



HAL
open science

Learning disconnected manifolds: no GAN's land

Ugo Tanielian, Thibaut Issenhuth, Elvis Dohmatob, Jérémie Mary

► **To cite this version:**

Ugo Tanielian, Thibaut Issenhuth, Elvis Dohmatob, Jérémie Mary. Learning disconnected manifolds: no GAN's land. International Conference on Machine Learning 2020, Jul 2020, Vienna, Austria. hal-02778441

HAL Id: hal-02778441

<https://inria.hal.science/hal-02778441v1>

Submitted on 5 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Learning disconnected manifolds: no GAN’s land

Ugo Tanielian^{1,2} Thibaut Issenhuth² Elvis Dohmatob² Jérémie Mary²

Abstract


Typical architectures of Generative Adversarial Networks make use of a unimodal latent/input distribution transformed by a continuous generator. Consequently, the modeled distribution always has connected support which is cumbersome when learning a disconnected set of manifolds. We formalize this problem by establishing a "no free lunch" theorem for the disconnected manifold learning stating an upper-bound on the precision of the targeted distribution. This is done by building on the necessary existence of a low-quality region where the generator continuously samples data between two disconnected modes. Finally, we derive a rejection sampling method based on the norm of generator’s Jacobian and show its efficiency on several generators including BigGAN.

1. Introduction

GANs (Goodfellow et al., 2014) provide a very effective tool for the unsupervised learning of complex probability distributions. For example, Karras et al. (2019) generate very realistic human faces while Yu et al. (2017) match state-of-the-art text corpora generation. Despite some early theoretical results on the stability of GANs (Arjovsky & Bottou, 2017) and on their approximation and asymptotic properties (Biau et al., 2018), their training remains challenging. More specifically, GANs raise a mystery formalized by Khayatkhoei et al. (2018): *how can they fit disconnected manifolds when they are trained to continuously transform a unimodal latent distribution?* While this question remains widely open, we will show that studying it can lead to some improvements in the sampling quality of GANs.

Indeed, training a GAN with the objective of continuously transforming samples from an unimodal distribution into a disconnected requires balancing between two caveats. On

¹Université Paris-Sorbonne, Paris, France ²Criteo AI Lab, France. Correspondence to: Ugo Tanielian <u.tanielian@criteo.com>.



(a) Heatmap of the generator’s Jacobian norm. White circles: quantiles of the latent distribution $\mathcal{N}(0, I)$.



(b) Green: target distribution. Coloured dots: generated samples colored w.r.t. the Jacobian Norm using same heatmap than (a).

Figure 1. Learning disconnected manifolds leads to the apparition of an area with high gradients and data sampled in between modes.

one hand, the generator could just ignore all modes but one, producing a very limited variety of high quality samples: this is an extreme case of the well known mode collapse (Arjovsky & Bottou, 2017). On the other hand, the generator could cover the different modes of the target distribution and necessarily generates samples out of the real data manifold as previously explained by Khayatkhoei et al. (2018).

As brought to the fore by Roth et al. (2017), there is a density mis-specification between the true distribution and the model distribution. Indeed, one cannot find parameters such that the model density function is arbitrarily close to the true distribution. To solve this issue, many empirical works have proposed to over-parameterize the generative distributions, as for instance, using a mixture of generators to better fit the different target modes. Tolstikhin et al. (2017) rely on boosting while Khayatkhoei et al. (2018) force each generator to target different sub-manifolds thanks to a criterion based on mutual information. Another direction is to add complexity in the latent space using a mixture of Gaussian distributions (Gurumurthy et al., 2017).

To better visualize this phenomenon, we consider a simple 2D motivational example where the real data lies on two disconnected manifolds. Empirically, when learning the distribution, GANs split the Gaussian latent space into two modes, as highlighted by the separation line in red in Figure 1a. More importantly, each sample drawn inside this red area in Figure 1a is then mapped in the output space in between the two modes (see Figure 1b). For the quantitative evaluation of the presence of out-of-manifold samples, a natural metric is the Precision-Recall (PR) proposed by Sajjadi et al. (2018) and its improved version (Improved PR) (Kynkäänniemi et al., 2019). A first contribution of this paper is to formally link them. Then, taking advantage of these metrics, we lower bound the measure of this out-of-manifold region and formalize the impossibility of learning disconnected manifolds with standard GANs. We also extend this observation to the multi-class generation case and show that the volume of off-manifold areas increases with the number of covered manifolds. In the limit, this increase drives the precision to zero.

To solve this issue and increase the precision of GANs, we argue that it is possible to remove out-of-manifold samples using a truncation method. Building on the work of Arvanitidis et al. (2017) who define a Riemaniann metric that significantly improves clustering in the latent space, our truncation method is based on information conveyed by the Jacobian’s norm of the generator. We empirically show that this rejection sampling scheme enables us to better fit disconnected manifolds without over-parametrizing neither the generative class of functions nor the latent distribution. Finally, in a very large high dimensional setting, we discuss the advantages of our rejection method and compare it to the truncation trick introduced by (Brock et al., 2019).

In a nutshell, our contributions are the following:

- We discuss evaluation of GANs and formally link the PR measure (Sajjadi et al., 2018) and its Improved PR version (Kynkäänniemi et al., 2019).
- We upper bound the precision of GANs with Gaussian latent distribution and formalize an impossibility result for disconnected manifolds learning.
- Using toy datasets, we illustrate the behavior of GANs when learning disconnected manifolds and derive a new truncation method based on the Jacobian’s Frobenius norm of the generator. We confirm its empirical performance on state-of-the-art models and datasets.

2. Related work

Fighting mode collapse. Goodfellow et al. (2014) were the first to raise the problem of mode collapse in the learning of disconnected manifolds with GANs. They observed that

when the generator is trained too long without updating the discriminator, the output distribution collapses to a few modes reducing the diversity of the samples. To tackle this issue, Salimans et al. (2016); Lin et al. (2018) suggested feeding several samples to the discriminator. Srivastava et al. (2017) proposed the use of a reconstructor network, mapping the data to the latent space to increase diversity.

In a different direction, Arjovsky & Bottou (2017) showed that training GANs using the original formulation (Goodfellow et al., 2014) leads to instability or vanishing gradients. To solve this issue, they proposed a Wasserstein GAN architecture (Arjovsky et al., 2017) where they restrict the class of discriminative functions to 1-Lipschitz functions using weight clipping. Pointing to issues with this clipping, Gulrajani et al. (2017); Miyato et al. (2018) proposed relaxed ways to enforce the Lipschitzness of the discriminator, either by using a gradient penalty or a spectral normalization. Albeit not exactly approximating the Wasserstein’s distance (Petzka et al., 2018), both implementations lead to good empirical results, significantly reducing mode collapse. Building on all of these works, we will further assume that generators are now able to cover most of the modes of the target distribution, leaving us the problem of out-of-manifold samples (*a.k.a.* low-quality pictures).

Generation of disconnected manifolds. When learning complex manifolds in high dimensional spaces using deep generative models, Fefferman et al. (2016) highlighted the importance of understanding the underlying geometry. More precisely, the learning of disconnected manifold requires the introduction of disconnectedness in the model. Gurumurthy et al. (2017) used a multi-modal entry distribution, making the latent space disconnected, and showed better coverage when data is limited and diverse. Alternatively, Khayatkhoei et al. (2018) studied the learning of a mixture of generators. Using a mutual information term, they encourage each generator to focus on a different submanifold so that the mixture covers the whole support. This idea of using an ensemble of generators is also present in the work of Tolstikhin et al. (2017) and Zhong et al. (2019), though they were primarily interested in the reduction of mode collapse.

In this paper, we propose a truncation method to separate the latent space into several disjoint areas. It is a way to learn disconnected manifolds without relying on the previously introduced over-parameterization techniques. As our proposal can be applied without retraining the whole architecture, we can use it successfully on very large nets. Close to this idea, Azadi et al. (2019) introduced a rejection strategy based on the output of the discriminator. However, this rejection sampling scheme requires the discriminator to be trained with a classification loss while our proposition can be applied to any generative models.

Evaluating GANs. The evaluation of generative models is an active area of research. Some of the proposed metrics only measure the quality of the generated samples such as the Inception score (Salimans et al., 2016) while others define distances between probability distributions. This is the case of the Frechet Inception distance (Heusel et al., 2017), the Wasserstein distance (Arjovsky et al., 2017) or kernel-based metrics (Gretton et al., 2012). The other main caveat for evaluating GANs lies in the fact that one does not have access to the true density nor the model density, prohibiting the use of any density based metrics. To solve this issue, the use of a third network that acts as an objective referee is common. For instance, the Inception score uses outputs from InceptionNet while the Fréchet Inception Distance compares statistics of InceptionNet activations. Since our work focuses on out-of-manifold samples, a natural measure is the PR measure (Sajjadi et al., 2018) and its Improved PR version (Kynkäänniemi et al., 2019), extensively discussed in the next section.

In the following, alongside precise definitions, we exhibit an upper bound on the precision of GANs with high recall (*i.e.* no mode collapse) and present a new truncation method.

3. Our approach

We start with a formal description of the framework of GANs and the relevant metrics. We later show a "no free lunch" theorem proving the necessary existence of an area in the latent space that generates out-of-manifold samples. We name this region the *no GAN’s land* since any data point sampled from this area will be in the frontier in between two different modes. We claim that dealing with it requires special care. Finally, we propose a rejection sampling procedure to avoid points out of the true manifold.

3.1. Notations

In the original setting of Generative Adversarial Networks (GANs), one tries to generate data that are "similar" to samples collected from some unknown probability measure μ_* . To do so, we use a parametric family of generative distribution where each distribution is the push-forward measure of a latent distribution Z and a continuous function modeled by a neural network.

Assumption 1. *The latent distribution Z is such that S_Z is a connected space.*

Note that for any distribution μ , S_μ refers to its support. Assumption 1 is common for GANs as in most of all practical applications, the random variable Z defined on a low dimensional space \mathbb{R}^d is either a multivariate Gaussian distribution or uniform distribution defined on a compact.

The measure μ_* is defined on a subset E of \mathbb{R}^D (potentially

a highly dimensional space), equipped with the norm $\|\cdot\|$. The generator has the form of a parameterized class of functions from \mathbb{R}^d (a space with a much lower dimension) to E , say $\mathcal{G} = \{G_\theta : \theta \in \Theta\}$, where $\Theta \subseteq \mathbb{R}^p$ is the set of parameters describing the model. Each function G_θ thus takes input from a d -dimensional random variable Z (Z is associated with probability distribution γ) and outputs "fake" observations with distribution μ_θ . Thus, the class of probability measures $\mathcal{P} = \{\mu_\theta : \theta \in \Theta\}$ is the natural class of distributions associated with the generator, and the objective of GANs is to find inside this class of candidates the one that generates the most realistic samples, closest to the ones collected from the unknown distribution μ_* .

Assumption 2. *Let $L > 0$. The generator G_θ takes the form of a neural network whose Lipschitz constant is smaller than L , *i.e.* for all (z, z') , we have $\|G_\theta(z') - G_\theta(z)\| \leq L\|z - z'\|$.*

This is a reasonable assumption, since Virmaux & Scaman (2018) present an algorithm that upper-bounds the Lipschitz constant of deep neural networks. Initially, 1-Lipschitzness was enforced only for the discriminator by clipping the weights (Arjovsky et al., 2017; Zhang et al., 2018), adding a gradient penalty (Gulrajani et al., 2017; Roth et al., 2017; Petzka et al., 2018), or penalizing the spectral norms (Miyato et al., 2018). Nowadays, state-of-the-art architectures for large scale generators such as SAGAN (Zhang et al., 2019) and BigGAN (Brock et al., 2019) also make use of spectral normalization for the generator.

3.2. Evaluating GANs with Precision and Recall

When learning disconnected manifolds, Srivastava et al. (2017) proved the need of measuring simultaneously the quality of the samples generated and the mode collapse. Sajjadi et al. (2018) proposed the use of a PR metric to measure the quality of GANs. The key intuition is that precision should quantify how much of the fake distribution can be generated by the true distribution while recall measures how much of the true distribution can be re-constructed by the model distribution. More formally, it is defined as follows:

Definition 1. (Sajjadi et al., 2018) *Let X, Y be two probability distributions. For $\alpha, \beta \in (0, 1]$, the probability distribution X is said to have an attainable precision α at recall β w.r.t. Y if there exists probability distributions μ, ν_X, ν_Y such that*

$$Y = \beta\mu + (1 - \beta)\nu_Y \quad \text{and} \quad X = \alpha\mu + (1 - \alpha)\nu_X$$

The component ν_Y denotes the part of Y that is "missed" by X , whereas, ν_X denotes the "noise" part of X . We denote $\bar{\alpha}$ (respectively $\bar{\beta}$) the maximum attainable precision (respectively recall). Th. 1 of (Sajjadi et al., 2018) states:

$$X(S_Y) = \bar{\alpha} \quad \text{and} \quad Y(S_X) = \bar{\beta}$$

Improved PR metric. Kynkäänniemi et al. (2019) highlighted an important drawback of the PR metric proposed by Sajjadi et al. (2018): it cannot correctly interpret situations when a large numbers of samples are packed together. To better understand this situation, consider a case where the generator slightly collapses on a specific data point, i.e. there exists $x \in E, \mu_\theta(x) > 0$. We show in Appendix A that if μ_\star is a non-atomic probability measure and μ_θ is highly precise (i.e. $\alpha = 1$), then the recall β must be 0.

To solve these issues, Kynkäänniemi et al. (2019) proposed an *Improved Precision-Recall* (Improved PR) metric built on a nonparametric estimation of support of densities.

Definition 2. (Kynkäänniemi et al., 2019) Let X, Y be two probability distributions. Let D_X, D_Y be two finite sample datasets: $D_X := \{x_i \sim X, i \in [1, n]\}$ and $D_Y := \{y_i \sim Y, i \in [1, n]\}$. For any $x \in D_X$ (respectively for any $y \in D_Y$), we consider $(x_{(1)}, \dots, x_{(n-1)})$, the re-ordering of elements in $D_X \setminus x$ given their euclidean distance with x . For any $k \in \mathbb{N}$ and $x \in D_X$, the precision $\alpha_k^n(x)$ of point x is defined as

$$\alpha_k^n(x) = 1 \iff \exists y \in D_Y, \|x - y\| \leq \|y_{(k)} - y\|$$

Similarly, the recall $\beta_k^n(y)$ of any given $y \in D_Y$ is

$$\beta_k^n(y) = 1 \iff \exists x \in D_X, \|y - x\| \leq \|x_{(k)} - x\|$$

Improved precision (respectively recall) are defined as the average over D_X (respectively D_Y) as follows

$$\alpha_k^n = \frac{1}{n} \sum_{x_i \in D_X} \alpha_k^n(x_i) \quad \beta_k^n = \frac{1}{n} \sum_{y_i \in D_Y} \beta_k^n(y_i)$$

A first contribution is to formalize the link between PR and Improved PR with the following theorem:

Theorem 1. Let μ_\star and μ_θ be two probability distributions with uniformly continuous probability density functions f_μ and f_ν . Assume that there exists constants $a_1 > 0, a_2 > 0$ such that for all $x \in E$ we have $a_1 < f_{\mu_\star}(x) \leq a_2$ and $a_1 < f_{\mu_\theta}(x) \leq a_2$ for some $c > 0$. Assume that k, n are such that $\frac{k}{\log(n)} \rightarrow +\infty$ and $\frac{k}{n} \rightarrow 0$. Then,

$$\alpha_k^n \rightarrow \bar{\alpha} \text{ in probability} \quad \text{and} \quad \beta_k^n \rightarrow \bar{\beta} \text{ in proba.}$$

This theorem, whose proof is delayed to Appendix B, underlines the nature of the Improved PR metric: the metric compares the supports of the modeled probability distribution μ_θ and of the true distribution μ_\star . This means that Improved PR is a tuple made of both maximum attainable precision $\bar{\alpha}$ and recall $\bar{\beta}$ (e.g. Theorem 1 of (Sajjadi et al., 2018)). As Improved PR is shown to have a better performance evaluating GANs sample quality, we use this metric for both the following theoretical results and experiments.

3.3. Learning disconnected manifolds

In this section, we aim to stress the difficulties of learning disconnected manifolds with standard GANs architectures. To begin with, we recall the following lemma.

Lemma 1. Assume that Assumptions 1 and 2 are satisfied. Then, for any $\theta \in \Theta$, the support S_{μ_θ} is connected.

There is consequently a discrepancy between the connectedness of S_{μ_θ} and the disconnectedness of S_{μ_\star} . In the case where the manifold lays on two disconnected components, our next theorem exhibit a no free lunch theorem:

Theorem 2. ("No free lunch" theorem) Assume that Assumptions 1 and 2 are satisfied. Assume also that true distribution μ_\star lays on two equally measured disconnected manifolds distant from a distance $D > 0$.

Then, any estimator μ_θ that samples equally in both modes must have a precision $\bar{\alpha}$ such that $\bar{\alpha} + \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\frac{\Phi^{-1}(\frac{\bar{\alpha}}{2})^2}{2}} \leq 1$. Besides, if $\bar{\alpha} \geq 3/4$, $\bar{\alpha} \lesssim 1 - \sqrt{\frac{2}{\pi}} W\left(\frac{D^2}{4L^2}\right)$ where W is the Lambert W function.

The proof of this theorem is delayed to Appendix C. It is mainly based on the Gaussian isoperimetric inequality (Borell, 1975; Sudakov & Tsirelson, 1978) that states that among all sets of given Gaussian measure in any finite dimensional Euclidean space, half-spaces have the minimal Gaussian boundary measure. If in Fig. 1, the generator has thus learned the optimal separation, it is yet not known, to the limit of our knowledge, how to enforce such geometrical properties in the latent space.

In real world applications, when the number of distinct sub-manifolds increases, we expect the volume of these boundaries to increase with respect to the number of different classes covered by the modeled distribution μ_θ . Going in this direction, we better formalize this situation, and show an extended "no free lunch theorem" by expliciting an upper-bound of the precision $\bar{\alpha}$ in this broader framework.

Assumption 3. The true distribution μ_\star lays on M equally-measured disconnected components at least distant from some constant $D > 0$.

This is likely to be true for datasets made of symbol designed to be highly distinguishable (e.g. digits in the MNIST dataset). In very high dimension, this assumption also holds for complex classes of objects appearing in many different contexts (e.g. the bubble class in ImageNet, see Appendix).

To better apprehend the next theorem, note A_m the pre-image in the latent space of mode m and A_m^r its r -enlargement: $A_m^r := \{z \in \mathbb{R}^d \mid \text{dist}(z, A_m) \leq r\}, r > 0$.

Theorem 3. (Generalized "no free lunch" theorem) Assume that Assumptions 1, 2, and 3 are satisfied, and that the pre-

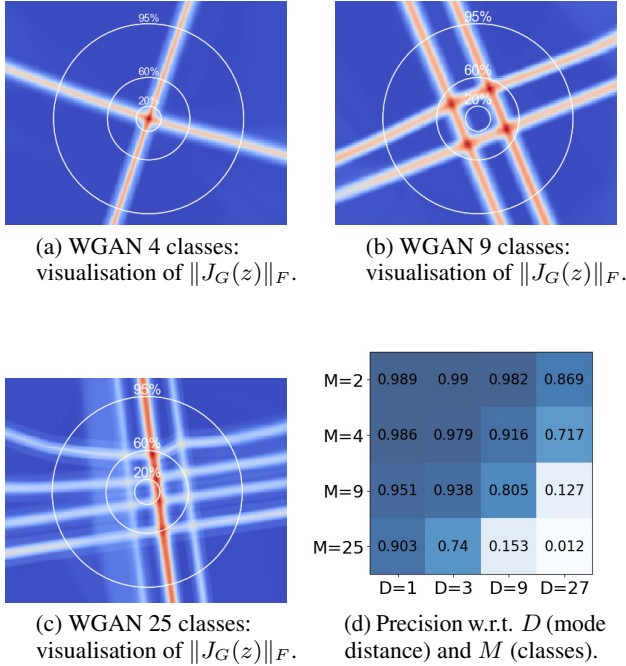


Figure 2. Illustration of Theorem 3. If the number of classes $M \rightarrow \infty$ or the distance $D \rightarrow \infty$, then the precision $\bar{\alpha} \rightarrow 0$. We provide in appendix heatmaps for more values of M .

image enlargements A_m^ε , with $\varepsilon = \frac{D}{2L}$, form a partition of the latent space with equally measured elements.

Then, any estimator μ_θ with recall $\bar{\beta} > \frac{1}{M}$ must have a precision $\bar{\alpha}$ at most $\frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}$ where $x = \Phi^{-1}(1 - \frac{1}{\bar{\beta}M})$ and Φ is the c.d.f. of a standard normal distribution.

Theorem 3, whose proof is delayed to Appendix D, states a lower-bound the measure of samples mapped out of the true manifold. We expect our bound to be loose since no theoretical results are known, to the best of our knowledge, on the geometry of the separation that minimizes the boundary between different classes (when $M \geq 3$). Finding this optimal cut would be an extension of the honeycomb theorem (Hales, 2001). In Appendix D.2 we give a more technical statement of Theorem 3 without assuming equality of measure of the sets A_m^ε .

The idea of the proof is to consider the border of an individual cell with the rest of the partition. It is clear that at least half of the frontier will be inside this specific cell. Then, to get to the final result, we sum the measures of the frontiers contained inside all of the different cells. Remark that our analysis is fine enough to keep a dependency in M which translates into a maximum precision that goes to zero when M goes to the infinity and all the modes are covered. More precisely, in this scenario where all pre-images have equal

measures in the latent space, one can derive the following bound for large values of K :

$$\bar{\alpha} \stackrel{K \rightarrow \infty}{\leq} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon\sqrt{2\log(K)}} \quad \text{where } \varepsilon = \frac{D}{2L} \quad (1)$$

For a fixed generator, Equation (1) illustrates that the precision $\bar{\alpha}$ decreases when either the distance D (equivalently ε) or the number of classes M increases. For a given ε , $\bar{\alpha}$ converges to 0 with a speed $O(\frac{1}{K\sqrt{2\varepsilon}})$. To better illustrate this asymptotic result, we provide results from a 2D synthetic setting. In this toy dataset, we control both the number M of disconnected manifolds and the distance D . Figure 2 clearly corroborates (1) as we can easily get the maximum precision close to 0 ($M = 25$, $D = 27$).

3.4. Jacobian-based truncation (JBT) method

The analysis of the deformation of the latent space offers a grasp on the behavior of GANs. For instance, Arvanitidis et al. (2017) propose a distance accounting for the distortions made by the generator. For any pair of points $(z_1, z_2) \sim Z^2$, the distance is defined as the length of the geodesic $d(z_1, z_2) = \int_{[0,1]} \|J_{G_\theta}(\gamma_t) \frac{d\gamma_t}{dt}\| dt$ where γ is the geodesic parameterized by $t \in [0, 1]$ and $J_{G_\theta}(z)$ denotes the Jacobian matrix of the generator at point z . Authors have shown that the use of this distance in the latent space improves clustering and interpretability. We make a similar observation that the generator’s Jacobian Frobenius norm provides meaningful information.

Indeed, the frontiers highlighted in Figures 2a, 2b, and 2c correspond to areas of low precision mapped out of the true manifold: this is the *no GAN’s land*. We argue that when learning disconnected manifolds, the generator tries to minimize the number of samples that do not belong to the support of the true distribution and that this can only be done by making paths steeper in the *no GAN’s land*. Consequently, data points $G_\theta(z)$ with high Jacobian Frobenius norm (JFN) are more likely to be outside the true manifold. To improve the precision of generative models, we thus define a new truncation method by removing points with highest JFN.

However, note that computing the generators’s JFN is expensive to compute for neural networks, since being defined as follows,

$$\|J_{G_\theta}(z)\|_F^2 = \sum_{i=1}^m \sum_{j=1}^n \left(\frac{\partial G_\theta(z)_i}{\partial z_j} \right)^2,$$

it requires a number of backward passes equal to the output dimension. To make our truncation method tractable, we use a stochastic approximation of the Jacobian Frobenius

norm based on the following result from Rifai et al. (2011):

$$\|J_{G_\theta}(z)\|^2 = \lim_{\substack{N \rightarrow \infty \\ \sigma \rightarrow 0}} \frac{1}{N} \sum_{\varepsilon_i} \frac{1}{\sigma^2} \|G_\theta(z + \varepsilon_i) - G_\theta(z)\|^2$$

where $\varepsilon_i \sim \mathcal{N}(0, \sigma^2 I)$ and I is the identity matrix of dimension d . The variance σ of the noise and the number of samples are used as hyper-parameters. In practice, σ in $[1e-4; 1e-2]$ and $N = 10$ give consistent results.

Based on the preceding analysis, we propose a new **Jacobian-based truncation (JBT)** method that rejects a certain ratio of the generated points with highest JFN. This truncation ratio is considered as an hyper-parameter for the model. We show in our experiments that our JBT can be used to detect samples outside the real data manifold and that it consequently improves the precision of the generated distribution as measured by the Improved PR metric.

4. Experiments

In the following, we show that our truncation method, JBT, can significantly improve the performances of generative models on several models, metrics and datasets. Furthermore, we compare JBT with over-parametrization techniques specifically designed for disconnected manifold learning. We show that our truncation method reaches or surpasses their performance, while it has the benefit of not modifying the training process of GANs nor using a mixture of generators, which is computationally expensive. Finally, we confirm the efficiency of our method by applying it on top of BigGAN (Brock et al., 2019).

Except for BigGAN, for all our experiments, we use Wasserstein GAN with gradient penalty (Gulrajani et al., 2017), called WGAN for conciseness. We give in Appendix K the full details of our experimental setting. The use of WGAN is motivated by the fact that it was shown to stabilize the training and significantly reduce mode collapse (Arjovsky & Bottou, 2017). However, we want to emphasise that our method can be plugged on top of any generative model fitting disconnected components.

4.1. Evaluation metrics

To measure performances of GANs when dealing with low dimensional applications - as with synthetic datasets - we equip our space with the standard Euclidean distance. However, for high dimensional applications such as image generation, Brock et al. (2019); Kynkäänniemi et al. (2019) have shown that embedding images into a feature space with a pre-trained convolutional classifier provides more semantic information. In this setting, we consequently use the euclidean distance between the images’ embeddings from a classifier. For a pair of images (a, b) , we define the distance $d(a, b)$ as $d(a, b) = \|\phi(a) - \phi(b)\|_2$ where




Figure 3. Mixture of 9 Gaussians in green, generated points in blue. Our truncation method (JBT) removes least precise data points as marginal precision plummets.

ϕ is a pre-softmax layer of a supervised classifier, trained specifically on each dataset. Doing so, they will more easily separate images sampled from the true distribution μ_* from the ones sampled by the distribution μ_θ .

We compare performances using Improved PR (Kynkäänniemi et al., 2019). We also report the *Marginal Precision* which is the precision of newly added samples when increasing the ratio of kept samples. Besides, for completeness, we report FID (Heusel et al., 2017) and recall precise definitions in Appendix G. Note that FID was not computed with InceptionNet, but a classifier pre-trained on each dataset.

4.2. Synthetic dataset

We first consider the true distribution to be a 2D Gaussian mixture of 9 components. Both the generator and the discriminator are modeled with a simple MLP with two hidden fully connected layers.

Interestingly, the generator tries to minimize the sampling of off-manifolds data during training until its JFN gets saturated. A visualization of this phenomenon is provided in Appendix H. One way to reduce the number of off-manifold samples is to use JBT. Indeed, off-manifold data points progressively disappear when being more and more selective, as illustrated in Figure 3c. We quantitatively confirm that our truncation method (JBT) improves the precision. On Fig. 3d, we observe that keeping the 70% of lowest JFN samples leads to an almost perfect precision of the support

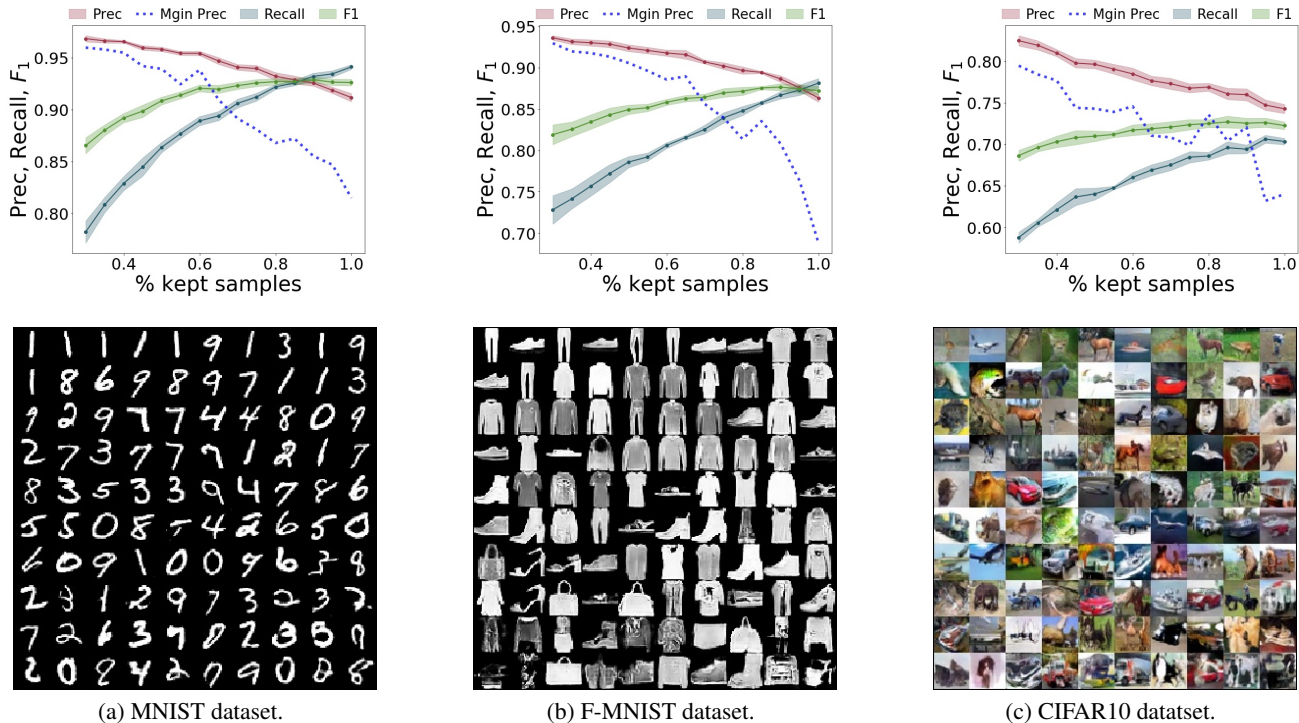


Figure 4. For high levels of kept samples, the marginal precision plummets of newly added samples, underlining the efficiency of our truncation method (JBT). Reported confidence intervals are 97% confidence intervals. On the second row, generated samples ordered by their JFN (left to right, top to bottom). In the last row, the data points generated are blurrier and outside the true manifold.

of the generated distribution. It proves that the off-manifold samples are in the 30% samples with highest JFN.

4.3. Image datasets

We further study JBT on three different datasets: MNIST (LeCun et al., 1998), FashionMNIST (Xiao et al., 2017) and CIFAR10 (Krizhevsky et al., 2009). Following (Khayatkhoei et al., 2018) implementation, we use a standard CNN architecture for MNIST and FashionMNIST while training a ResNet-based model for CIFAR10 (Gulrajani et al., 2017).

Figure 4 highlights that JBT also works on high dimensional datasets as the marginal precision plummets for high truncation ratios. Furthermore, when looking at samples ranked by increasing order of their JFN, we notice that samples with highest JFN are standing in-between manifolds. For example, those are ambiguous digits resembling both a "0" and a "6" or shoes with unrealistic shapes.

To further assess the efficiency of our truncation method, we also compare its performances with two state-of-the-art over-parameterization techniques that were designed for disconnected manifold learning. First, (Gurumurthy et al., 2017) propose DeliGAN, a reparametrization trick to transform the unimodal Gaussian latent distribution into a mixture. The different mixture components are later learnt by gradient descent. For fairness, the re-parametrization trick is

MNIST	Prec.	Rec.	FID
WGAN	91.2±0.3	93.7±0.5	24.3±0.3
WGAN JBT 90%	92.5±0.5	92.9±0.3	26.9±0.5
WGAN JBT 80%	93.3±0.3	91.8±0.4	33.1±0.3
W-Deligan	89.0±0.6	93.6±0.3	31.7±0.5
DMLGAN	93.4±0.2	92.3±0.2	16.8±0.4
F-MNIST			
WGAN	86.3±0.4	88.2±0.2	259.7±3.5
WGAN JBT 90%	88.6±0.6	86.6±0.5	257.4±3.0
WGAN JBT 80%	89.8±0.4	84.9±0.5	396.2±6.4
W-Deligan	88.5±0.3	85.3±0.6	310.9±3.1
DMLGAN	87.4±0.3	88.1±0.4	253.0±2.8

Table 1. JBT $x\%$ means we keep the $x\%$ samples with lowest Jacobian norm. Our truncation method (JBT) matches over-parameterization techniques. \pm is 97% confidence interval.

used on top of WGAN. Second, (Khayatkhoei et al., 2018) define DMLGAN, a mixture of generators to better learn disconnected manifolds. In this architecture, each generator is encouraged to target a different submanifold by enforcing high mutual information between generated samples and generator’s ids. Keep in mind that for DeliGAN (respectively DMLGAN), the optimal number of components (respectively generators) is not known and is a hyper-parameter of the model that has to be cross-validated.




Figure 5. On the first row, per-class precision-recall curves comparing Brock et al. (2019)’s truncation trick and our truncation method (JBT), on three ImageNet classes generated by BigGAN. We show better results on complex and disconnected classes (e.g. bubble). Reported confidence intervals are 97% confidence intervals. On the second row, generated samples ordered by their JFN (left to right, top to bottom). We observe a concentration of off-manifold samples for images on the bottom row, confirming the soundness of JBT.

The results of the comparison are presented in Table 1. In both datasets, JBT 80 % outperforms DeliGAN and DMLGAN in terms of precision while keeping a reasonable recall. This confirms our claim that over-parameterization techniques are unnecessary. As noticed by Kynkäänniemi et al. (2019), we also observe that FID does not correlate properly with the Improved PR metric. Based on the Frechet distance, only a distance between multivariate Gaussians, we argue that FID is not suited for disconnected manifold learning as it approximates distributions with unimodal ones and loses many information.

4.4. Spurious samples rejections on BigGAN

Thanks to the simplicity of JBT, we can also apply it on top of any trained generative model. In this subsection, we use JBT to improve the precision of a pre-trained BigGAN model (Brock et al., 2019), which generates class-conditioned ImageNet (Deng et al., 2009) samples. The class-conditioning lowers the problem of off-manifold samples, since it reduces the disconnectedness in the output distribution. However, we argue that the issue can still exist on high-dimensional natural images, in particular complex classes can still be multi-modal (e.g. the bubble class). The bottom row in Figure 5 shows a random set of 128 images for three different classes ranked by their JFN in ascending order (left to right, top to bottom). We observe a clear concentration of spurious samples on the bottom row images.

To better assess the Jacobian based truncation method, we compare it with the truncation trick from Brock et al. (2019). This truncation trick aims to reduce the variance of the latent space distribution using truncated Gaussians. While easy and effective, this truncation has some issues: it requires to complexify the loss to enforce orthogonality in weight matrices of the network. Moreover, as explained by Brock et al. (2019) "only 16% of models are amenable to truncation, compared to 60% when trained with Orthogonal Regularization". For fairness of comparison, the pre-trained network we use is optimized for their truncation method. On the opposite, JBT is simpler to apply since 100% of the tested models were amenable to the proposed truncation.

Results of this comparison are shown in the upper row of Figure 5. Our method can outperform their truncation trick on difficult classes with high intra-class variation, e.g. bubble and house finch. This confirms our claim that JBT can detect outliers within a class. However, one can note that their trick is particularly well suited for simpler unimodal classes, e.g. parachute and reaches high precision levels.

5. Conclusion

In this paper, we provide insights on the learning of disconnected manifolds with GANs. Our analysis shows the necessary existence of an off-manifold area with low precision. Besides, we empirically show on several datasets and

models that we can detect these areas and remove samples located in between two modes thanks to a newly proposed truncation method. In future work, we will study how the use and understanding of latent space geometry can further help us to better fit the interior of disconnected manifolds. We suspect the *no GAN's land* to be involved in the training instability of GANs since the presence of high gradients in a low probability region might induce a lot of variance.

References

- Arjovsky, M. and Bottou, L. Towards principled methods for training generative adversarial networks. In *International Conference on Learning Representations*, 2017.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223, 2017.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. In *ICLR*, 2017.
- Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. Discriminator rejection sampling. In *International Conference on Learning Representations*, 2019.
- Biau, G. and Devroye, L. *Lectures on the nearest neighbor method*. Springer, 2015.
- Biau, G., Cadre, B., Sangnier, M., and Tanielian, U. Some theoretical properties of gans. *arXiv:1803.07819*, 2018.
- Borell, C. The brunn-minkowski inequality in gauss space. *Inventiones mathematicae*, 30(2):207–216, 1975.
- Boucheron, S., Lugosi, G., and Massart, P. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. OUP Oxford, 2013. ISBN 9780199535255.
- Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Devroye, L. and Wise, G. Detection of abnormal behavior via nonparametric estimation of the support. *SIAM Journal on Applied Mathematics*, 38:480–488, 1980.
- Dowson, D. and Landau, B. The fréchet distance between multivariate normal distributions. *Journal of multivariate analysis*, pp. 450–455, 1982.
- Dudley, R. M. *Real Analysis and Probability*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2 edition, 2002. doi: 10.1017/CBO9780511755347.
- Fefferman, C., Mitter, S., and Narayanan, H. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 2016.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, J. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. 2014.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *Journal of Machine Learning Research*, pp. 723–773, 2012.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems*, pp. 5767–5777, 2017.
- Gurumurthy, S., Kiran Sarvadevabhatla, R., and Venkatesh Babu, R. Deligan: Generative adversarial networks for diverse and limited data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- Hales, T. C. The honeycomb conjecture. *Discrete & Computational Geometry*, pp. 1–22, 2001.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6626–6637, 2017.
- Kallenberg, O. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.
- Khayatkhoei, M., Singh, M. K., and Elgammal, A. Disconnected manifold learning for generative adversarial networks. In *Advances in Neural Information Processing Systems*, pp. 7343–7353, 2018.
- Krizhevsky, A., Hinton, G., et al. Learning multiple layers of features from tiny images. 2009.
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., and Aila, T. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, pp. 3929–3938, 2019.

- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pp. 2278–2324, 1998.
- Ledoux, M. Isoperimetry and gaussian analysis. In *Lectures on probability theory and statistics*, pp. 165–294. Springer, 1996.
- Lin, Z., Khetan, A., Fanti, G., and Oh, S. Pacgan: The power of two samples in generative adversarial networks. In *Advances in Neural Information Processing Systems*, pp. 1498–1507, 2018.
- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Pandeva, T. and Schubert, M. Mmgan: Generative adversarial networks for multi-modal distributions. *arXiv:1911.06663*, 2019.
- Petzka, H., Fischer, A., and Lukovnikov, D. On the regularization of wasserstein GANs. In *International Conference on Learning Representations*, 2018.
- Rifai, S., Mesnil, G., Vincent, P., Muller, X., Bengio, Y., Dauphin, Y., and Glorot, X. Higher order contractive auto-encoder. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 645–660. Springer, 2011.
- Roth, K., Lucchi, A., Nowozin, S., and Hofmann, T. Stabilizing training of generative adversarial networks through regularization. In *Advances in Neural Information Processing Systems*, pp. 2018–2028, 2017.
- Sajjadi, M. S., Bachem, O., Lucic, M., Bousquet, O., and Gelly, S. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems*, pp. 5228–5237, 2018.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pp. 2234–2242, 2016.
- Srivastava, A., Valkov, L., Russell, C., Gutmann, M. U., and Sutton, C. Veegan: Reducing mode collapse in gans using implicit variational learning. In *Advances in Neural Information Processing Systems*, pp. 3308–3318, 2017.
- Sudakov, V. N. and Tsirelson, B. S. Extremal properties of half-spaces for spherically invariant measures. *Journal of Mathematical Sciences*, pp. 9–18, 1978.
- Tolstikhin, I. O., Gelly, S., Bousquet, O., Simon-Gabriel, C.-J., and Schölkopf, B. Adagan: Boosting generative models. In *Advances in Neural Information Processing Systems*, pp. 5424–5433, 2017.
- Virmaux, A. and Scaman, K. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In *Advances in Neural Information Processing Systems*, pp. 3835–3844, 2018.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. 2017.
- Yu, L., Zhang, W., Wang, J., and Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Zhang, H., Goodfellow, I., Metaxas, D., and Odena, A. Self-attention generative adversarial networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7354–7363, 2019.
- Zhang, P., Liu, Q., Zhou, D., Xu, T., and He, X. On the discriminative-generalization tradeoff in GANs. In *International Conference on Learning Representations*, 2018.
- Zhong, P., Mo, Y., Xiao, C., Chen, P., and Zheng, C. Rethinking generative mode coverage: A pointwise guaranteed approach. In *Advances in Neural Information Processing Systems*, pp. 2086–2097, 2019.

A. Highlighting drawbacks of the PR metric by Sajjadi et al. (2018)

Lemma 2. Assume that the modeled distribution μ_θ slightly collapses on a specific data point, i.e. there exists $x \in E$, $\mu_\theta(x) > 0$. Assume also that μ_\star is a continuous probability measure and that μ_θ has a recall $\beta = 1$. Then the precision must be such that $\alpha = 0$.

Proof. Using Definition 1, we have that there exists μ such that

$$\mu_\star = \alpha\mu + (1 - \alpha)\nu_{\mu_\star} \quad \text{and} \quad \mu_\theta = \mu$$

Thus, $0 = \mu_\star(x) \geq \alpha\mu(x) = \alpha\mu_\theta(x)$. Which implies that $\alpha = 0$. \square

B. Proof of Theorem 1

The proof of Theorem 1 relies on theoretical results from non-parametric estimation of the supports of probability distribution studied by Devroye & Wise (1980).

For the following proofs, we will require the following notation: let φ be a strictly monotonous function be such that $\lim_{n \rightarrow \infty} \frac{\varphi(n)}{n} = 0$ and $\lim_{n \rightarrow \infty} \frac{\varphi(n)}{\log(n)} = \infty$. We note $B(x, r) \subseteq E$, the open ball centered in x and of radius r . For a given probability distribution μ , S_μ refers to its support. We recall that for any x in a dataset D , $x_{(k)}$ denotes its k nearest neighbor in D . Finally, for a given probability distribution μ and a dataset D_μ sampled from μ^n , we note R_{\min} and R_{\max} the following:

$$R_{\min} = \min_{x \in E} \|x - x_{(\varphi(n))}\|, \quad R_{\max} = \max_{x \in E} \|x - x_{(\varphi(n))}\| \quad (2)$$

In the following lemma, we show asymptotic behaviours for both R_{\min} and R_{\max} .

Lemma 3. Let μ be a probability distribution associated with a uniformly continuous probability density function f_μ . Assume that there exists constants $a_1 > 0, a_2 > 0$ such that for all $x \in E$, we have $a_1 < f_\mu(x) \leq a_2$. Then,

$$R_{\min} \xrightarrow[n \rightarrow \infty]{} 0 \text{ a.s.} \quad \text{and} \quad R_{\min}^d \xrightarrow[n \rightarrow \infty]{} \infty \text{ a.s.}$$

$$R_{\max} \xrightarrow[n \rightarrow \infty]{} 0 \text{ a.s.} \quad \text{and} \quad R_{\max}^d \xrightarrow[n \rightarrow \infty]{} \infty \text{ a.s.}$$

Proof. We will only prove that $R_{\max} \xrightarrow[n \rightarrow \infty]{} 0$ a.s. and and $R_{\min}^d \xrightarrow[n \rightarrow \infty]{} \infty$ a.s. as the rest follows.

The result is based on a nearest neighbor result from Biau & Devroye (2015). Considering the $\varphi(n)$ nearest neighbor density estimate $f_n^{\varphi(n)}$ based on a finite sample dataset D_μ , Theorem 4.2 states that if f_μ is uniformly continuous then:

$$\sup_{x \in E} \|f_n^{\varphi(n)}(x) - f_\mu(x)\| \rightarrow 0.$$

where $f_n^{\varphi(n)}(x) = \frac{\varphi(n)}{nV_d\|x - x_{\varphi(n)}\|^d}$ with V_d being the volume of the unit ball in \mathbb{R}^d .

Let $\varepsilon > 0$ such that $\varepsilon < a_1/2$. There exists $N \in \mathbb{N}$ such that for all $n \geq N$, we have, almost surely, for all $x \in E$:

$$a_1 - \varepsilon \leq f_n^{\varphi(n)}(x) \leq a_2 + \varepsilon$$

$$a_1 - \varepsilon \leq \frac{\varphi(n)}{nV_d\|x - x_{\varphi(n)}\|^d} \leq a_2 + \varepsilon$$

Consequently, for all $n \geq N$, for all $x \in E$ almost surely:

$$\|x - x_{\varphi(n)}\| \leq \left(\frac{\varphi(n)}{nV_d(a_1 - \varepsilon)} \right)^{1/d}$$

Thus, $\sup_{x \in E} \|x - x_{\varphi(n)}\| \rightarrow 0$ a.s.

Also, almost surely

$$n\|x - x_{\varphi(n)}\|^d \geq \frac{\varphi(n)}{V_d(a_2 + \varepsilon)}$$

Thus, $\inf_{x \in E} \|x - x_{\varphi(n)}\| \rightarrow \infty$ a.s. \square

Lemma 4. Let μ, ν be two probability distributions associated with uniformly continuous probability density functions f_μ and f_ν . Assume that there exists constants $a_1 > 0, a_2 > 0$ such that for all $x \in E$, we have $a_1 < f_\mu(x) \leq a_2$ and $a_1 < f_\nu \leq a_2$. Also, let D_μ, D_ν be datasets sampled from μ^n, ν^n . If μ is an estimator for ν , then

(i) for all $x \in D_\mu$, $\alpha_{\varphi(n)}^n(x) \xrightarrow[n \rightarrow \infty]{} \mathbb{1}_{\text{supp}(\nu)}(x)$ in proba.

(ii) for all $y \in D_\nu$, $\beta_{\varphi(n)}^n(y) \xrightarrow[n \rightarrow \infty]{} \mathbb{1}_{\text{supp}(\mu)}(x)$ in proba.

Proof. We will only show the result for (i), since a similar proof holds for (ii).

Thus, we want to show that

$$\text{for all } x \in D_\mu, \alpha_{\varphi(n)}^n(x) \xrightarrow[n \rightarrow \infty]{} \mathbb{1}_{\text{supp}(\nu)}(x) \text{ a. s.}$$

First, let's assume that $x \notin S_\nu$. Biau & Devroye (2015, Lemma 2.2) have shown that

$$\lim_{n \rightarrow \infty} \|x_{(\varphi(n))} - x\| = \inf\{\|x - y\| \mid y \in S_\nu\} \text{ a.s.}$$

As S_ν is a closed set - e.g. (Kallenberg, 2006) - we have

$$\lim_{n \rightarrow \infty} \|x - x_{(\varphi(n))}\| > 0 \text{ a.s.}$$

and

$$\text{for all } y \in D_\nu, \lim_{n \rightarrow \infty} \|y - y_{(\varphi(n))}\| = 0 \text{ a.s.}$$

Thus, $\lim_{n \rightarrow \infty} \alpha_{\varphi(n)}^n(x) = 0$ a.s..

Now, let's assume that $x \in S_\nu$. Using Definition 2, the precision of a given data point x can be rewritten as follows:

$$\alpha_{\varphi(n)}^n(x) = 1 \iff \exists y \in D_\nu, x \in B(y, \|y - y_{\varphi(n)}\|)$$

Using notation from (2), we note

$$R_{\min} = \min_{y \in E} \|y - y_{\varphi(n)}\|, \quad R_{\max} = \max_{y \in E} \|y - y_{\varphi(n)}\|.$$

It is clear that :

$$\bigcup_{y \in D_\nu} B(y, R_{\min}) \subseteq S_\nu^n \subseteq \bigcup_{y \in D_\nu} B(y, R_{\max}), \quad (3)$$

where $S_\nu^n = \bigcup_{y \in D_\nu} B(y, \|y - y_{\varphi(n)}\|)$.

Besides, combining Lemma 3 with Devroye & Wise (1980, Theorem 1), we have that:

$$\begin{aligned} \nu(S_\nu \Delta \bigcup_{y \in D_\nu} B(y, R_{\min})) &\xrightarrow[n \rightarrow 0]{} 0 \quad \text{in proba.} \\ \nu(S_\nu \Delta \bigcup_{y \in D_\nu} B(y, R_{\max})) &\xrightarrow[n \rightarrow 0]{} 0 \quad \text{in proba.} \end{aligned}$$

where Δ here refers to the symmetric difference.

Thus, using (3), it is now clear that, $\mu(S_\nu \Delta S_\nu^n) \rightarrow 0$ in probability. Finally, given $x \in S_\mu$, we have $\mu(x \in S_\nu^n) = \nu(\alpha_{\varphi(n)}^n(x) = 1) \rightarrow 1$ in probability. \square

We can now finish the proof for Theorem 1. Recall that $\bar{\alpha} = \mu(S_\nu)$ and similarly, $\bar{\beta} = \nu(S_\mu)$.

Proof. We have that

$$|\alpha_{\varphi(n)}^n - \bar{\alpha}| = \left| \frac{1}{n} \sum_{x_i \in D_\mu} \alpha_{\varphi(n)}^n(x_i) - \int_E \mathbb{1}_{x \in S_\nu} \mu(dx) \right|$$

Then,

$$\begin{aligned} |\alpha_{\varphi(n)}^n - \bar{\alpha}| &= \left| \frac{1}{n} \sum_{x_i \in D_\mu} (\alpha_{\varphi(n)}^n(x_i) - \mathbb{1}_{x_i \in S_\nu}) \right. \\ &\quad \left. + \left(\frac{1}{n} \sum_{x_i \in D_\mu} \mathbb{1}_{x_i \in S_\nu} - \int_E \mathbb{1}_{x \in S_\nu} \mu(dx) \right) \right| \\ &= |\mathbb{E}_{x_i \sim \mu_n} (\alpha_{\varphi(n)}^n(x_i) - \mathbb{1}_{x_i \in S_\nu})| \quad (4) \\ &\quad + (|\mathbb{E}_{\mu_n} \mathbb{1}_{S_\nu} - \mathbb{E}_\mu \mathbb{1}_{S_\nu}|) \quad (5) \end{aligned}$$

where μ_n is the empirical distribution of μ . As μ_n converges weakly to μ almost surely (e.g. Dudley (2002, Theorem 11.4.1)) and since $\mathbb{1}_{x \in S_\nu}$ is bounded, we can bound (5) as follows:

$$\lim_{n \rightarrow \infty} \mathbb{E}_{x \sim \mu_n} \mathbb{1}_{x \in \text{supp}(\mu)} - \mathbb{E}_{x \sim \mu} \mathbb{1}_{x \in \text{supp}(\mu)} = 0 \quad \text{a. s.}$$

Now, to bound (4), we use the fact that for any $x \in D_\mu$, the random variable $\alpha_{\varphi(n)}^n(x)$ converges to $\mathbb{1}_{x \in S_\nu}$ in probability (Lemma 4) and that for all $x \in D_\mu$, both $\alpha_{\varphi(n)}^n(x) \leq 1$ and $\mathbb{1}_{x \in S_\nu} \leq 1$. Consequently, using results from the weak law for triangular arrays, we have that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{x_i \in D_\mu} (\alpha_{\varphi(n)}^n(x_i) - \mathbb{1}_{x_i \in S_\nu}) = 0 \quad \text{in proba.}$$

Finally,

$$|\alpha_{\varphi(n)}^n - \bar{\alpha}| \xrightarrow[n \rightarrow \infty]{} 0 \quad \text{in proba.,}$$

which proves the result. The same proof works for $\lim_{k \rightarrow \infty} \beta_k^n = \bar{\beta}$. \square

C. Proof of Theorem 2

This proof is based on the Gaussian isoperimetric inequality historically shown by (Borell, 1975; Sudakov & Tsirelson, 1978).

Proof. Let μ_\star be a distribution defined on E laying on two disconnected manifolds M_1 and M_2 such that $\mu_\star(M_1) = \mu_\star(M_2) = \frac{1}{2}$ and $d(M_1, M_2) = D$. Note that for any subsets $A \subseteq E$ and $B \subseteq E$, $d(A, B) := \inf_{(x,y) \in A \times B} \|x - y\|$.

Let $G_\theta^{-1}(M_1)$ (respectively $G_\theta^{-1}(M_2)$) be the subset in \mathbb{R}^d be the pre-images of M_1 (respectively M_2).

Consequently, we have for all $k \in [1, n]$

$$\gamma(G_\theta^{-1}(M_1)) = \mu_\theta(M_1) = \gamma(G_\theta^{-1}(M_2)) \geq \frac{\bar{\alpha}}{2}$$

We consider $(G_\theta^{-1}(M_1))^\varepsilon$ (respectively $(G_\theta^{-1}(M_2))^\varepsilon$) the ε enlargement of $G_\theta^{-1}(M_1)$ (respectively $G_\theta^{-1}(M_2)$) where $\varepsilon = \frac{D}{2L}$. We know that $(G_\theta^{-1}(M_1))^\varepsilon \cap (G_\theta^{-1}(M_2))^\varepsilon = \emptyset$.

Thus, we have that:

$$\gamma((G_\theta^{-1}(M_1))^\varepsilon) + \gamma((G_\theta^{-1}(M_2))^\varepsilon) \leq 1$$

Besides, by denoting Φ the function defined for any $t \in \mathbb{R}$ by $\Phi(t) = \int_{-\infty}^t \frac{\exp(-s^2/2)}{\sqrt{2\pi}} ds$, we have

$$\begin{aligned} \gamma((G_\theta^{-1}(M_1))^\varepsilon) + \gamma((G_\theta^{-1}(M_2))^\varepsilon) &\geq 2\Phi(\Phi^{-1}(\frac{\alpha}{2}) + \varepsilon) \\ &\quad \text{(using Theorem 1.3 from (Ledoux, 1996))} \\ &\geq \alpha + \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\Phi^{-1}(\frac{\alpha}{2})^2/2} \\ &\quad \text{(since } \Phi^{-1}(\frac{\alpha}{2}) + \varepsilon < 0 \text{ and } \Phi \text{ convex on }]-\infty, 0]) \end{aligned}$$

Thus, we have that

$$\alpha + \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\Phi^{-1}(\frac{\alpha}{2})^2/2} \leq 1$$

Thus, by noting

$$\alpha^* = \sup\{\alpha \in [0, 1] \mid \alpha + \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\frac{\Phi^{-1}(\frac{\alpha}{2})^2}{2}} \leq 1\},$$

we have our result.

For $\alpha \geq 3/4$. By noting $\alpha = 1 - x$, we have

$$\begin{aligned} \Phi^{-1}\left(\frac{\alpha}{2}\right) &= \frac{\sqrt{2\pi}x}{2} + O(x^3) \\ \text{And, } e^{-\frac{\Phi^{-1}(\frac{\alpha}{2})^2}{2}} &= e^{-\frac{\pi x^2}{4}} + O(e^{-x^4}) \\ \text{Thus, } 1 - x + \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\frac{\pi x^2}{4}} + O(e^{-x^4}) &\leq 1 \\ \iff x &\geq \frac{2\varepsilon}{\sqrt{2\pi}} e^{-\frac{\pi x^2}{4}} + O(e^{-x^4}) \\ \implies x &\geq \sqrt{\frac{2}{\pi}} W(\varepsilon^2) \end{aligned}$$

where W is the product log function. Thus, $\alpha \leq 1 - \sqrt{\frac{2}{\pi}} W(\varepsilon^2)$. \square

As an example, in the case where $\varepsilon = 1$, we have that $W(1) \approx 0.5671$, $x > 0.4525$ and $\alpha < 0.5475$.

D. Proof of Theorem 3

D.1. Equitable setting

This result is a consequence of Prop. 1 that we will assume true in this section.

We consider that the unknown true distribution μ_* lays on M disjoint manifolds of equal measure. As specified in Section 3, the latent distribution γ is a multivariate Gaussian defined on \mathbb{R}^d . For each $k \in [1, M]$, we consider in the latent space, the pre-images A_k .

It is clear that A_1, \dots, A_M are pairwise disjoint Borel subsets of \mathbb{R}^d . We denote \bar{M} , the number of classes covered by the estimator μ_θ , such that for all $i \in [1, \bar{M}]$, we have $\gamma(A_i) > 0$. We know that $\bar{M} \geq M\bar{\beta} > 1$.

For each $i \in [1, \bar{M}]$, we denote A_i^ε , the ε -enlargement of A_i . For any pair (i, j) it is clear that $A_i^\varepsilon \cap A_j^\varepsilon = \emptyset$ where $\varepsilon = \frac{D}{2L}$ (D being the minimum distance between two sub-manifolds and L being the Lipschitz constant of the generator).

As assumed, we know that $A_i^\varepsilon, i \in [1, \bar{M}]$ partition the latent space in equal measure, consequently, we assume that

$$\sum_{i=1}^{\bar{M}} \gamma(A_i^\varepsilon) = 1 \quad \text{and} \quad \gamma(A_1) = \dots = \gamma(A_M) = 1/\bar{M} \quad (6)$$

Thus, we have that

$$\bar{\alpha} = \sum_{i=1}^{\bar{M}} \gamma(A_i^\varepsilon) = 1 - \gamma(\Delta^{-\varepsilon}(A_1^\varepsilon, \dots, A_{\bar{M}}^\varepsilon))$$

Using Proposition 1, we have

$$\begin{aligned} \gamma(\Delta^{-\varepsilon}(A_1^\varepsilon, \dots, A_n^\varepsilon)) &\geq 1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x} \\ \text{Thus, } \bar{\alpha} &\leq \frac{1+y^2}{y^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon y} \end{aligned}$$

where $y = \Phi^{-1}(1 - \max_{k \in [\bar{M}]} \gamma(A_k^\varepsilon)) = \Phi^{-1}(\frac{\bar{M}-1}{\bar{M}})$ and $\Phi(t) = \int_{-\infty}^t \frac{\exp(-s^2/2)}{\sqrt{2\pi}} ds$.

Knowing that $\bar{M} \geq \bar{\beta}M$ we have that

$$\Phi^{-1}\left(1 - \frac{1}{\bar{M}}\right) \geq \Phi^{-1}\left(1 - \frac{1}{\bar{\beta}M}\right)$$

We conclude by saying that the function $x \mapsto \frac{1+x^2}{x^2} e^{-\varepsilon x}$ is decreasing for $x > 0$. Thus,

$$\bar{\alpha} \leq \frac{1+y^2}{y^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon y} \quad (7)$$

where $y = \Phi^{-1}\left(1 - \frac{1}{\bar{\beta}M}\right)$ and $\Phi(t) = \int_{-\infty}^t \frac{\exp(-s^2/2)}{\sqrt{2\pi}} ds$.

For the further analysis when $K \rightarrow \infty$, please refer to subsection E, note using the result in (16), one obtains the desired upper-bound on $\bar{\alpha}$

$$\bar{\alpha} \stackrel{K \rightarrow \infty}{\leq} 1 - e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon \sqrt{2 \log(K)}}$$

D.2. More general setting

As done previously, we denote \bar{M} , the number of classes covered by the estimator μ_θ , such that for all $i \in [1, \bar{M}]$, we have $\gamma(A_i) > 0$. We still assume that $\bar{M} > 1$. However, we now relax the previous assumption made in (6) and assume the milder assumption that there exists $w_1, \dots, w_M \in [0, 1]^M$ such that for all $m \in [1, M]$, $\gamma(A_m^\varepsilon) = w_m$, $\sum_m w_m \leq 1$ and $\max_{i \in [1, M]} w_m = w^{\max} < 1$.

Consider, $A^{\mathbb{G}} = \left(\bigcup_{i=1}^{\bar{M}} A_i^\varepsilon\right)^{\mathbb{G}}$ and denote $w^{\mathbb{G}} = \gamma(A^{\mathbb{G}}) \leq 1 - \bar{\alpha}$. Consequently, we have

$$\begin{aligned} \sum_{i=1}^{\bar{M}} \gamma(A_i^\varepsilon) + \gamma(A^{\mathbb{G}}) &= 1 \\ \gamma(\Delta^{-\varepsilon}(A_1^\varepsilon, \dots, A_M^\varepsilon, A^{\mathbb{G}})) + \sum_{i=1}^{\bar{M}} \gamma(A_i^\varepsilon) &= 1 - \gamma(A^{\mathbb{G}}) \\ \bar{\alpha} &= 1 - w^{\mathbb{G}} - \gamma(\Delta^{-\varepsilon}(A_1^\varepsilon, \dots, A_M^\varepsilon, A^{\mathbb{G}})) \end{aligned}$$

In this setting, it is clear that $A_1, \dots, A_M, A^{\mathbb{G}}$ is a partition of \mathbb{R}^d under the measure γ . Using, result from Proposition 1, we have

$$\gamma(\Delta^{-\varepsilon}(A_1^\varepsilon, \dots, A_M^\varepsilon, A^{\mathbb{G}})) \geq 1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}$$

where $x = \Phi^{-1}\left(1 - \max(w^{\mathbb{G}}, w^{\max})\right)$ and $\Phi(t) = \int_{-\infty}^t \frac{\exp(-t^2/2)}{\sqrt{2\pi}} ds$.

Finally, we have that

$$\bar{\alpha} \leq \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x} - w^{\mathbb{G}} \quad (8)$$

In the case where $\gamma(A^{\mathbb{G}}) = 0$, we find a result similar to (7).

E. Lower-bounding boundaries of partitions in a Gaussian space

Notations and preliminaries Given $\varepsilon \geq 0$ and a subset A of euclidean space $\mathbb{R}^d = (\mathbb{R}^d, \|\cdot\|)$, let $A^\varepsilon := \{z \in \mathbb{R}^d \mid \text{dist}(z, A) \leq \varepsilon\}$ be its ε -enlargement, where $\text{dist}(z, A) := \inf_{z' \in A} \|z' - z\|_2$ is the distance of the point $z \in \mathbb{R}^d$ from A . Let γ be the standard Gaussian distribution in \mathbb{R}^d and let A_1, \dots, A_K be $K \geq 2$ pairwise disjoint Borel subsets of \mathbb{R}^d whose union has unit (i.e full) Gaussian measure $\sum_{k=1}^K w_k = 1$, where $w_k := \gamma(A_k)$. Such a collection $\{A_1, \dots, A_K\}$ will be called an (w_1, \dots, w_K) -partition of standard d -dimensional Gaussian space (\mathbb{R}^d, γ) .

For each $k \in [K]$, define the compliment $A_{-k} := \cup_{k' \neq k} A_{k'}$, and let $\partial^{-\varepsilon} A_k := \{z \in A_k \mid \text{dist}(z, A_{-k}) \leq \varepsilon\}$ be the *inner ε -boundary* of A_k , i.e the points of A_k which are within distance ε of some other $A_{k'}$. For every $(k, k') \in [K]^2$ with $k' \neq k$, it is an easy exercise to show that

$$\begin{aligned} \partial^{-\varepsilon} A_k \cap \partial^{-\varepsilon} A_{k'} &= \emptyset \\ \partial^{-\varepsilon} A_k \cap A_{-k} &= \emptyset \\ A_{-k}^\varepsilon &= \partial^{-\varepsilon} A_k \cup A_{-k} \end{aligned} \quad (9)$$

Now, let $\Delta^{-\varepsilon}(A_1, \dots, A_K) := \cup_{k=1}^K \partial^{-\varepsilon} A_k$ be the union of all the inner ε -boundaries. This is $\Delta^{-\varepsilon}(A_1, \dots, A_K)$ the set of points of $\cup_{k=1}^K A_k$ which are on the boundary between some two distinct A_k and $A_{k'}$. We want to find a lower bound in the measure $\gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K))$.

Proposition 1. *Given $K \geq 4$ and $w_1, \dots, w_K \in (0, 1/4]$ such that $\sum_{k=1}^K w_k = 1$, we have the bound:*

$$\inf_{A_1, \dots, A_K} \gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) \geq 1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}$$

where the infimum is taken over all (w_1, \dots, w_K) -partitions of standard Gaussian space (\mathbb{R}^d, γ) , and $x := \Phi^{-1}\left(1 - \max_{k \in [K]} w_k\right)$.

Proof. By (9), we have the formula

$$\gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) = \sum_{k=1}^K \gamma(\partial^{-\varepsilon} A_k) \quad (10)$$

$$= \sum_{k=1}^K \gamma(A_{-k}^\varepsilon) - \gamma(A_{-k}). \quad (11)$$

Let $w_{-k} := \gamma(A_{-k}) = 1 - w_k$, and assume $w_{-k} \geq 3/4$, i.e $w_k \leq 1/4$, for all $k \in [K]$.

For example, this condition holds in the equitable scenario where $w_k = 1/K$ for all k .

Now, by standard *Gaussian Isoperimetric Inequality* (see (Boucheron et al., 2013) for example), one has

$$\begin{aligned} \gamma(A_{-k}^\varepsilon) &\geq \Phi(\Phi^{-1}(\gamma(A_{-k})) + \varepsilon) \\ &= \Phi(\Phi^{-1}(1 - w_k) + \varepsilon). \end{aligned} \quad (12)$$

Using the bound $\frac{x}{1+x^2} \varphi(x) < 1 - \Phi(x) < \frac{1}{x} \varphi(x) \forall x > 0$ where φ is the density of the standard Gaussian law. We can further find that

$$\begin{aligned} \Phi(\Phi^{-1}(1 - w_k) + \varepsilon) &\geq 1 - w_k \frac{1 + \Phi^{-1}(1 - w_k)^2}{\Phi^{-1}(1 - w_k)^2} \times \\ &\quad e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon \Phi^{-1}(1 - w_k)} \\ &\geq 1 - w_k \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x} > 0 \end{aligned} \quad (13)$$

(since the function $x \mapsto \frac{1+x^2}{x^2} e^{-\varepsilon x}$ is decreasing for $x > 0$)

where $x := \min_{k \in [K]} \Phi^{-1}(1 - w_k) = \Phi^{-1}(1 - \max_{k \in [K]} w_k) \geq \Phi^{-1}(3/4) > 0.67$. Combining (10), (12), and (13) yields the following

$$\begin{aligned} \gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) &\geq \sum_{k=1}^K \left(1 - w_k \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x} - (1 - w_k)\right) \\ &= \sum_{k=1}^K \left(1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}\right) w_k \\ &= 1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}, \end{aligned}$$

Asymptotic analysis In the limit, it is easy to check that in the case where $\max_{k \in [K]} w_k \rightarrow 0$, we have that $x \rightarrow \infty$.

In this setting, we thus have $\frac{1+x^2}{x^2} \rightarrow 1$ and can now derive the following bound:

$$\inf_{A_1, \dots, A_K} \gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) \xrightarrow{\max_{k \in [K]} w_k \rightarrow 0} 1 - e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}.$$

Equitable scenario In the equitable scenario where $w_k = 1/K$ for all k , we have

$$\inf_{A_1, \dots, A_K} \gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) \geq 1 - \frac{1+x^2}{x^2} e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon x}$$

where

$$x = \Phi^{-1}(1 - 1/K) \quad (14)$$

When $K \geq 8$ we additionally have

$$\Phi^{-1}(1 - 1/K) \geq \sqrt{2 \log \left(\frac{K(q(K)^2 - 1)}{\sqrt{2\pi}q(K)^3} \right)} \quad (15)$$

where $q(K) = \sqrt{2 \log(\sqrt{2\pi}K)}$.

Consequently, we have when $K \rightarrow \infty$, the following behavior:

$$\gamma(\Delta^{-\varepsilon}(A_1, \dots, A_K)) \xrightarrow{K \rightarrow \infty} 1 - e^{-\frac{1}{2}\varepsilon^2} e^{-\varepsilon \sqrt{2 \log(K)}} \quad (16)$$

□

Proof of the inequality (15). Set $p := 1/K$. First, for any $x > 0$, we have the following upper:

$$\int_x^\infty e^{-y^2/2} dy = \int_x^\infty \frac{y}{y} e^{-y^2/2} dy \leq \frac{1}{x} \int_x^\infty y e^{-y^2/2} dy = \frac{e^{-x^2/2}}{x}.$$

For a lower bound:

$$\int_x^\infty e^{-y^2/2} dy = \int_x^\infty \frac{y}{y} e^{-y^2/2} dy = \frac{e^{-x^2/2}}{x} - \int_x^\infty \frac{1}{y^2} e^{-y^2/2} dy$$

and

$$\int_x^\infty \frac{1}{y^2} e^{-y^2/2} dy = \int_x^\infty \frac{y}{y^3} e^{-y^2/2} dy \leq \frac{e^{-x^2/2}}{x^3}$$

and combining these gives

$$\int_x^\infty e^{-y^2/2} dy \geq \left(\frac{1}{x} - \frac{1}{x^3} \right) e^{-x^2/2}.$$

Thus

$$\frac{1}{\sqrt{2\pi}} \left(\frac{1}{x} - \frac{1}{x^3} \right) e^{-x^2/2} \leq 1 - \Phi(x) \leq \frac{1}{\sqrt{2\pi}} \frac{1}{x} e^{-x^2/2},$$

from where

$$\frac{1}{\sqrt{2\pi}} \left(\frac{1}{\Phi^{-1}(1-p)} - \frac{1}{\Phi^{-1}(1-p)^3} \right) e^{-\Phi^{-1}(1-p)^2/2} \quad (17)$$

$$\leq p \leq \frac{1}{\sqrt{2\pi}} \frac{1}{\Phi^{-1}(1-p)} e^{-\Phi^{-1}(1-p)^2/2} \quad (18)$$

Using (18), when $\Phi^{-1}(1-p) \geq 1$ (that is $p \leq 0.15$ or equivalently $K \geq 8$), we have the following upper bound $\Phi^{-1}(1-p) \leq q(p)$ where $q(p) := \sqrt{2 \log(\sqrt{2\pi}/p)}$. Then, injecting $q(p)$ in (17):

$$\frac{1}{\sqrt{2\pi}} \left(\frac{1}{q(p)} - \frac{1}{q(p)^3} \right) e^{-\Phi^{-1}(1-p)^2/2} \leq p.$$

Now when $q(p) \geq 1$ you have:


$$e^{-\Phi^{-1}(1-p)^2/2} \leq \frac{\sqrt{2\pi}p q(p)^3}{q(p)^2 - 1}$$

and


$$\Phi^{-1}(1-p) \geq \sqrt{2 \log \left(\frac{q(p)^2 - 1}{\sqrt{2\pi}p q(p)^3} \right)}. \quad (19)$$

There is one additional requirement on p which is simply that the argument of the log should be ≥ 1 i.e. $q(p)^2 - 1 \geq \sqrt{2\pi}p q(p)^3$, which is true as soon as $K \geq 8$. □


F. Visualization of Theorem 3




(a) WGAN 4 classes: visualisation of $\|J_G(z)\|_F$.




(c) WGAN 9 classes: visualisation of $\|J_G(z)\|_F$.




(e) WGAN 3 classes: visualisation of $\|J_G(z)\|_F$.




(g) WGAN 5 classes: visualisation of $\|J_G(z)\|_F$.




(b) Green blobs: true densities. Dots: generated points.



(d) Green blobs: true densities. Dots: generated points.



(f) Green blobs: true densities. Dots: generated points.



(h) Green blobs: true densities. Dots: generated points.

Figure 6. Learning disconnected manifolds: visualization of the gradient of the generator (JFN) in the latent space and densities in the output space.

G. Definition of the different metrics used

In the sequel, we present the different metrics used in Section 4 of the paper to assess performances of GANs. We have:

- Improved Precision/Recall (PR) metric (Kynkäänniemi et al., 2019): it has been presented in Definition 2. Intuitively, Based on a k-NN estimation of the manifold of real (resp. generated) data, it assesses whether generated (resp. real) points belong in the real (resp. generated) data manifold or not. The proportion of generated (resp. real) points that are in the real (resp. generated) data manifold is the precision (resp. recall).
- the Hausdorff distance: it is defined by

$$\text{Haus}(A, B) = \max \left\{ \max_{a \in A} \min_{b \in B} \|a - b\|, \max_{b \in B} \min_{a \in A} \|a - b\| \right\}$$

Such a distance is useful to evaluate the closeness of two different supports from a metric space, but is sensitive to outliers because of the max operation. It has been recently used for theoretical purposes by Pandevara & Schubert (2019).

- the Frechet Inception distance: first proposed by Dowson & Landau (1982), the Frechet distance was applied in the setting of GANs by Heusel et al. (2017). This distance between multivariate Gaussians compares statistic of generated samples to real samples as follows

$$\text{FID} = \|\nu_\star - \nu_\theta\|^2 + \text{Tr}(\Sigma_\star + \Sigma_\theta + 2(\Sigma_\star \Sigma_\theta)^{\frac{1}{2}})$$

where $X_\star = \mathcal{N}(\nu_\star, \Sigma_\star)$ and $X_\theta = \mathcal{N}(\nu_\theta, \Sigma_\theta)$ are the activations of a pre-softmax layer. However, when dealing with disconnected manifolds, we argue that this distance is not well suited as it approximates the distributions with unimodal one, thus losing many information.

The choice of such metrics is motivated by the fact that metrics measuring the performances of GANs should not rely on relative densities but should rather be point sets based metrics.

H. Saturation of a MLP neural network

In Section 4.2, we claim that the generator reduces the sampling of off-manifold data points up to a saturation point. Figure 7 below provides a visualization of this phenomenon. In this synthetic case, we learn a 9-component mixture of Gaussians using simple GANs architecture (both the generator and the discriminator are MLP with two hidden layers). The minimal distance between two modes is set to 9. We clearly see in Figure 7d that the precision saturates around 80%.

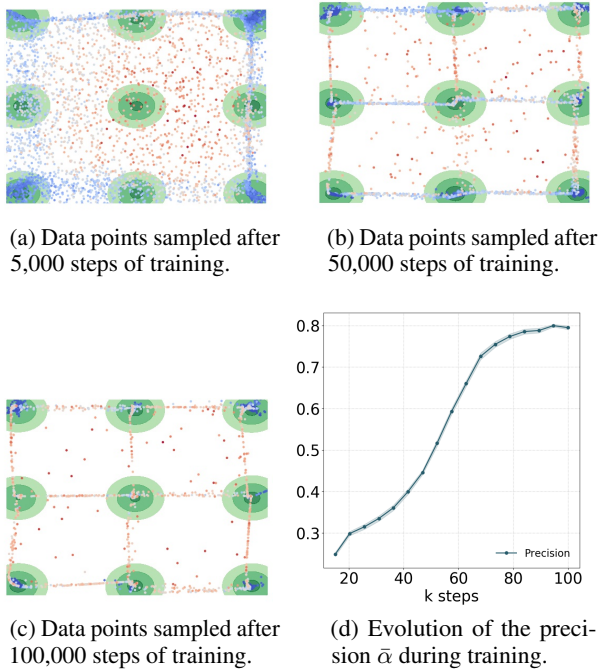


Figure 7. Learning 9 disconnected manifolds with a standard GANs architecture.

I. More results and visualizations on MNIST/F-MNIST/CIFAR10

Additionally to those in Section 4.3, we provide in Figure 9 and Table 2 supplementary results for MNIST, F-MNIST and CIFAR-10 datasets.



Figure 8. Visualization of our truncation method on CIFAR10.



Figure 9. Visualization of our truncation method (JBT) on three real-world datasets: MNIST, F-MNIST and CIFAR-10.




Figure 10. For high levels of kept samples, the marginal precision plummets of newly added samples, underlining the efficiency of our truncation method (JBT). Reported confidence intervals are 97% confidence intervals. On the second row, generated samples ordered by their JFN (left to right, top to bottom). In the last row, the data points generated are blurrier and outside the true manifold.

MNIST	Prec.	Rec.	F1	Haus.	FID	EMD
WGAN	91.2 ± 0.3	93.7 ± 0.5	92.4 ± 0.4	49.7 ± 0.2	24.3 ± 0.3	21.5 ± 0.1
WGAN 90% lowest JFN	92.5 ± 0.5	92.9 ± 0.3	92.7 ± 0.4	48.1 ± 0.2	26.9 ± 0.5	21.3 ± 0.2
WGAN 80% lowest JFN	93.3 ± 0.3	91.8 ± 0.4	92.6 ± 0.4	50.6 ± 0.4	33.1 ± 0.3	21.4 ± 0.4
W-Deligan	89.0 ± 0.6	93.6 ± 0.3	91.2 ± 0.5	50.7 ± 0.3	31.7 ± 0.5	22.4 ± 0.1
DMLGAN	93.4 ± 0.2	92.3 ± 0.2	92.8 ± 0.2	48.2 ± 0.3	16.8 ± 0.4	20.7 ± 0.1
Fashion-MNIST						
WGAN	86.3 ± 0.4	88.2 ± 0.2	87.2 ± 0.3	140.6 ± 0.7	259.7 ± 3.5	61.9 ± 0.3
WGAN 90% lowest JFN	88.6 ± 0.6	86.6 ± 0.5	87.6 ± 0.5	138.7 ± 0.9	257.4 ± 3.0	61.3 ± 0.6
WGAN 80% lowest JFN	89.8 ± 0.4	84.9 ± 0.5	87.3 ± 0.4	146.3 ± 1.1	396.2 ± 6.4	63.3 ± 0.7
W-Deligan	88.5 ± 0.3	85.3 ± 0.6	86.9 ± 0.4	141.7 ± 1.1	310.9 ± 3.1	60.9 ± 0.4
DMLGAN	87.4 ± 0.3	88.1 ± 0.4	87.7 ± 0.4	141.9 ± 1.2	253.0 ± 2.8	60.9 ± 0.4
CIFAR10						
WGAN	74.3 ± 0.5	70.3 ± 0.4	72.3 ± 0.5	334.7 ± 3.5	634.8 ± 4.6	151.2 ± 0.2
WGAN 90% lowest JFN	76.0 ± 0.7	69.4 ± 0.5	72.5 ± 0.6	318.1 ± 3.7	631.3 ± 4.5	150.7 ± 0.2
WGAN 80% lowest JFN	76.9 ± 0.5	68.6 ± 0.5	72.5 ± 0.5	323.5 ± 4.0	725.0 ± 3.5	150.1 ± 0.3
W-Deligan	71.5 ± 0.7	69.8 ± 0.7	70.6 ± 0.7	328.7 ± 2.1	727.8 ± 3.9	154.0 ± 0.3
DMLGAN	74.1 ± 0.5	65.7 ± 0.6	69.7 ± 0.6	328.6 ± 2.7	967.2 ± 4.1	152.0 ± 0.4

Table 2. Scores on MNIST and Fashion-MNIST. JFN stands for Jacobian Frobenius norm. ± is 97% confidence interval.

J. More results on BigGAN and ImageNet

In Figure 11, we show images from the Bubble class of ImageNet. It supports our claim of manifold disconnectedness, even within a class, and outlines the importance of studying the learning of disconnected manifolds in generative models. Then, in Figure 12, we give more examples from BigGAN 128x128 class-conditioned generator. We plot in the same format than in 4.4. Specifically, for different classes, we plot 128 images ranked by JFN. Here again, we see a concentration of off-manifold samples on the last row, proving the efficiency of our method. Example of classes responding particularly well to our ranking are House Finch (c), Monnarch Butterfly (i) or Wood rabbit (m). For each class, we also show an histogram of JFN based on 1024 samples. It shows that the JFN is a good indicator of the complexity of the class. For example, classes such as Cornet (q) or Football helmet (s) are very diverse and disconnected, resulting in high JFNs.




Figure 11. Images from the Bubble class of ImageNet showing that the class is complex and slightly multimodal.


Learning disconnected manifolds: no GAN's land




(a) 'Black swan' class.



(b) 'Black swan' class histogram.




(c) 'House finch' class.




(d) 'House finch' class histogram.




(e) 'Indigo bunting' class.



(f) 'Indigo bunting' class histogram.




(g) 'Cheetah' class.




(h) 'Cheetah' class histogram.




(i) 'Monarch butterfly' class.




(j) 'Monarch butterfly' class histogram.




(k) 'Loggerhead turtle' class.




(l) 'Loggerhead turtle' class histogram.




(m) 'Wood rabbit' class.




(n) 'Wood rabbit' class histogram.




(o) 'Trash can' class.



(p) 'Trash can' class histogram.




(q) 'Cornet/Horn' class.




(r) 'Cornet/Horn' class histogram.




(s) 'Football helmet' class.




(t) 'Football helmet' class histogram.




(u) 'Harmonica' class.



(v) 'Harmonica' class histogram.




(w) 'Parachute' class.



(x) 'Parachute' class histogram.



(y) 'Peacock' class.



(z) 'Peacock' class histogram.

Figure 12. For several classes with BigGAN model.

K. Network Architecture and Hyperparameters

Table 3. Models for Synthetic datasets

Operation	Feature Maps	Activation
G(z): $z \sim \mathcal{N}(0, 1)$	2	
Fully Connected - layer1	20	ReLU
Fully Connected - layer2	20	ReLU
D(x)		
Fully Connected - layer1	20	ReLU
Fully Connected - layer2	20	ReLU
Batch size	32	
Leaky ReLU slope	0.2	
Gradient Penalty weight	10	
Learning Rate	0.0002	
Optimizer	Adam: $\beta_1 = 0.5$	$\beta_2 = 0.5$

Table 4. WGAN for MNIST/Fashion MNIST

Operation	Kernel	Strides	Feature Maps	Activation
G(z): $z \sim \mathcal{N}(0, Id)$			100	
Fully Connected			$7 \times 7 \times 128$	
Convolution	3×3	1×1	$7 \times 7 \times 64$	LReLU
Convolution	3×3	1×1	$7 \times 7 \times 64$	LReLU
Nearest Up Sample			$14 \times 14 \times 64$	
Convolution	3×3	1×1	$14 \times 14 \times 32$	LReLU
Convolution	3×3	1×1	$14 \times 14 \times 32$	LReLU
Nearest Up Sample			$14 \times 14 \times 64$	
Convolution	3×3	1×1	$28 \times 28 \times 16$	LReLU
Convolution	5×5	1×1	$28 \times 28 \times 1$	Tanh
D(x)			$28 \times 28 \times 1$	
Convolution	4×4	2×2	$14 \times 14 \times 32$	LReLU
Convolution	3×3	1×1	$14 \times 14 \times 32$	LReLU
Convolution	4×4	2×2	$7 \times 7 \times 64$	LReLU
Convolution	3×3	1×1	$7 \times 7 \times 64$	LReLU
Fully Connected			1	-
Batch size	256			
Leaky ReLU slope	0.2			
Gradient Penalty weight	10			
Learning Rate	0.0002			
Optimizer	Adam	$\beta_1 : 0.5$	$\beta_2 : 0.5$	

For DeliGan, we use the same architecture and simply add 50 Gaussians for the reparametrization trick. For DMLGAN, we re-use the architecture of the authors.

Table 5. DMLGAN for MNIST/Fashion MNIST

Operation	Kernel	Strides	Feature Maps	BN	Activation
G(z): $z \sim N(0, Id)$			100		
Fully Connected			$7 \times 7 \times 128$	-	
Convolution	3×3	1×1	$7 \times 7 \times 64$	-	Leaky ReLU
Convolution	3×3	1×1	$7 \times 7 \times 64$	-	Leaky ReLU
Nearest Up Sample			$14 \times 14 \times 64$	-	
Convolution	3×3	1×1	$14 \times 14 \times 32$	-	Leaky ReLU
Convolution	3×3	1×1	$14 \times 14 \times 32$	-	Leaky ReLU
Nearest Up Sample			$14 \times 14 \times 64$	-	
Convolution	3×3	1×1	$28 \times 28 \times 16$	-	Leaky ReLU
Convolution	5×5	1×1	$28 \times 28 \times 1$	-	Tanh
Encoder Q(x), Discriminator D(x)			$28 \times 28 \times 1$		
Convolution	4×4	2×2	$14 \times 14 \times 32$	-	Leaky ReLU
Convolution	3×3	1×1	$14 \times 14 \times 32$	-	Leaky ReLU
Convolution	4×4	2×2	$7 \times 7 \times 64$	-	Leaky ReLU
Convolution	3×3	1×1	$7 \times 7 \times 64$	-	Leaky ReLU
D Fully Connected			1	-	-
Q Convolution	3×3		$7 \times 7 \times 64$	Y	Leaky ReLU
Q Convolution	3×3		$7 \times 7 \times 64$	Y	Leaky ReLU
Q Fully Connected			$n_g = 10$	-	Softmax
Batch size	256				
Leaky ReLU slope	0.2				
Gradient Penalty weight	10				
Learning Rate	0.0002				
Optimizer	Adam	$\beta_1 = 0.5$	$\beta_2 = 0.5$		

Table 6. WGAN for CIFAR10, from (Gulrajani et al., 2017)

Operation	Kernel	Strides	Feature Maps	BN	Activation
G(z): $z \sim N(0, Id)$			128		
Fully Connected			$4 \times 4 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$4 \times 4 \times 128$	Y	ReLU
Nearest Up Sample			$8 \times 8 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	Y	ReLU
Nearest Up Sample			$16 \times 16 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$16 \times 16 \times 128$	Y	ReLU
Nearest Up Sample			$32 \times 32 \times 128$	-	
Convolution	3×3	1×1	$32 \times 32 \times 3$	-	Tanh
Discriminator D(x)			$32 \times 32 \times 3$		
ResBlock	$[3 \times 3] \times 2$	1×1	$32 \times 32 \times 128$	-	ReLU
AvgPool	2×2	1×1	$16 \times 16 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$16 \times 16 \times 128$	-	ReLU
AvgPool	2×2	1×1	$8 \times 8 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	-	ReLU
ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	-	ReLU
Mean pooling (spatial-wise)	-	-	128	-	
Fully Connected			1	-	-
Batch size	64				
Gradient Penalty weight	10				
Learning Rate	0.0002				
Optimizer	Adam	$\beta_1 = 0.$	$\beta_2 = 0.9$		
Discriminator steps	5				

Table 7. DMLGAN for CIFAR10, from (Gulrajani et al., 2017)

Operation	Kernel	Strides	Feature Maps	BN	Activation
G(z): $z \sim N(0, Id)$			128		
Fully Connected			$4 \times 4 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$4 \times 4 \times 128$	Y	ReLU
Nearest Up Sample			$8 \times 8 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	Y	ReLU
Nearest Up Sample			$16 \times 16 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$16 \times 16 \times 128$	Y	ReLU
Nearest Up Sample			$32 \times 32 \times 128$	-	
Convolution	3×3	1×1	$32 \times 32 \times 3$	-	Tanh
Encoder Q(x), Discriminator D(x)			$32 \times 32 \times 3$		
ResBlock	$[3 \times 3] \times 2$	1×1	$32 \times 32 \times 128$	-	ReLU
AvgPool	2×2	1×1	$16 \times 16 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$16 \times 16 \times 128$	-	ReLU
AvgPool	2×2	1×1	$8 \times 8 \times 128$	-	
ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	-	ReLU
D ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	-	ReLU
D Mean pooling (spatial-wise)	2×2	1×1	128	-	
D Fully Connected			1	-	-
Q ResBlock	$[3 \times 3] \times 2$	1×1	$8 \times 8 \times 128$	-	ReLU
Q Mean pooling (spatial-wise)	2×2	1×1	128	-	
Q Fully Connected			$n_g = 10$	-	Softmax
Batch size	64				
Gradient Penalty weight	10				
Learning Rate	0.0002				
Optimizer	Adam	$\beta_1 = 0.$	$\beta_2 = 0.9$		
Discriminator steps	5				