2-Dimensional Dataset and Colored-MNIST. The first application of MaGNET comes when the DGN is employed on manifolds where it is known that the samples should be sampled uniformly

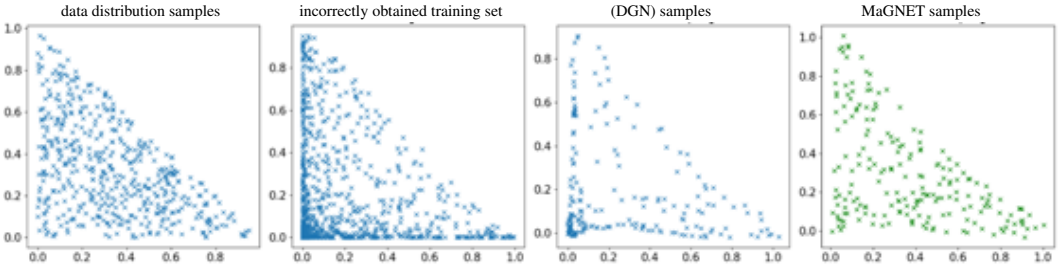

Figure 5: Depiction of employing a naive sampling strategy (third column) on a GAN that has been trained on a curated dataset (first column) obtained by oversampling points near the origin (second column). After training, the GAN learned the training set distribution, which in this case would lead to incorrect data generation and exploration. In contrast, our proposed uniform sampling methodology MaGNET from Sec. 3.2 recovers the uniform data sampling on the support as followed by the true data density (fourth column). Note that the same number of samples are used for the third (standard) and fourth (MaGNET) sampling strategies.

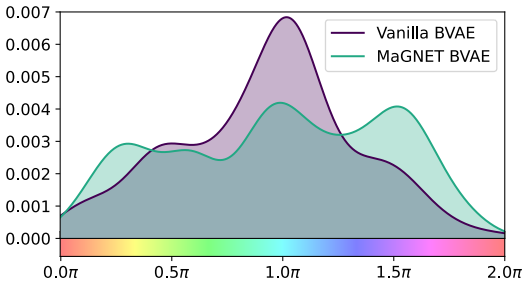

Figure 6: Depiction of the hue (color) distribution of samples obtained from standard and MaGNET sampling from a trained BetaVAE model. Original dataset was purposefully favoring cyan coloring to represent training set inconsistency, uniform sampling on the true color MNIST manifold should recover uniform hue distribution, MaGNET is able to approximately recover uniform coloring. The remaining gaps around the red color represent regions where the DGN manifold approximation was incorrect, hence uniform sampling can not recover those samples.

but the training set contained inconsistencies. Such inconsistencies can be due to cost constraints or challenges related to dataset compilation. We provide an illustrative examples in Fig. 5 where we demonstrate that unless uniform sampling is employed, the trained DGN reproduces the inconsistencies present in the training set, as expected. This toy 2-dimensional visualization validates our method from Sec. 3.2.

Going further, we propose to take the MNIST dataset (in this case, only considering the digit 8) and apply imbalanced coloring based on the hue distribution provided in Fig. 12 in the Appendix that favors cyan color. We then trained on that cyan-inclined dataset a β-VAE DGN. Once trained, we compare standard sampling against MaGNET sampling of that same DGN to obtain the sampled dataset hue distributions that we depict in Fig. 6. We observe that MaGNET corrects the hue distribution back to uniformity. More details on how to guarantee that uniform sampling implies uniform attributes in Sec. 4.4. We now turn to high-quality image generation.

Uniform Face Generation: CelebA-HQ and Flickr-Faces-HQ with progGAN and StyleGAN2. The first experiment that we propose concerns sampling from the StyleGAN2 (Karras et al., 2020) pre-trained on the Flickr-Faces-HQ (FFHQ) dataset. StyleGAN2 has two DGNs, one that maps to an intermediate latent space, termed style-space and another DGN that maps style-space vectors to the pixels-space (output of StyleGAN2). Implementation details are deferred to Appendix F. We focus here on applying MaGNET onto the entire StyleGAN2 model (the composition of both DGNs), Sec. 4.4 will propose to apply MaGNET on the style-space DGN only. In Fig. 1 & Appendix Fig. 15 we provide random samples from standard and MaGNET sampling. Upon qualitative evaluation it can be seen that the MaGNET batch of faces have a significantly larger variety of age distribution, background variations and wearable accessories compared to standard sampling. For experiments with CelebA-HQ dataset, we adopt the DGN of Progressively Growing GAN (progGAN) (Karras et al., 2017), trained on 1024×1024 resolution images. In Fig. 9 & Appendix Fig. 16 we provide random samples from standard and MaGNET sampling, the latter are qualitatively more diverse than standard sampling. In particular, we see that uniform manifold sampling via MaGNET recovers samples containing a large number of attributes that are generally underrepresented in the vanilla DGN distribution. To quantify the diversity induced by MaGNET, we use Microsoft Azure Face API to tag attributes for $5,000$ samples for different sampling methods and report those in

Figure 7: Random batch of samples generated from BigGAN (left) and MaGNET BigGAN (right), conditioned on the *Samoyed* class of ImageNet. While BigGAN samples contain homogeneous postures, MaGNET samples represent the true span of the data manifold learned by BigGAN.

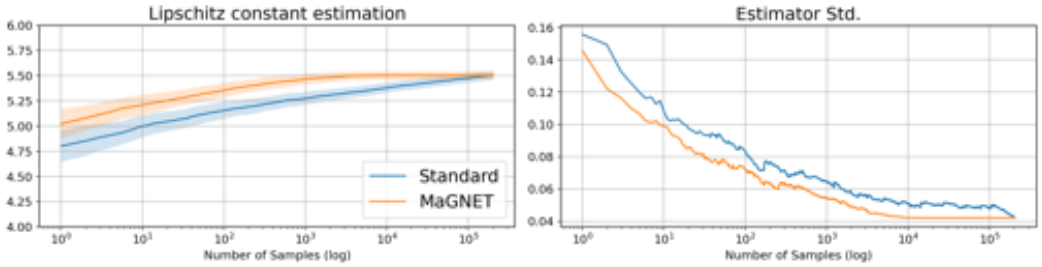

Figure 8: Depiction of Lipschitz constant estimation using standard DGN sampling and the proposed uniform sampling using MaGNET (**left**). Each line represents the mean over 200 Monte-Carlo runs. As expected, estimation of such statistics from samples converges faster when employed on uniformly distributed samples. The standard deviation of the Monte-Carlo estimations are also provided (**right**) where it is clear that the uniform sampling via MaGNET reaches smaller standard deviation at earlier steps i.e. is faster to converge. This is a key application of MaGNET: speeding-up convergence of statistical estimation of quantities, such as the Lipschitz constant in this case. The x-axis represents the number of samples used for estimation in log-scale.

Fig. 22. Notice that uniform sampling not only recovers under-represented groups e.g. age < 30, head-wear, and bald hair, it also increases the presence of neutral emotion and black hair.

Conditionally Uniform Generation: ImageNet with BigGAN. We propose experiments on the state-of-the-art conditional generative model BigGAN (Brock et al., 2019) using MaGNET sampling. In Fig. 7 we provide random samples from standard and MaGNET sampling. More experiments on different classes are presented in Appendix E. We see that uniform sampling on the learned data manifold yields a large span of backgrounds and textures, including humans. This is quite understandable considering that ImageNET was curated using a large number of images scraped from the internet. While standard sampling shows a small number of modes that were present during training. MaGNET therefore could possibly be used for more diverse data exploration/model interpretation and also as a diagnostic tool to assess the quality of the learned manifold a posteriori of training.

4.4 Application: Monte-Carlo Estimation and Attribute Rebalancing

We conclude this section with two more practical aspects offered by MaGNET.

Reduced-Variance Monte-Carlo Estimator. The first is to speed-up (in term of number of required samples) basic Monte-Carlo estimation of arbitrary topological quantities of the generated manifold. In fact, suppose one's goal is to estimate the Lipschitz constant of a DGN. It is clear that a direct estimation method consists of using the known bound given by $\max_{\boldsymbol{z}} \|\boldsymbol{J}_{\boldsymbol{S}}(\boldsymbol{z})\|_F$ (Wood & Zhang, 1996). This estimation can be done by repeatedly sampling latent vectors $\boldsymbol{z}$ from the same distribution that one used for training a DGN. However this implies that the produced samples will not be uniformly distributed on the manifold in turn leading to slower convergence of the estimator. Instead, we propose to use MaGNET, and report our findings in Fig. 8. Further practical applications of MaGNET to reduce estimators variance can eb found in Baggenstoss (2017).

Style-space MaGNET sampling rebalances attributes. When thinking of uniform sampling on a manifold it might seem natural to expect fairness i.e. fair representation of different attributes such as equal representation of gender, ethnicity, hair color, etc. This is however not necessarily true in all cases. In fact, it is trivial to show that each category will be equally represented iff their support on the true data manifold is of equal volume (integrated with respect to the data manifold).

Figure 9: Random batch of samples generated from vanilla progGAN (**left**) and MaGNET progGAN (**right**). Some samples contain fusion of attributes which are not frequently represented in the data distribution, e.g. female w/ facial hair, male w/ makeup, etc. Larger batches in Appendix E.

Fortunately, as we mentioned in Sec. 4.3, some architectures have explicitly built a style-space as with StyleGAN2. This style-space is a latent space in which attributes are better organized along affine subspaces and occupy nearly the same volume (Karras et al., 2019) i.e. MaGNET applied on the style-space DGN should improve fairness. By applying MaGNET sampling on the style-space, we are able to reduce gender bias from 67–33% (female-male) in standard StyleGAN2 to 60–40%. This simple result demonstrates the importance of our proposed sampling and how it can be used to increase fairness for DGNs trained on biased training sets. MaGNET in the style-space also yields improvements in terms of recall and precision (Sajjadi et al., 2018). Given a reference distribution (e.g. FFHQ dataset) and a learned distribution, precision measures the fidelity of samples from the learned distribution while recall measures diversity via the proportion of the reference distribution that is covered by the learned distribution. We compare the metrics for face images generated via latents $z \sim N(0, aI)$ where $a \in 0.5, 1, 1.5, 2$, $z \sim U[-2, 2]$ and MaGNET sampling on style-space. For 70K samples generated for each case, MaGNET sampling obtains a recall and precision of $(0.822, 0.92)$ with a 4.12% relative increase in recall and 3.01% relative increase in precision compared to other latent $\boldsymbol{m}z$ sampling methods (metrics are averaged for 10 seeds).

5 Conclusions, Limitations and Future Works

We have demonstrated how the affine spline formulation of DGNs offers a rich portal to obtain theoretical results to provably provide uniform sampling on the manifold learned by a DGN. This allows to become robust to possibly incorrect training set distribution that any DGN would learn to replicate after during its training. We provided various experiments using pretrained state-of-the-art generative models and demonstrated that uniform sampling on the manifold offers many benefits from data exploration to statistical estimation. Beyond the sole goal of uniform sampling on a manifold, MaGNET opened many avenues, yet MaGNET is not a "fit-all" solution.

When to not sample uniformly. We can identify to general cases in which one should not employ uniform sampling of the DGN manifold. The first case occurs whenever the true manifold is known to be discontinuous and one imperatively needs to avoid sampling in those regions of discontinuities. In fact, in the discontinuous case, DGN training, will adapt its weight to put zero (or near zero) density in those discontinuous regions preventing standard sampling to reach those regions (Balestriero et al., 2020b). However, MaGNET will reverse this process and introduce samples back in those regions. The second case occurs if one aims to produce samples from the exact same distribution as the training set distribution (assuming training of the DGN was successful). In fact, in this scenario one should use the exact same latent distribution at evaluation time than the one used during training.

Future works Currently, there are two main limitations of our MaGNET sampling strategy. The first one lies in the assumption that the trained DGN was able to learn a good enough approximation of the true underlying data manifold. In future work, we plan to explore how MaGNET could be used to test such an assumption. One fruitful directions is as follows. Train a DGN using several different sub-sampled datasets (similar to bootstrap methods) and then study if MaGNET samples populate manifolds that all coincide between the different DGNs. If training was successful, those sampled manifolds should coincide. The second one lies in understanding the relation between uniform sampling and uniform attribute representation. We demonstrated how uniform sampling in the style-space of StyleGAN2 does ensures that relationship -by construction-. However, doing MaGNET sampling arbitrarily on any DGN might not always produce better distribution attributes. Hence, developing novel theoretical results to characterize those two effects is one crucial future direction to explore to equip MaGNET with theoretical fairness guarantees.

6 REPRODUCIBILITY STATEMENT

Reproducible code for the various experiments and figures will be provided upon completion of the review process. Computation and software details are provided in Appendix F, with the proofs of our results in Appendix G. Discussion on the settings in which MaGNET is desirable and on possible limitations is provided in (Sec. 5).

REFERENCES

Helena Andrés-Terré and Pietro Lió. Perturbation theory approach to study the latent space degeneracy of variational autoencoders, 2019.

M Arjovsky and L Bottou. Towards principled methods for training generative adversarial networks. arxiv, 2017. *arXiv preprint arXiv:1701.04862*.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.

Paul M Baggenstoss. Uniform manifold sampling (ums): Sampling the maximum entropy pdf. *IEEE Transactions on Signal Processing*, 65(9):2455–2470, 2017.

R. Balestriero and R. Baraniuk. Mad max: Affine spline insights into deep learning. *arXiv preprint arXiv:1805.06576*, 2018a.

R. Balestriero and R. G. Baraniuk. A spline theory of deep networks. In *Proc. Int. Conf. Mach. Learn.*, volume 80, pp. 374–383, Jul. 2018b.

Randall Balestriero and Richard G Baraniuk. Mad max: Affine spline insights into deep learning. *Proceedings of the IEEE*, 2020.

Randall Balestriero, Romain Cosentino, Behnaam Aazhang, and Richard Baraniuk. The geometry of deep networks: Power diagram subdivision. In *Advances in Neural Information Processing Systems 32*, pp. 15806–15815. 2019.

Randall Balestriero, Sébastien Paris, and Richard Baraniuk. Analytical probability distributions and exact expectation-maximization for deep generative networks. *Advances in Neural Information Processing Systems*, 33, 2020a.

Randall Balestriero, Sebastien Paris, and Richard Baraniuk. Max-affine spline insights into deep generative networks. *arXiv preprint arXiv:2002.11912*, 2020b.

Matan Ben-Yosef and Daphna Weinshall. Gaussian mixture generative adversarial networks for diverse datasets, and the unsupervised clustering of images. *arXiv preprint arXiv:1808.10356*, 2018.

Rianne van den Berg, Leonard Hasenclever, Jakub M Tomczak, and Max Welling. Sylvester normalizing flows for variational inference. *arXiv preprint arXiv:1803.05649*, 2018.

Gérard Biau, Benoît Cadre, MAXIME Sangnier, and Ugo Tanielian. Some theoretical properties of gans. *arXiv preprint arXiv:1803.07819*, 2018.

Leo Breiman, William Meisel, and Edward Purcell. Variable kernel estimates of multivariate densities. *Technometrics*, 19(2):135–144, 1977.

Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *In Proc. ICLR*, 2019.

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.

Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in Neural Information Processing Systems*, 31: 2610–2620, 2018.

Elliott Ward Cheney and William Allan Light. *A course in approximation theory*, volume 101. American Mathematical Soc., 2009.

Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018.

Laurent Dinh, David Krueger, and Yoshua Bengio. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516*, 2014.

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.

Ishan Durugkar, Ian Gemp, and Sridhar Mahadevan. Generative multi-adversarial networks. *arXiv preprint arXiv:1611.01673*, 2016.

Gintare Karolina Dziugaite, Daniel M Roy, and Zoubin Ghahramani. Training generative neural networks via maximum mean discrepancy optimization. *arXiv preprint arXiv:1505.03906*, 2015.

Magnus Egerstedt and Clyde Martin. *Control theoretic splines: optimal control, statistics, and path planning*. Princeton University Press, 2009.

Otto Fabius and Joost R van Amersfoort. Variational recurrent auto-encoders. *arXiv preprint arXiv:1412.6581*, 2014.

Cesare Fantuzzi, Silvio Simani, Sergio Beghelli, and Riccardo Rovatti. Identification of piecewise affine models in noisy environment. *International Journal of Control*, 75(18):1472–1485, 2002.

I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*, volume 1. MIT Press, 2016.

I. J Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, pp. 2672–2680. MIT Press, 2014.

Will Grathwohl, Ricky TQ Chen, Jesse Betterncourt, Ilya Sutskever, and David Duvenaud. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*, 2018.

Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017.

L. A. Hannah and D. B. Dunson. Multivariate convex regression with adaptive partitioning. *J. Mach. Learn. Res.*, 14(1):3261–3294, 2013.

Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2(5):6, 2017.

Sunhee Hwang, Sungho Park, Dohyung Kim, Mirae Do, and Hyeran Byun. Fairfacegan: Fairness-aware facial image-to-image translation. In *2020 31st British Machine Vision Conference (BMVC)*, pp. 1309–1320, 2020.

S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

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.

Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.

Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proc. CVPR*, pp. 8110–8119, 2020.

D. P Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pp. 10215–10224, 2018.

Naveen Kodali, Jacob Abernethy, James Hays, and Zsolt Kira. On convergence and stability of gans. *arXiv preprint arXiv:1705.07215*, 2017.

Jerry Li, Aleksander Madry, John Peebles, and Ludwig Schmidt. Towards understanding the dynamics of generative adversarial networks. *arXiv preprint arXiv:1706.09884*, 2017.

Shuang Liu, Olivier Bousquet, and Kamalika Chaudhuri. Approximation and convergence properties of generative adversarial learning. In *Advances in Neural Information Processing Systems*, pp. 5545–5553, 2017.

A. Magnani and S. P. Boyd. Convex piecewise-linear fitting. *Optim. Eng.*, 10(1):1–17, 2009.

Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2794–2802, 2017.

Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.

Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.

Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. In *Advances in Neural Information Processing systems*, pp. 2924–2932, 2014.

Gabriel Peyré. Manifold models for signals and images. *Computer vision and image understanding*, 113(2):249–260, 2009.

Murray Rosenblatt. Remarks on some nonparametric estimates of a density function. *The Annals of Mathematical Statistics*, pp. 832–837, 1956.

Aurko Roy, Ashish Vaswani, Arvind Neelakantan, and Niki Parmar. Theory and experiments on vector quantized autoencoders. *arXiv preprint arXiv:1805.11063*, 2018.

Walter Rudin. *Real and Complex Analysis*. Tata McGraw-hill education, 2006.

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.

Veit Sandfort, Ke Yan, Perry J Pickhardt, and Ronald M Summers. Data augmentation using generative adversarial networks (cyclegan) to improve generalizability in ct segmentation tasks. *Scientific reports*, 9(1):1–9, 2019.

Kevin Scaman and Aladin Virmaux. Lipschitz regularity of deep neural networks: analysis and efficient estimation. *arXiv preprint arXiv:1805.10965*, 2018.

Michael Spivak. *Calculus on manifolds: a modern approach to classical theorems of advanced calculus*. CRC press, 2018.

Shuhan Tan, Yujun Shen, and Bolei Zhou. Improving the fairness of deep generative models without retraining. *arXiv preprint arXiv:2012.04842*, 2020.

Akinori Tanaka. Discriminator optimal transport. *arXiv preprint arXiv:1910.06832*, 2019.

Saravanan Thirumuruganathan, Shohedul Hasan, Nick Koudas, and Gautam Das. Approximate query processing for data exploration using deep generative models. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*, pp. 1309–1320. IEEE, 2020.

Jakub M Tomczak and Max Welling. Improving variational auto-encoders using householder flow. *arXiv preprint arXiv:1611.09630*, 2016.

Jakub M Tomczak and Max Welling. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*, 2017.

Michael Unser. A representer theorem for deep neural networks. *arXiv preprint arXiv:1802.09210*, 2018.

Arash Vahdat and Jan Kautz. Nvae: A deep hierarchical variational autoencoder. *arXiv preprint arXiv:2007.03898*, 2020.

Aaron van den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.

Shuning Wang and Xusheng Sun. Generalization of hinging hyperplanes. *IEEE Transactions on Information Theory*, 51(12):4425–4431, 2005.

GR Wood and BP Zhang. Estimation of the lipschitz constant of a function. *Journal of Global Optimization*, 8(1):91–103, 1996.

Jiacheng Xu and Greg Durrett. Spherical latent spaces for stable variational autoencoders. *arXiv preprint arXiv:1808.10805*, 2018.

Dingdong Yang, Seunghoon Hong, Yunseok Jang, Tianchen Zhao, and Honglak Lee. Diversity-sensitive conditional generative adversarial networks, 2019.

Pengchuan Zhang, Qiang Liu, Dengyong Zhou, Tao Xu, and Xiaodong He. On the discrimination-generalization tradeoff in gans. *arXiv preprint arXiv:1711.02771*, 2017.

Junbo Zhao, Michael Mathieu, and Yann LeCun. Energy-based generative adversarial network. *arXiv preprint arXiv:1609.03126*, 2016.

SUPPLEMENTARY MATERIAL

The following appendices support the main paper and are organized as follows: Appendix A addresses some details on the background of continuous piecewise affine deep networks, that were omitted in the main paper. Appendices E and F provide additional figures and training details for all the experiments that were studied in the main text, and Appendix G provides the proofs of all the theoretical results.

A BACKGROUND ON CONTINUOUS PIECEWISE AFFINE DEEP NETWORKS

A *max-affine spline operator* (MASO) concatenates independent *max-affine spline* (MAS) functions, with each MAS formed from the point-wise maximum of R affine mappings (Magnani & Boyd, 2009; Hannah & Dunson, 2013). For our purpose each MASO will express a DN layer and is thus an operator producing a D^ℓ dimensional vector from a $D^{\ell-1}$ dimensional vector and is formally given by

$$\text{MASO}(\boldsymbol{v}; \{\boldsymbol{A}_r, \boldsymbol{b}_r\}_{r=1}^R) = \max_{r=1,\ldots,R} \boldsymbol{A}_r \boldsymbol{v} + \boldsymbol{b}_r, \tag{5}$$

where $\boldsymbol{A}_r \in \mathbb{R}^{D^\ell \times D^{\ell-1}}$ are the slopes and $\boldsymbol{b}_r \in \mathbb{R}^{D^\ell}$ are the offset/bias parameters and the maximum is taken coordinate-wise. For example, a layer comprising a fully connected operator with weights $\boldsymbol{W}^\ell$ and biases $\boldsymbol{b}^\ell$ followed by a ReLU activation operator corresponds to a (single) MASO with $R = 2, \boldsymbol{A}_1 = \boldsymbol{W}^\ell, \boldsymbol{A}_2 = \boldsymbol{0}, \boldsymbol{b}_1 = \boldsymbol{b}^\ell, \boldsymbol{b}_2 = \boldsymbol{0}$. Note that a MASO is a *continuous piecewise-affine* (CPA) operator (Wang & Sun, 2005).

The key background result for this paper is that *the layers of DNs constructed from piecewise affine operators (e.g., convolution, ReLU, and max-pooling) are MASOs* (Balestriero & Baraniuk, 2018b;a):

$$\exists R \in \mathbb{N}^*, \exists \{\boldsymbol{A}_r, \boldsymbol{b}_r\}_{r=1}^R \text{ s.t. } \text{MASO}(\boldsymbol{v}; \{\boldsymbol{A}_r, \boldsymbol{b}_r\}_{r=1}^R) = g^\ell(\boldsymbol{v}), \forall \boldsymbol{v} \in \mathbb{R}^{D^{\ell-1}}, \tag{6}$$

making the entire DGN a composition of MASOs.
The CPA spline interpretation enabled from a MASO formulation of DGNs provides a powerful global geometric interpretation of the network mapping based on a partition of its input space $\mathbb{R}^S$ into polyhedral regions and a per-region affine transformation producing the network output. The partition regions are built up over the layers via a *subdivision* process and are closely related to Voronoi and power diagrams (Balestriero et al., 2019). We now propose to greatly extend such insights to carefully characterize and understand DGNs as well as provide theoretical justifications to various observed behaviors e.g. mode collapse.

B UNIFORM AND GAUSSIAN MANIFOLD DISTRIBUTIONS

We now demonstrate the use of the above result by considering practical examples for which we are able to gain insights into the DGN data modeling and generation. We consider the two most common cases: (i) the latent distribution is set as $\boldsymbol{z} \sim \mathcal{N}(0, 1)$ and (ii) the latent distribution is set as $\boldsymbol{z} \sim \mathcal{U}(0, 1)$ (on the hypercube of dimension S). We obtain the following result by direct application of Thm. 1.

Corollary 1. *The generated density distribution $p_{\boldsymbol{S}}$ of the Gaussian and uniform densities are given by*

$$p_{\boldsymbol{S}}(\boldsymbol{x}) = \sum_{\omega \in \Omega} \frac{e^{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{b}_\omega)^T (\boldsymbol{A}_\omega^+)^T \boldsymbol{A}_\omega^+ (\boldsymbol{x}-\boldsymbol{b}_\omega)}}{\sqrt{(2\pi)^S \det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} \mathbb{1}_{\{\boldsymbol{x} \in \boldsymbol{G}(\omega)\}}, \qquad \textit{(Gaussian)} \tag{7}$$

$$p_{\boldsymbol{S}}(\boldsymbol{x}) = \sum_{\omega \in \Omega} \frac{1}{\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} \mathbb{1}_{\{\boldsymbol{x} \in \boldsymbol{S}(\omega)\}}. \qquad \textit{(Uniform)} \tag{8}$$

The two above formulae provide a precise description of the produced density given that the latent space density is Gaussian or Uniform. In the Gaussian case, the per-region slope matrices act upon

the ℓ_2 distance by rescaling it from the coordinates of $\boldsymbol{A}_\omega$ and the per-region offset parameters $\boldsymbol{b}_\omega$ are the mean against which the input $\boldsymbol{x}$ is compared against. In the Uniform case, the change of volume (recall Eq. 3) is the only quantity that impacts the produced density. We will heavily rely on this observation for the next section where we study how to produce a uniform sampling onto the CPA manifold of an affine spline.

We derive the analytical form for the case of Gaussian and Uniform latent distribution in Appendix G.3. From the analytical derivation of the generator density distribution, we obtain its differential entropy.

Corollary 2. *The differential Shannon entropy of the output distribution $\boldsymbol{p_G}$ of the DGN is given by* $E(\boldsymbol{p_G}) = E(\boldsymbol{p_z}) + \sum_{\omega\in\Omega} P(\boldsymbol{z} \in \omega) \log(\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)})$.

As a result, the differential entropy of the output distribution $\boldsymbol{p_G}$ corresponds to the differential entropy of the latent distribution $\boldsymbol{p_z}$ plus a convex combination of the per-region volume changes. It is thus possible to optimize the latent distribution $\boldsymbol{p_z}$ to better fit the target distribution entropy as in Ben-Yosef & Weinshall (2018) and whenever the prior distribution is fixed, any gap between the latent and output distribution entropy imply the need for high change in volumes between ω and $\boldsymbol{G}(\omega)$.

C Per-Region Affine Mappings

For completeness we also provide that analytical form of the per-region affine mappings

$$\boldsymbol{A}_\omega = \left(\prod_{i=0}^{L-1} \text{diag}\left(\dot{\sigma}^{L-i}(\omega)\right) \boldsymbol{W}^{L-i} \right), \tag{9}$$

$$\boldsymbol{b}_\omega = \boldsymbol{b}^L + \sum_{\ell=1}^{L-1} \left[\left(\prod_{i=0}^{L-\ell-1} \text{diag}\left(\dot{\sigma}^{L-i}(\omega)\right) \boldsymbol{W}^{L-i} \right) \text{diag}\left(\dot{\sigma}^{\ell}(\omega)\right) \boldsymbol{b}^\ell \right], \tag{10}$$

where $\dot{\sigma}^\ell(\boldsymbol{z})$ is the pointwise derivative of the activation function of layer ℓ based on its input $\boldsymbol{W}^\ell \boldsymbol{z}^{\ell-1} + \boldsymbol{b}^\ell$, which we note as a function of $\boldsymbol{z}$ directly. For precise definitions of those operators see Balestriero & Baraniuk (2020). The diag operator simply puts the given vector into a diagonal square matrix. For convolutional layers (or else) one can simply replace the corresponding $\boldsymbol{W}^\ell$ with the correct slope matrix parametrization as discussed in Sec. 2. Notice that since the employed activation functions $\sigma^\ell, \forall \ell \in \{1, \ldots, L\}$ are piecewise affine, their derivative is piecewise constant, in particular with values $[\dot{\sigma}^\ell(\boldsymbol{z})]_k \in \{\alpha, 1\}$ with $\alpha = 0$ for ReLU, $\alpha = -1$ for absolute value, and in general with $\alpha > 0$ for Leaky-ReLU for $k \in \{1, \ldots, D^\ell\}$.

D Number of Samples and Uniformity

Exact uniformity is reached when the Monte Carlo samples have covered each region of the DGN partition boundary. For large state-of-the-art models this condition requires sampling on the order of millions. However, we conducted an experiment to see how the number of samples really impacted the uniformity of the generated manifold as follows. We compute precision and recall metrics [4] for StyleGAN2 with K generated samples obtained from N Monte Carlo samples based on our sampling strategy by varying N. We use $K = 5000$ and N ranging from $10,000$ to $500,000$. Based on the metrics, we identify that increasing beyond $K = 250,000$ no longer impacts the metrics, showing that this number of monte carlo samples is enough to converge (approximately) to the uniform sampling in that case, see Fig. 10.

E Extra Figures

This section contains samples from our proposed methods, more samples along with attribute data and pretrained weights are available at our project `link`.

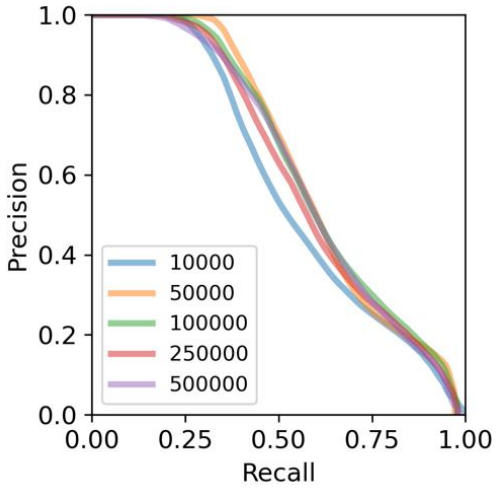

Figure 10: Evolution of the precision/recall curves for varying number of samples N form the monte-carlo sampling against the number of samples $K = 5k$ for StyleGAN2.

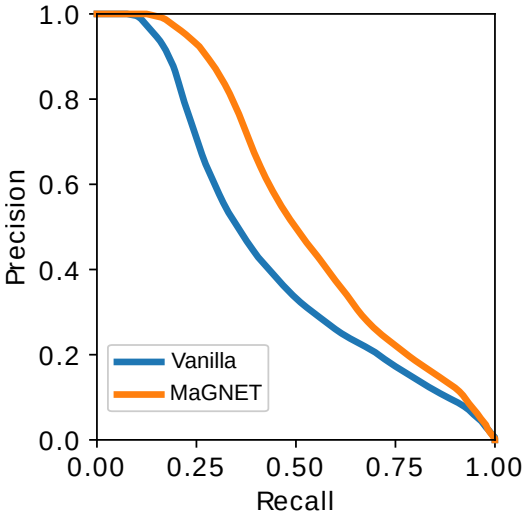

Figure 11: Precision-recall curves for $K = 70k$ samples from Vanilla StyleGAN2 and MaGNET StyleGAN2

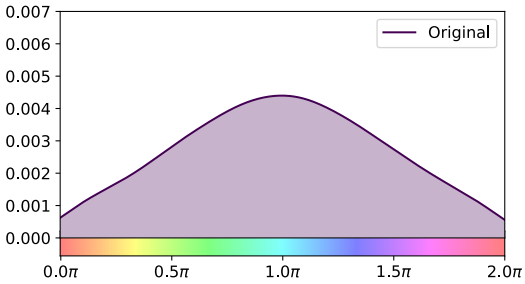

Figure 12: Depiction of the imbalance hue distribution applied to color the MNIST digits.

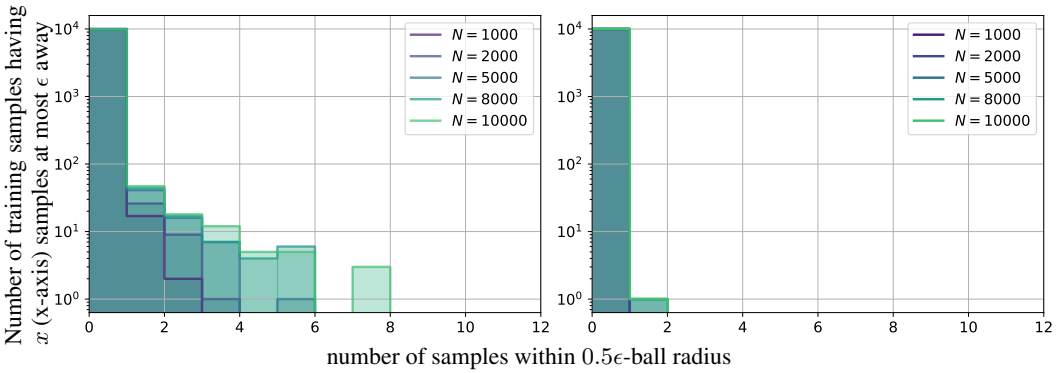

Figure 13: Reprise of Fig. 3. Vanilla NVAE **Left**, MaGNET NVAE **Right**

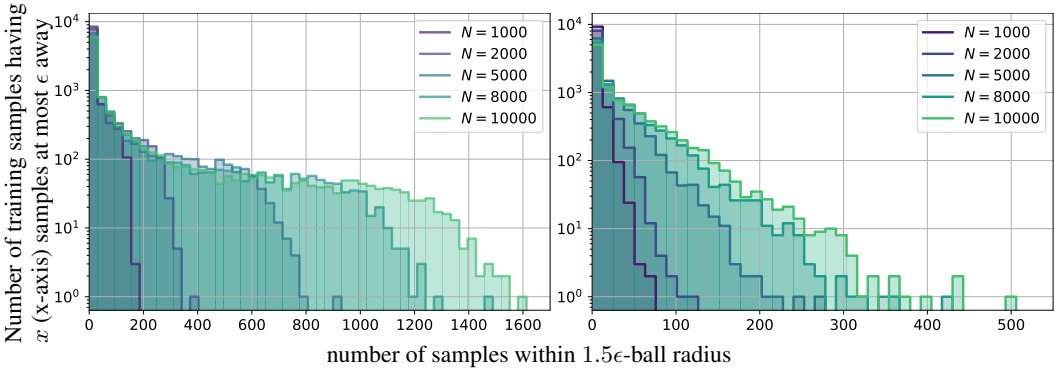

Figure 14: Reprise of Fig. 3. Vanilla NVAE **Left**, MaGNET NVAE **Right**

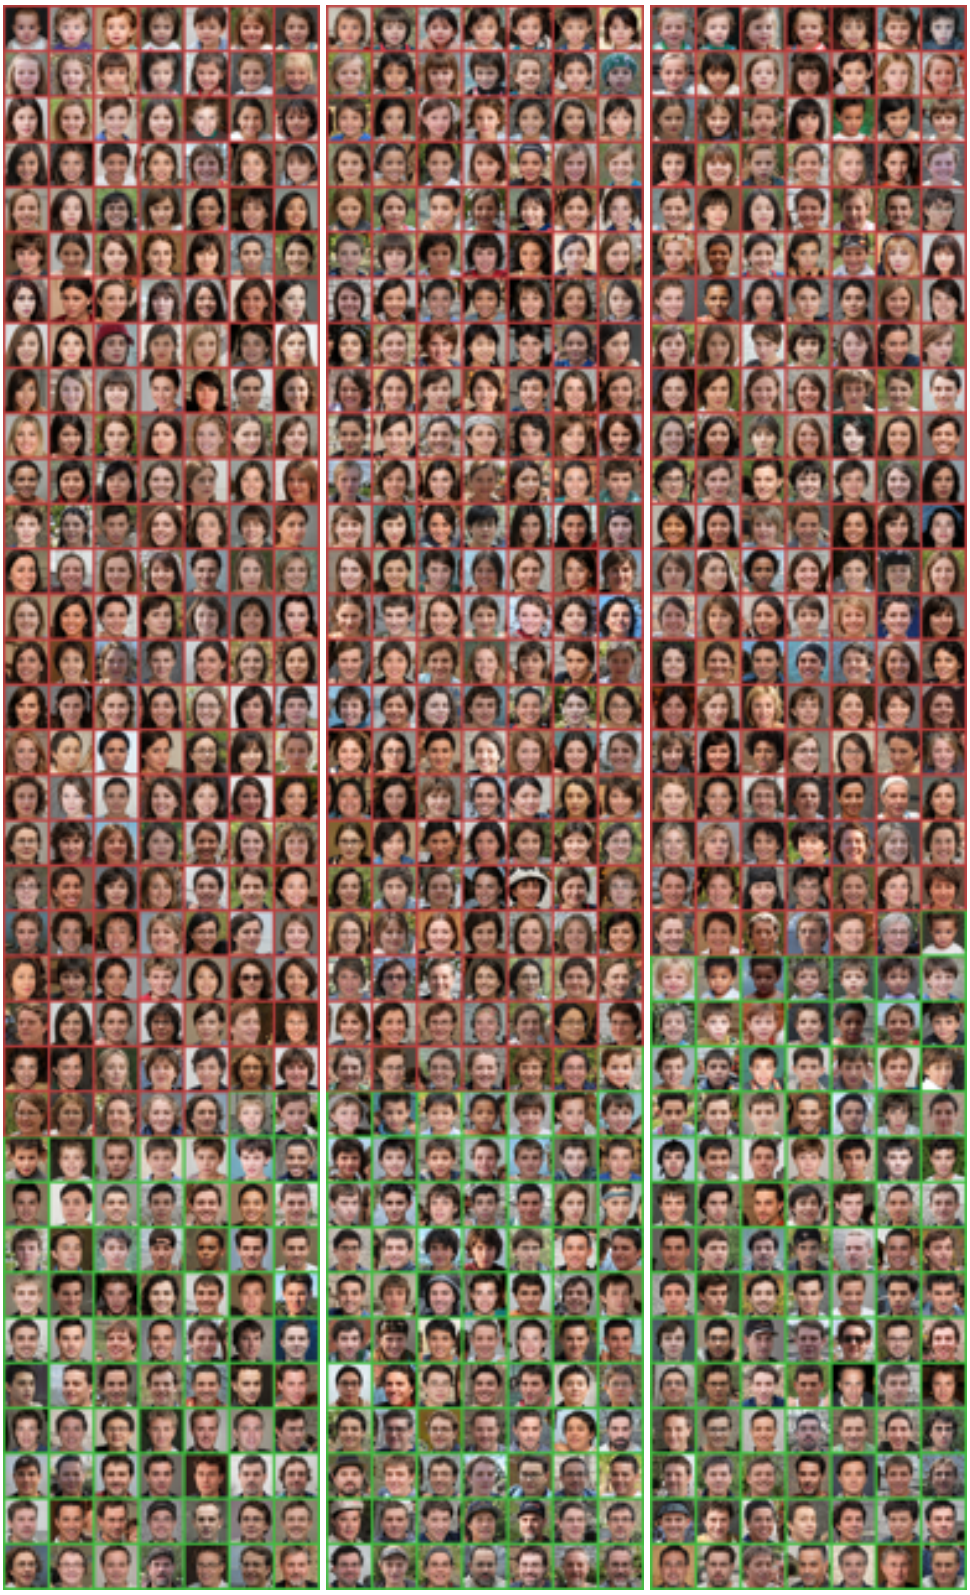

Figure 15: Random batches of 245 samples from a StyleGAN2 trained on FFHQ, generated via standard sampling (left), MaGNET sampling in the pixel-space (middle) and MaGNET sampling in the style-space.

Figure 16: Random Samples from vanilla progGAN (left) and MaGNET progGAN (right) trained on the CelebA-HQ dataset. Samples are sorted by gender & age and color coded by gender as visually predicted by the Microsoft Cognitive API. Samples not recognized by the API are color coded as white at the bottom.

Figure 17: Random Samples from vanilla BigGAN (left) and MaGNET BigGAN (right) from the *Collie* class of Imagenet.

Figure 18: Random Samples from vanilla BigGAN (left) and MaGNET BigGAN (right) from the *Wolf* class of Imagenet.

Figure 19: Random Samples from vanilla BigGAN (left) and MaGNET BigGAN (right) from the *Siamese* class of Imagenet.

Figure 20: Random Samples from vanilla BigGAN (left) and MaGNET BigGAN (right) from the *Tabby* class of Imagenet.

Figure 21: Random Samples from vanilla BigGAN (left) and MaGNET BigGAN (right) from the *Tiger* class of Imagenet.

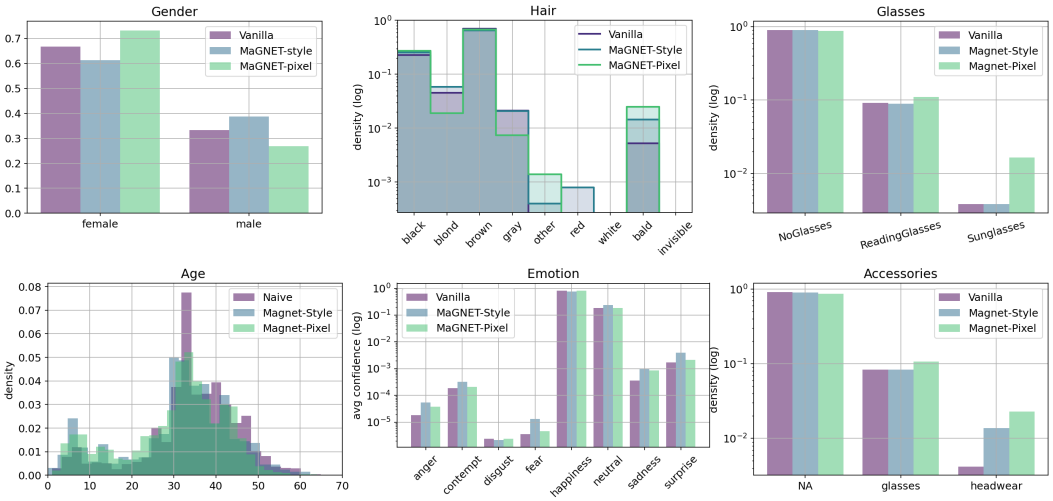

Figure 22: Facial Attributes of 5000 StyleGAN2 samples using vanilla sampling, MaGNET style-space sammpling and MaGNET pixel-space sampling. We see that MaGNET style-space increases uniformity in gender and age distributins whereas MaGNET pixel-space yields more variations in physical attributes and accessories.

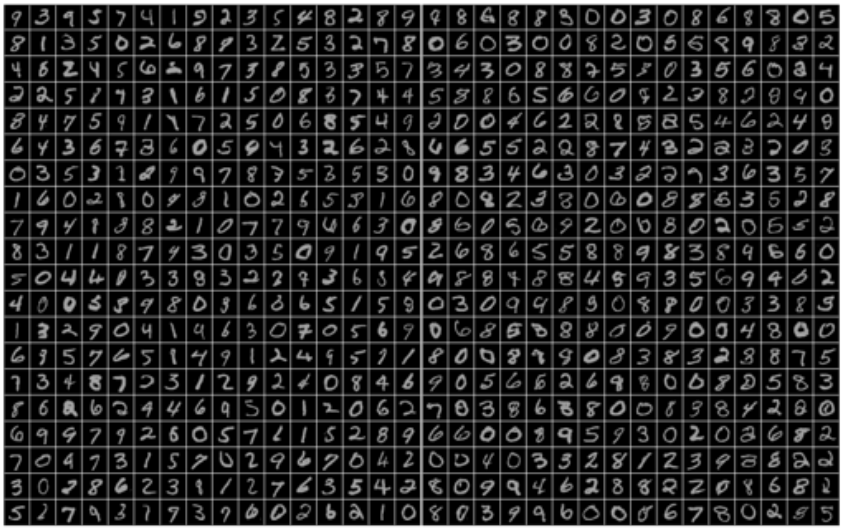

Figure 23: Random Samples from vanilla NVAE (left) and MaGNET NVAE (right) trained on the MNIST dataset.

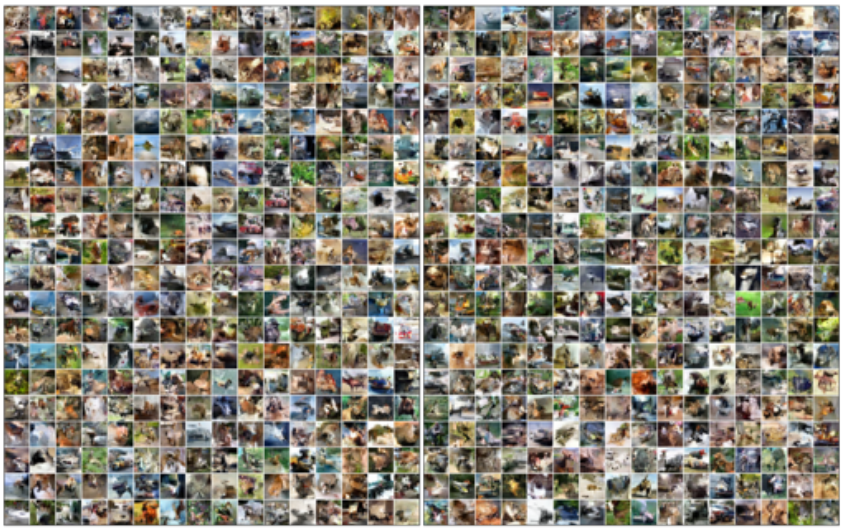

Figure 24: Random Samples from vanilla NVAE (left) and MaGNET NVAE (right) trained on the CIFAR dataset.

F Architecture and Hardware Details

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.

For styleGAN2 we use the official config-e of provided in the github StyleGAN2 repo. We use default interpolating truncation of 0.5 in the stylespace to ensure generation quality of faces. For BigGAN we use the biggan-deep arhcitecture with no truncation. We also use the official NVAE and progGAN models and weights from their github repos for the respective implementations. For the jacobian determinant of calulcation of images w.r.t latents we do a random projection into a lower dimensional space and take the top-k singular values for StyleGAN2-pixel, NVAE, progGAN and bigGAN.

G Proofs

G.1 Proof of Lemma 1

Proof. In the special case of an affine transform of the coordinate given by the matrix $A \in \mathbb{R}^{D \times D}$ the well known result from demonstrates that the change of volume is given by $|\det(A)|$ (see Theorem 7.26 in Rudin (2006)). However in our case the mapping is a rectangular matrix as we span an affine subspace in the ambient space $\mathbb{R}^D$ making $|\det(A)|$ not defined.

First, we shall note that in the case of a Riemannian manifold (as is the produced surface from the per-region affine mapping) the volume form used in the usual change of variable formula can be defined via the square root of the determinant of the metric tensor. Now, for a surface of intrinsic dimension n embedded in Euclidean space of dimension m (in our case, the per-region affine mapping produces an affine subspace) parametrized by the mapping $M : \mathbb{R}^n \mapsto \mathbb{R}^m$ (in our case this mapping is simply the affine mapping $M(\boldsymbol{z}) = \boldsymbol{z}_\omega z + \boldsymbol{b}_\omega$ for each region) the metric tensor is given by $g = DM^T DM$ with D the jacobian/differential operator (in our case $g = \boldsymbol{A}_\omega^T \boldsymbol{A}_\omega$ for each region). This result is also known as Sard's theorem (Spivak, 2018). We thus obtain that the change of volume from the region ω to the affine subspace $\boldsymbol{G}(\omega)$ is given by $\sqrt{\det(A^T A)}$ which can also

be written as follows with USV^T the svd-decomposition of the matrix A:

$$\begin{aligned}\sqrt{\det(A^TA)} = \sqrt{\det((USV^T)^T(USV^T))} =& \sqrt{\det((VS^TU^T)(USV^T))}\\ =& \sqrt{\det(VS^TSV^T)}\\ =& \sqrt{\det(S^TS)}\\ =& \prod_{i:\sigma_i\neq 0}\sigma_i(A)\end{aligned}$$

leading to

$$\int_{\text{Aff}(\omega,\boldsymbol{A},\boldsymbol{b})} d\boldsymbol{x} = \frac{1}{\sqrt{\det(\boldsymbol{A}^T\boldsymbol{A})}}\int_{\omega} d\boldsymbol{z}$$

□

G.2 Proof of Theorem 1

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

$$\begin{aligned}P_{\boldsymbol{G}}(\boldsymbol{x}\in w) =& \sum_{\omega\in\Omega}\int_{\omega\cap w} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x}))\sqrt{\det(J_{\boldsymbol{G}^{-1}}(\boldsymbol{x})^TJ_{\boldsymbol{G}^{-1}}(\boldsymbol{x}))}d\boldsymbol{x}\\ =& \sum_{\omega\in\Omega}\int_{\omega\cap w} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x}))\sqrt{\det((\boldsymbol{A}_\omega^+)^T\boldsymbol{A}_\omega^+)}d\boldsymbol{x}\\ =& \sum_{\omega\in\Omega}\int_{\omega\cap w} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x}))(\prod_{i:\sigma_i(\boldsymbol{A}_\omega^+)>0}\sigma_i(\boldsymbol{A}_\omega^+))d\boldsymbol{x}\\ =& \sum_{\omega\in\Omega}\int_{\omega\cap w} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x}))(\prod_{i:\sigma_i(\boldsymbol{A}_\omega)>0}\sigma_i(\boldsymbol{A}_\omega))^{-1}d\boldsymbol{x} \quad \text{Etape 1}\\ =& \sum_{\omega\in\Omega}\int_{\omega\cap w} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x}))\frac{1}{\sqrt{\det(\boldsymbol{A}_\omega^T\boldsymbol{A}_\omega)}}d\boldsymbol{x}\end{aligned}$$

Let now prove the Etape 1 step by proving that $\sigma_i(A^+) = (\sigma_i(A))^{-1}$ where we lighten notations as $A := \boldsymbol{A}_\omega$ and USV^T is the svd-decomposition of A:

$$\begin{aligned}A^+ = (A^TA)^{-1}A^T =& ((USV^T)^T(USV^T))^{-1}(USV^T)^T\\ =& (VS^TU^TUSV^T)^{-1}(USV^T)^T\\ =& (VS^TSV^T)^{-1}VS^TU^T\\ =& V(S^TS)^{-1}S^TU^T\\ \implies& \sigma_i(A^+) = (\sigma_i(A))^{-1}\end{aligned}$$

with the above it is direct to see that $\sqrt{\det((\boldsymbol{A}_\omega^+)^T\boldsymbol{A}_\omega^+)} = \frac{1}{\sqrt{\det(\boldsymbol{A}_\omega^T\boldsymbol{A}_\omega)}}$ as follows

$$\begin{aligned}\sqrt{\det((\boldsymbol{A}_\omega^+)^T\boldsymbol{A}_\omega^+)} = \prod_{i:\sigma_i\neq 0}\sigma_i(\boldsymbol{A}_\omega^+) =& \prod_{i:\sigma_i\neq 0}\sigma_i(\boldsymbol{A}_\omega)^{-1}\\ =& \left(\prod_{i:\sigma_i\neq 0}\sigma_i(\boldsymbol{A}_\omega)\right)^{-1}\\ =& \frac{1}{\sqrt{\det(\boldsymbol{A}_\omega^T\boldsymbol{A}_\omega)}}\end{aligned}$$

which gives the desired result. □

G.3 PROOF OF COROLLARY 1

Proof. We now demonstrate the use of Thm. 1 where we consider that the latent distribution is set as $\boldsymbol{z} \sim \mathcal{N}(0,1)$. We obtain that

$$\begin{aligned}
\boldsymbol{p}_{\boldsymbol{G}}(\boldsymbol{x} \in w) &= \sum_{\omega \in \Omega} \int_{\omega \cap w} \mathbb{1}_{\boldsymbol{x} \in \boldsymbol{G}(\omega)} p_{\boldsymbol{z}}(\boldsymbol{G}^{-1}(\boldsymbol{x})) \det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)^{-\frac{1}{2}} d\boldsymbol{x} \\
&= \sum_{\omega \in \Omega} \int_{\omega \cap w} \mathbb{1}_{\boldsymbol{x} \in \boldsymbol{G}(\omega)} \frac{1}{(2\pi)^{S/2}\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} e^{-\frac{1}{2}\|\boldsymbol{G}^{-1}(\boldsymbol{x})\|_2^2} d\boldsymbol{x} \\
&= \sum_{\omega \in \Omega} \int_{\omega \cap w} \mathbb{1}_{\boldsymbol{x} \in \boldsymbol{G}(\omega)} \frac{1}{(2\pi)^{S/2}\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} e^{-\frac{1}{2}((\boldsymbol{A}_\omega^+(\boldsymbol{x}-\boldsymbol{b}_\omega))^T((\boldsymbol{A}^+(\boldsymbol{x}-\boldsymbol{b}_\omega))} d\boldsymbol{x} \\
&= \sum_{\omega \in \Omega} \int_{\omega \cap w} \mathbb{1}_{\boldsymbol{x} \in \boldsymbol{G}(\omega)} \frac{1}{(2\pi)^{S/2}\sqrt{\det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)}} e^{-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{b}_\omega)^T(\boldsymbol{A}_\omega^+)^T\boldsymbol{A}_\omega^+(\boldsymbol{x}-\boldsymbol{b}_\omega)} d\boldsymbol{x}
\end{aligned}$$

giving the desired result that is reminiscent of Kernel Density Estimation (KDE) (Rosenblatt, 1956) and in particular adaptive KDE (Breiman et al., 1977), where a partitioning of the data manifold is performed on each cell (ω in our case) different kernel parameters are used. □

Proof. We now turn into the uniform latent distribution case on a bounded domain U in the DGN input space. By employing again Thm. 1, the given formula one can directly obtain that the output density is given by

$$p_{\boldsymbol{G}}(\boldsymbol{x}) = \frac{\sum_{\omega \in \Omega} \mathbb{1}_{\boldsymbol{x} \in \omega} \det(\boldsymbol{A}_\omega^T \boldsymbol{A}_\omega)^{-\frac{1}{2}}}{Vol(U)} \tag{11}$$

□

G.4 PROOF OF THM. 2

As we assume successful training, then regardless of the actual distribution p_x, the DGN will learn the correct underlying manifold, and learn the best approximation to p_x as possible onto this manifold. Now, applying MaGNET sampling i.e. Sec. 3.2 is equivalent to sampling from a distribution $p_{\boldsymbol{m}z}$ such that after DGN mapping, that distribution is uniform on the learned manifold (see Thm. 1). As we assumed that regardless of p_x the DGN approximates correctly the true manifold, and as we then adapt the sampling distribution $p_{\boldsymbol{m}z}$ to always obtain uniform sampling on that manifold, we see that this final sampling becomes invariant upon the data distribution (on the manifold) leading to the desired result.