



HAL
open science

Latent reweighting, an almost free improvement for GANs

Thibaut Issenhuth, Ugo Tanielian, David Picard, Jeremie Mary

► **To cite this version:**

Thibaut Issenhuth, Ugo Tanielian, David Picard, Jeremie Mary. Latent reweighting, an almost free improvement for GANs. 2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), Jan 2022, Waikoloa, United States. pp.3574-3583, 10.1109/WACV51458.2022.00363 . hal-03791961

HAL Id: hal-03791961

<https://hal.science/hal-03791961>

Submitted on 29 Sep 2022

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.

Latent reweighting, an almost free improvement for GANs

Thibaut Issenhuth^{1,2}, Ugo Tanielian¹, David Picard², Jérémie Mary¹

¹ Criteo AI Lab, Paris, France.

² LIGM, Ecole des Ponts, Marne-la-vallée, France.

{thibaut.issenhuth, ugo.tanielian, jeremie.mary}@criteo.com, david.picard@enpc.fr

Abstract

Standard formulations of GANs, where a continuous function deforms a connected latent space, have been shown to be misspecified when fitting different classes of images. In particular, the generator will necessarily sample some low-quality images in between the classes. Rather than modifying the architecture, a line of works aims at improving the sampling quality from pre-trained generators at the expense of increased computational cost. Building on this, we introduce an additional network to predict latent importance weights and two associated sampling methods to avoid the poorest samples. This idea has several advantages: 1) it provides a way to inject disconnectedness into any GAN architecture, 2) since the rejection happens in the latent space, it avoids going through both the generator and the discriminator, saving computation time, 3) this importance weights formulation provides a principled way to reduce the Wasserstein’s distance to the target distribution. We demonstrate the effectiveness of our method on several datasets, both synthetic and high-dimensional.

1. Introduction

GANs [10] are an effective way to learn complex and high-dimensional distributions, leading to state-of-the-art models for image synthesis in both unconditional [17] and conditional settings [6]. However, it is well-known that a single generator with an unimodal latent variable cannot recover a distribution composed of disconnected sub-manifolds [19]. This leads to a common problem for practitioners: the existence of very low-quality samples when covering different modes. This is formalized by [31] which refers to this area as the no GAN’s land and provides impossibility theorems on the learning of disconnected manifolds with standard formulations of GANs. Fitting a disconnected target distribution requires an additional mechanism inserting disconnectedness in the modeled distribution. A first solution is to add some expressivity to the model: [19] propose to train a mixture of generators,

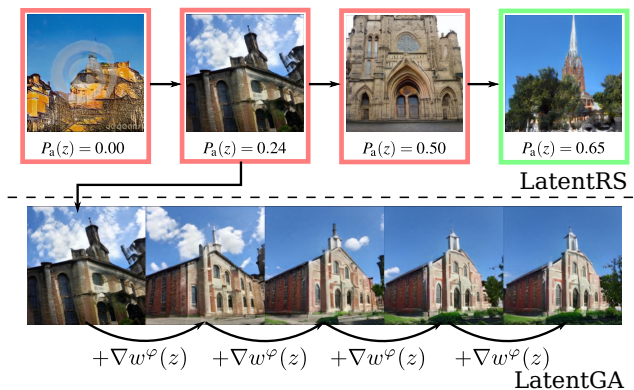


Figure 1: Overview of the proposed method. GANs tend to produce poor images for unlucky draws of the latent variable (top row, left). We introduce importance weights $w^\varphi(z)$ in the latent space that allow us to use rejection sampling and accept a given latent variable z with probability $P_a(z) \propto w^\varphi(z)$ (LatentRS, top row), or to perform a simple gradient ascent over the importance weight (LatentGA, bottom row), leading to better images. Both strategies can be combined for improved image quality. Images generated with StyleGAN2 trained on LSUN Church.

while [13] make use of a multi-modal latent distribution.

A second line of research relies heavily on a variety of Monte-Carlo algorithms, such as Rejection Sampling [3] or Metropolis-Hastings [32]. Monte-Carlo methods aim at sampling from a target distribution, while having only access to samples generated from a proposal distribution. Using the previously learned generative distribution as a proposal distribution, this idea was successfully applied to GANs. However, one of the main drawbacks is that Monte-Carlo algorithms only guarantee to sample from the target distribution under strong assumptions. First, we need access to the density ratios between the proposal and target distributions or equivalently to a perfect discriminator [3]. Second, these methods are efficient only if the support of the proposal distribution fully covers the one of the target distribution. This is unlikely to be the case when dealing with high-dimensional datasets [1].

To tackle this issue, we propose a novel method aiming at reducing the Wasserstein distance between the previously trained generative model and the target distribution. This is done via the adversarial training of a third network that learns importance weights in the latent space. Note that this network does not aim at increasing the support of the proposal distribution but at re-weighting the latent distribution, under a Wasserstein criterion. Thus, these importance weights define a new distribution in the latent space, from which we propose to sample with two complementary methods: latent rejection sampling (latentRS) and latent gradient ascent (latentGA). To better understand our approach, we illustrate its efficiency with simple examples. On the top of the Figure 1, we show samples coming from a pre-trained StyleGAN2 [17] and their respective acceptance probability (latentRS). At the bottom, we exhibit a sequence of generated images while following a gradient ascent on the learned importance weights (latentGA).

Our contributions are the following:

- We propose a novel approach that trains a neural network to directly modify the latent space of a GAN. This provides a principled way to reduce the Wasserstein distance to the target distribution.
- We show how to sample from this new generative model with different methods: latent Rejection Sampling (latentRS), latent Gradient Ascent (latentGA), and latentRS+GA, a method that leverages the complementarity between the two previous solutions.
- We run a large empirical comparison between our proposed methods and previous approaches on a variety of datasets and distributions. We empirically show that all of our proposed solutions significantly reduce the computational cost of inference. More interestingly, our solutions propose a wide span of performances ranging from latentRS, optimizing speed, that matches state-of-the-art almost for free (computational cost divided by 15) and latentRS+GA (computational cost divided by 3) that outperforms previous approaches.

Notation. Before moving to the related work section, we shortly present the notation needed in the paper. The goal of the generator is to generate data points that are “similar” to samples collected from some target probability measure μ_* . The measure μ_* is defined on a potentially high-dimensional space \mathbb{R}^D , equipped with the euclidean norm $\|\cdot\|$. We call μ_n the empirical measure. To approach μ_* , 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. In most applications, the random variable Z defined on a low-dimensional space \mathbb{R}^d is either a multivariate

Gaussian distribution or uniform distribution. The generator is a parameterized class of functions from \mathbb{R}^d to \mathbb{R}^D , 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 takes input from Z and outputs “fake” observations with distribution $\mu_\theta = G_\theta \# Z$. On the other hand, the discriminator is described by a family of functions from \mathbb{R}^D to \mathbb{R} , say $\mathcal{D} = \{D_\alpha : \alpha \in \Lambda\}$, $\Lambda \subseteq \mathbb{R}^Q$. Finally, for any given distribution μ , we note S_μ its support.

2. Related Work

[10] already stated that when training vanilla GANs, the generator could ignore modes of the target distribution: this is called mode collapse. A significant step towards understanding this phenomenon was made by [1] who explained that the standard formulation of GANs leads to vanishing or unstable gradients. The authors proposed the Wasserstein GANs (WGANs) architecture [2] where, in particular, discriminative functions are restricted to the class of 1-Lipschitz functions. WGANs aim at solving the following:

$$\sup_{\alpha \in \Lambda} \inf_{\theta \in \Theta} \mathbb{E}_{x \sim \mu_*} D_\alpha(x) - \mathbb{E}_{z \sim \gamma} D_\alpha(G_\theta(z)) \quad (1)$$

2.1. Learning disconnected manifolds with GANs: training and evaluation

The broader drawback of standard GANs is that, since any modeled distribution is the push-forward of a unimodal distribution by a continuous transformation, it has a connected support. This means that when the generator covers multiple disconnected modes of the target distribution, it necessarily generates samples out of the real data manifold [19]. Consequently, any thorough evaluation of GANs should assess simultaneously both the quality and the variety of the generated samples. To solve this issue, [27] and [20] propose a Precision/Recall metric that aims at measuring both the *mode dropping* and the *mode inventing*. The precision refers to the portion of generated points that belongs to the target manifold, while the recall measures how much of the target distribution can be reconstructed by the model distribution.

Building on this metric, [31] highlighted the trade-off property of GANs deriving upper-bounds on the precision of standard GANs. To solve this problem, a common direction of research consists in over-parameterizing the generative model. [19] enforces diversity by using a mixture of generators, while [13] suggests that a mixture of Gaussians in the latent space is efficient to learn diverse and limited data. Similarly, [4] propose importance weights that aim at robustifying the training of GANs and make it less sensitive to the target distribution’s outliers.

2.2. Improving the quality of GANs post-training

Another line of research consists in improving the sampling quality of pre-trained GANs. [31] proposed a heuristic to insert disconnectedness and remove the samples mapped out of the true manifold. [30] designed Discriminator Optimal Transport (DOT), a gradient ascent driven by a Wasserstein discriminator to improve every single sample. Similarly, [7] follow a discriminator-driven Langevin dynamic.

Another well-studied possibility would be to use Monte-Carlo (MC) methods [25]. Following this path, [3] were the first to use a rejection sampling method to improve the quality of the proposal distribution μ_θ . The authors use the fact that the optimal vanilla discriminator trained with binary cross-entropy is equal to $\mu_\star/(\mu_\star + \mu_\theta)$. Thus, a parametric discriminator $D_\alpha : \mathbb{R}^D \rightarrow [0, 1]$ can be used to approximate the density ratios r_α as follows:

$$r_\alpha(x) := \frac{\mu_\star(x)}{\mu_\theta(x)} = \frac{D_\alpha(x)}{1 - D_\alpha(x)}. \quad (2)$$

This density ratio can then be plugged in the Rejection Sampling (RS) algorithm. Doing so, it can be shown that sampling from μ_θ and accepting samples probabilistically is equivalent to sample from the target distribution μ_\star . The acceptance probability of a given sample x is $\mathbb{P}_a(x) = \frac{r_\alpha(x)}{k}$. This is valid as long as there is a constant $k \in \mathbb{R}^+$ such that $\mu_\star(x) \leq k\mu_\theta(x)$ for all x .

[32] use similar density ratios and derive MH-GAN, by using the independent Metropolis-Hasting algorithm [14]. Finally, [11] use these density ratios r_α as importance weights and perform discrete sampling relying on the Sampling-Importance-Resampling (SIR) algorithm [26]. Given $X_1, \dots, X_n \sim \mu_\theta^n$, we have:

$$\mu_{\theta, \alpha}^{\text{SIR}}(X_i) = \frac{r_\alpha(X_i)}{\sum_{j=1}^n r_\alpha(X_j)}. \quad (3)$$

Note that these algorithms all rely on similar density ratios and differ by the acceptance-rejection scheme chosen. Interestingly, in RS, the acceptance rate is not controlled, but we are guaranteed to sample from μ_\star . Conversely, with SIR and MH, the acceptance rate is a chosen parameter, but we are sampling from an approximation of the target distribution.

2.3. Drawbacks of density-ratio-based methods

Even though these methods have the advantage of being straightforward, they suffer from one main drawback. In practice, because both the target and the proposal manifold do not have full dimension in \mathbb{R}^D [9], [1, Lemma 3] show that it is highly likely that $\mu_\theta(S_{\mu_\theta} \cap S_{\mu_\star}) = 0$ and

$\mu_\star(S_{\mu_\theta} \cap S_{\mu_\star}) = 0$. Consequently, when dealing with high-dimensional datasets, the proposal distribution μ_θ and the target distribution μ_\star might intersect on a null set. Thus, one would have $r_\alpha(x) = 0$ almost everywhere on S_{μ_θ} . In this setting, the assumptions of MC methods are broken, and these algorithms will not allow sampling from μ_\star .

In order to correct this drawback, our method proposes to avoid the computation of density ratios from a classifier and to directly learn how to re-weight the proposal distribution. Our proposed scheme aims at minimizing the Wasserstein distance to the empirical measure while controlling the range of these importance weights.

3. Adversarial Learning of Latent Importance weights

Similar to previous works, our method aims at improving the performance of a generative model, post-training. We assume the existence of a WGAN model (G_θ, D_α) pre-trained using equation 1. The pushforward generative distribution μ_θ is assumed to be an imperfect approximation of the target distribution. The goal is now to learn how to redistribute the mass of the modeled distribution so that it best fits the target distribution.

3.1. Definition of the method

To improve the sampling quality of our pre-trained GANs, we propose to learn an importance weight function that directly learns how to avoid low-quality images and focus on very realistic ones. More formally, we over-parameterize the class of generative distributions and define a parametric class $\Omega = \{w^\varphi, \varphi \in \Phi\}$ of importance weight functions. Each function w^φ associates importance weights to latent space variables and is defined from \mathbb{R}^d to \mathbb{R}^+ . For a given latent space distribution γ and a network w^φ , a new measure γ^φ is defined on \mathbb{R}^d :

$$\text{for all } z \in \mathbb{R}^d, d\gamma^\varphi(z) = w^\varphi(z)d\gamma(z) \quad (4)$$

Using this formulation, we can prove the following lemma:

Lemma 1 *Assume that $\mathbb{E}_\gamma w^\varphi = 1$, then the measure γ^φ is a probability distribution defined on \mathbb{R}^d .*

Consequently, we now propose a new modeled generative distribution μ_θ^φ , the pushforward distribution $\mu_\theta^\varphi = G_\theta \# \gamma^\varphi$. The objective is to find the optimal importance weights w^φ that minimizes the Wasserstein distance between the true distribution μ_\star and the new class of generative distributions. The proposed method can thus be seen as minimizing the Wasserstein distance to the target distribution, over an increased class of generative distributions. Denoting by Lip_1 the set of 1-Lipschitz real-valued functions on \mathbb{R}^D , i.e.,

$$\text{Lip}_1 = \left\{ f : \mathbb{R}^D \rightarrow \mathbb{R} : \frac{|f(x) - f(y)|}{\|x - y\|} \leq 1, (x \neq y) \in (\mathbb{R}^D)^2 \right\},$$

we want, given a pre-trained model μ_θ , to solve:

$$\begin{aligned} \arg \min_{\varphi \in \Phi} W(\mu_*, \mu_\theta^\varphi) &= \arg \min_{w^\varphi \in \Omega} \sup_{D \in \text{Lip}_1} \mathbb{E}_{\mu_*} D - \mathbb{E}_{\mu_\theta^\varphi} D \\ &= \arg \min_{w^\varphi \in \Omega} \sup_{D \in \text{Lip}_1} \mathbb{E}_{\mu_*} D - \mathbb{E}_{\mu_\theta} w^\varphi D \end{aligned}$$

The network w^φ , parameterized using a feed-forward neural network, thus learns how to redistribute the mass of μ_θ such that μ_θ^φ is closer to μ_* in terms of Wasserstein distance. Similarly to the WGANs training, the discriminator D_α approximates the Wasserstein distance. D_α and w^φ are trained adversarially, whilst keeping the weights of G_θ frozen, using the following optimization scheme:

$$\inf_{\varphi \in \Phi} \sup_{\alpha \in \Lambda} \mathbb{E}_{x \sim \mu_*} D_\alpha(x) - \mathbb{E}_{z \sim Z} w^\varphi(z) \times D_\alpha(G_\theta(z)) \quad (5)$$

Note that our formulation can also be plugged on top of any objective function used for GANs.

3.2. Optimization procedure

However, as in the field of counterfactual estimation, a naive optimization of importance weights by gradient descent can lead to trivial solutions.

1. First, if for example, the Wasserstein critic D_α outputs negative values for any generated sample, the network w^φ could simply learn to avoid the dataset and output 0 everywhere [28].
2. Second, another problem comes from the fact that equation 5 can be minimized not only by putting large importance weights $w^\varphi(z)$ on the examples with high likelihoods $D_\alpha(G_\theta(z))$ but also by maximizing the sum of the weights: this is the propensity overfitting [29].
3. For the objective defined in equation 5 to be a valid Wasserstein distance minimization scheme, the measure μ_θ^φ must be a probability distribution, *i.e.* $\mathbb{E}_\gamma w^\varphi = 1$.

To tackle this, we first add a penalty term in the loss to enforce the expectation of the importance weights to be close to 1. This is similar to the self-normalization proposed by [29]. However, one still has to cope with the setting where the distribution γ^φ collapses to discrete data points:

Theorem 1 *Given a pre-trained generative distribution μ_θ absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^D . Let Φ be the non-parametric class of continuous functions satisfying $\mathbb{E}_\gamma w^\varphi = 1$. We have that:*

$$W(\mu_n, \frac{1}{n} \sum_{i=1}^n \delta(\tilde{X}_i)) \leq \inf_{\varphi \in \Phi} W(\mu_n, \mu_\theta^\varphi)$$

where δ refers to the Dirac probability distribution and $\tilde{X}_i = \arg \min_{x \in S_{\mu_\theta}} \|x - X_i\|$.

For clarity, the proof is delayed in Appendix. Intuitively, this theorem shows that the best way to approximate the empirical measure μ_n would be by considering a mixture of Diracs with each mode being the projection of a training data point on the support of the learned manifold S_{μ_θ} . The network w^φ could thus be tempted to approximate this mixture of Diracs defined in Theorem 1 and collapse on some specific latent data points. This could lead to an increased time complexity at inference (see [3, Section 3]). More importantly, this would mean *a mode collapse* and a lack of diversity in the generated samples.

To avoid such cases where small areas of z have really high $w^\varphi(z)$ values (mode collapse), we enforce a soft-clipping on the weights [5, 11]. Note that this constraint on $w^\varphi(z)$ could also be implemented with a bounded activation function on the final layer, such as a re-scaled sigmoid or tanh activation. Finally, we get the following objective function for the network w^φ :

$$\begin{aligned} \sup_{\varphi \in \Phi} \mathbb{E}_{z \sim Z} \underbrace{w^\varphi(z) (D_\alpha(G_\theta(z)) - \Delta)}_{\text{discriminator reward}} - \lambda_1 \underbrace{(\mathbb{E}_{z \sim Z} w^\varphi(z) - 1)^2}_{\text{self-normalization}} \\ - \lambda_2 \underbrace{\mathbb{E}_{z \sim Z} \max(0, (w^\varphi(z) - m))^2}_{\text{soft-clipping}}, \quad (6) \end{aligned}$$

where $\Delta = \min_{z \sim Z} D_\alpha(G_\theta(z))$. λ_1 , λ_2 , and m are hyper-parameters (values displayed in Appendix). For more details, we refer the reader to Algorithm 1.

Algorithm 1: Adversarial learning of w^φ

Require: Data μ_n , Prior Z , Gen. G_θ , Disc. D_α , number of D_α updates n_d , soft-clipping param. m , regularization weights λ_1 and λ_2 , batch size b ;

while φ has not converged **do**

for $i = 0, \dots, n_d$ **do**

Sample real data $\{x_i\}_{i=1}^b \sim \mu_n$;

Sample latent vectors $\{z_i\}_{i=1}^b \sim Z$;

$\text{EMD} \leftarrow \frac{1}{b} \sum_{i=1}^b D_\alpha(x_i) - w^\varphi(z_i) D_\alpha(G_\theta(z_i))$;

$\text{GP} \leftarrow \text{Gradient-Penalty}(D_\alpha, x, G_\theta(z))$;

$\text{grad}_\alpha \leftarrow \nabla_\alpha(-\text{EMD} + \text{GP})$;

Update α with grad_α ;

end

Sample $\{z_i\}_{i=1}^b \sim Z$;

$\Delta \leftarrow \min_i [D_\alpha(G_\theta(z_i))]$;

$\text{EMD} \leftarrow \frac{1}{b} \sum_{i=1}^b w(z_i) [D_\alpha(G_\theta(z_i)) - \Delta]$;

$R_{\text{norm}} \leftarrow (\frac{1}{b} \sum_{i=1}^b w(z_i) - 1)^2$;

$R_{\text{clip}} \leftarrow \frac{1}{b} \sum_{i=1}^b \max(0, w^\varphi(z_i) - m)^2$;

$\text{grad}_\varphi \leftarrow \nabla_\varphi(\text{EMD} + \lambda_1 R_{\text{norm}} + \lambda_2 R_{\text{clip}})$;

Update φ with grad_φ ;

end

3.3. Sampling from the latent importance weights

Given a pre-trained generator G_θ and an importance network w^ϕ , we now present the three proposed sampling algorithms associated with our model:

1) Latent Rejection Sampling (latentRS, Algorithm 2).

The first proposed method aims at sampling from the newly learned latent distribution γ^ϕ defined in equation 4. Since the learned importance weights are capped by m defined in equation 6, this setting fits in the Rejection Sampling (RS) algorithm [25]. Any sample $z \sim \gamma$ is now accepted with probability $\mathbb{P}_a(z) = w^\phi(z)/m$. Interestingly, by actively capping the importance weights as it is done in counterfactual estimation [5, 8], one controls the acceptance rates $\mathbb{P}_a(z)$ of the rejection sampling algorithm:

$$\mathbb{E}_\gamma \mathbb{P}_a(z) = \int_{\mathbb{R}^d} \frac{w^\phi(z)}{m} d\gamma(z) = \frac{1}{m}.$$

Algorithm 2: LatentRS

Requires: Prior Z , Gen. G_θ , Importance weight network w^ϕ , maximum importance weight m ;

while True do

 Sample $z \sim Z$;

 Sample $\alpha \sim \text{Uniform}[0, 1]$;

if $\frac{w^\phi(z)}{m} \geq \alpha$ **then**

 break;

end

end

$x \leftarrow G_\theta(z)$;

Result: Selected point x

2) Latent Gradient Ascent (latentGA). Inspired from [30, Algorithm 2], we propose a second method, latentGA, where we perform gradient ascent in the latent space (see the algorithm in Appendix). For any given sample in the latent space, we follow the path maximizing the learned importance weights. This method is denoted latentGA. Note that the learning rate and the number of updates used for this method are hyper-parameters that need to be tuned.

3) Combining latentRS with Gradient Ascent (latentRS+GA, see Appendix).

Finally, we propose to combine sequentially both methods. In a first step, we avoid low-quality samples with latentRS. Then, we use latentGA to further improve the remaining generated samples. See algorithm in Appendix.

3.4. Advantages of the proposed approach

We now discuss two advantages of our method compared to previous density-ratio-based Monte-Carlo methods.

Computational cost. By using sampling algorithms in the latent space, we avoid going through both the generator and the discriminator, leading to a significant computational speed-up. This is of particular interest when dealing with high-dimensional spaces, since we do not need to pass through deep CNNs generator and discriminator [17]. In the next experimental section, we observe a computational cost decreased by a factor of 10.

Monte-Carlo methods do not properly work when the support S_{μ_θ} does not fully cover the support S_{μ_*} .

To better illustrate this claim, we consider a simple 2D motivational example where the real data lies on four disconnected manifolds. We start with a proposal distribution (in blue) that does not fully recover the target distribution (Figure 2a). In this setting, we see in Figure 2b that the discriminator’s density-ratio-based methods [3] avoids half of the proposal distribution, while our proposed method learns a very different re-weighting (see Figure 2c).

This illustration is important since [1, Theorem 2.2] have shown that in high-dimension the intersection $S_{\mu_*} \cap S_{\mu_\theta}$ is likely to be a negligible set under μ_θ . Knowing that S_{μ_θ} does not fully recover S_{μ_*} , there is thus no theoretical guarantee that using a sampling algorithm will improve the estimation of μ_* . On the opposite, our method looks for the optimal re-weighting of μ_θ under a well-defined criterion: the Wasserstein distance. This results in a better fit of the real data distribution (see next section).

4. Experiments

In this section, we illustrate the efficiency of the proposed methods, latentRS, latentGA, and latentRS+GA on both synthetic and natural image datasets. On image generation tasks, we empirically stress that latentRS slightly surpasses density-ratio-based methods with respect to the Earth Mover’s distance while reducing the time complexity by a factor of around 10. The use of latentGA also gives interesting experimental visualizations and improves image quality. More importantly, when combined, we show that latentRS+GA surpasses the concurrent methods, while still being less computationally intensive. Finally, we show results with different models such as Progressive GAN [16] and StyleGAN2 [18].

4.1. Evaluation metrics

To measure the performances of GANs when dealing with low-dimensional applications, we equip our space with the standard Euclidean distance. However, for the case of image generation, we follow [6, 20] and consider the euclidean distance between embeddings of a pre-trained network, that convey more semantic information. Thus, for a pair of images (a, b) , we define the distance $d(a, b)$ as




Figure 2: Synthetic experiment mimicking the setting of GANs in high-dimension, where data and generated manifolds are close but do not perfectly intersect. While DRS only selects the intersection of manifolds and ignores the rest, the latent importance weights define a rejection mechanism that minimizes the Wasserstein distance. For conciseness, WGAN stands for WGAN-GP.

	EMD Swiss Roll	EMD 25 Gaussians
WGAN	0.030±0.002	0.044±0.001
WGAN: DRS	0.036±0.004	0.038±0.002
WGAN: SIR	0.037±0.003	0.041±0.001
WGAN: DOT	0.029±0.003	0.035±0.002
WGAN: latentRS (*)	0.025±0.002	0.036±0.001

Table 1: Comparison of latentRS with concurrent methods on two synthetic datasets in the same setting as DOT [30]. Our method enables a consistent gain in EMD, surpassing other methods on Swiss Roll and slightly behind DOT on Mixture of 25 Gaussians. For conciseness, WGAN stands for WGAN-GP.

$d(a, b) = \|\phi(a) - \phi(b)\|_2$ where ϕ is a pre-softmax layer of a supervised classifier. On MNIST and F-MNIST, the classifier is pre-trained on the given dataset. On CelebA and LSUN Church, we use VGG-16 pre-trained on ImageNet.

To begin with, we report the FID [15]. We also compare the performance of the different methods with the Precision/Recall (PR) metric [20]. It is a more robust version of the Precision/Recall metric, which was first applied in the context of GANs by [27]. Finally, we approximate the Wasserstein distance using the Earth Mover’s Distance (EMD) between generated and real data points. This measure is particularly suited to the study of WGANs, since it is linked to their objective function. Letting $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_n\}$ be two collections of n data points and \mathcal{S} be the set of permutations of $[1, n]$, the Earth Mover’s distance between X and Y is defined by:

$$\text{EMD}(X, Y) = \min_{\sigma \in \mathcal{S}} \sum_{i=1}^n \|x_i - y_{\sigma_i}\|$$

4.2. Synthetic datasets

To begin the experimental study, we test our method on 2D synthetic datasets in the same setting as [30]. Table 1 compares the latentRS method with previous approaches

on the Swiss roll dataset and on a mixture of 25 Gaussians. We see that the network w^ϕ efficiently redistributes the pre-trained distribution μ_θ since $\text{EMD}(\mu_n, \mu_\theta^\phi)$ is significantly smaller than $\text{EMD}(\mu_n, \mu_\theta)$.

4.3. Image datasets

Implementation of baselines. We now compare latentRS, latentGA, and latentRS+GA with previous works leveraging discriminator’s information on high-dimensional data. In particular, we implemented a wide set of post-processing methods for GANs: DRS [3], MH-GAN [32], SIR [11] and DOT [30]. DRS, MH-GAN and SIR use the same density ratios, and we did not see significant differences between those three methods in our experiments. Consequently, for the following experiments, we compare our algorithms to SIR and DOT. For SIR, we take the discriminator at the end of the adversarial training, fine-tune it with the binary cross-entropy loss and select the best model in terms of EMD. Overall, we explicitly follow the framework used by [3, 11]: we keep the gradient penalty [12], spectral normalization [24] during fine-tuning and do not include an explicit mechanism to calibrate the classifier.

4.3.1 Description of datasets and neural architectures

We first consider two well-known image datasets that are MNIST [21] and FashionMNIST (F-MNIST). We follow [19] and use a standard CNN architecture composed of a sequence of blocks made of 3x3 convolution layer and ReLU activations with nearest neighbor upsampling. For these datasets, the discriminator is trained using the hinge loss [22] with gradient penalty (Hinge-GP). Finally, the architecture used for the network w^ϕ is very simple: an MLP with 4 fully-connected layers and ReLU activation (with a width = $4 \times d$).

CelebA [23] is a large-scale dataset of faces covering a variety of poses. We use a pre-trained model of Progressive GAN [16] at 128x128 resolution. The discriminator is

CelebA 128x128	Prec. (\uparrow)	Rec. (\uparrow)	EMD (\downarrow)	FID (\downarrow)	Inference (ms)
ProGAN	74.2 \pm 0.9	60.7 \pm 1.4	25.4 \pm 0.1	11.30 \pm 0.02	3.6
ProGAN: SIR	79.5 \pm 0.4	57.3\pm1.0	24.9 \pm 0.2	12.01 \pm 0.04	49.0
ProGAN: DOT	81.3 \pm 1.0	52.9 \pm 1.4	25.0 \pm 0.1	11.01 \pm 0.03	67.6
ProGAN: latentRS (\star)	80.4 \pm 0.9	55.7 \pm 1.0	24.7 \pm 0.1	10.77 \pm 0.04	4.5
ProGAN: latentRS+GA (\star)	83.3\pm1.0	52.7 \pm 0.9	24.5\pm0.1	10.75\pm0.04	20.5
LSUN Church 256x256					
StyleGAN2	55.6 \pm 1.2	62.4 \pm 1.1	23.6 \pm 0.1	6.91 \pm 0.02	11.7
StyleGAN2: SIR	60.5 \pm 1.4	58.1\pm1.3	23.4 \pm 0.1	7.36 \pm 0.01	130.0
StyleGAN2: DOT	67.4 \pm 1.4	48.3 \pm 1.0	23.1 \pm 0.1	6.85 \pm 0.02	196.7
StyleGAN2: latentRS (\star)	63.3 \pm 0.7	57.7 \pm 1.0	23.1 \pm 0.1	6.31 \pm 0.02	16.2
StyleGAN2: latentRS+GA (\star)	72.6\pm1.1	43.2 \pm 1.3	22.6\pm0.1	6.27\pm0.03	43.2

Table 2: latentRS+GA is the best performer, and latentRS matches SOTA with a significantly reduced inference cost (by an order of at least 10). \pm is 97% confidence interval. Inference refers to the time in milliseconds needed to compute one image on a NVIDIA V100 GPU.

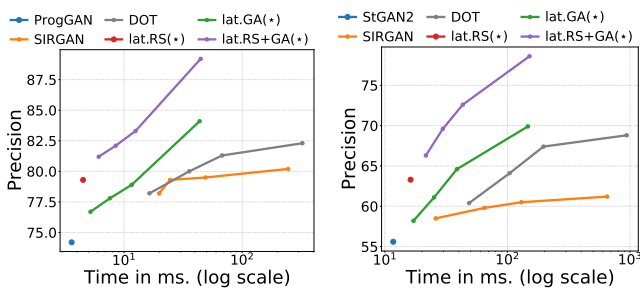


Figure 3: Visualization of the trade-off between the time spent to generate an image and its average precision. Interestingly, latentRS+GA has the best Pareto front. Left: ProGAN trained on CelebA. Right: StyleGAN2 trained on LSUN Church.

trained using a Wasserstein loss with gradient-penalty. Also, the architecture used for the network w^ϕ is really standard: a 5 hidden-layer MLP with a width of the same size than the latent space dimension.

LSUN Church [33] is a dataset of church images with a lot of variety. We use a pre-trained model of StyleGAN2 [18] at 256x256 resolution. Similarly to the CelebA dataset, the discriminator is trained using a Wasserstein loss with gradient-penalty. Also, the architecture used for the network w^ϕ is a 3 hidden-layer MLP with width equal to the latent space dimension. Note that the StyleGAN architecture already contains an 8-layer MLP network $M_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^d$ that transforms a latent space variable to an intermediate latent variable [17]. We consequently leverage this pre-trained M_θ and train the network w^ϕ on top of it.

4.3.2 Results

The main results of this comparison are shown in Table 2 and Figure 3. On all studied datasets, our latentRS+GA

outperforms every other method on the EMD with lower computational cost. Interestingly, latentRS achieves good performance on FID while being more than 15 times faster. Figure 3 is particularly interesting since it gives a good visualization of the trade-off between computational cost and quality of the generated samples. On this experiment ran on CelebA and LSUN, we observe that latentRS+GA can achieve a significantly better precision than both SIR and DOT while being much faster. Interestingly, even though these datasets are high-dimensional, contain only one-class, and w^ϕ has a low capacity, our proposed methods still produce interesting results.

To visualize the efficiency of the proposed method, Figure 4 shows generated samples along with their acceptance probabilities. As expected, we observe that higher acceptance probabilities correlate with higher quality images. Figure 5 stresses how generated images improve when performing latent gradient ascent on the importance weights. Finally, we provide more qualitative results and details on the experiments in supplementary material.

5. Conclusion

This paper deals with improving the quality of pre-trained GANs. Conversely, to concurrent methods which leverage the discriminator at inference time, we propose to train adversarially a neural network which learns importance weights in the latent space of GANs. These latent importance weights are then used with two complementary sampling methods: latentRS and latentGA. We experimentally show that this latent reweighting consistently enhances the quality of the pre-trained model. When these two methods are combined in latentRS+GA, it surpasses concurrent post-training methods while being less computationally intensive.




Figure 4: Images drawn from the generative model and their acceptance probabilities with the latentRS algorithm, given by the network w^ϕ . As expected, the quality of images correlates with higher acceptance rates on all datasets: MNIST, F-MNIST, CelebA, and LSUN.



Figure 5: Gradient ascent on latent importance weights (latentGA): the quality is gradually improved as we move to larger importance weights. Each image is generated only for visualization, and one can run this gradient ascent directly in the latent space using w^ϕ . Interestingly, this gradient ascent only involves a simple MLP network which is computationally cheap.

References

- [1] M. Arjovsky and L. Bottou. Towards principled methods for training generative adversarial networks. In *International Conference on Learning Representations*, 2017. 1, 2, 3, 5
- [2] M. Arjovsky, S. Chintala, and L. Bottou. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pages 214–223. PMLR, 2017. 2
- [3] Samaneh Azadi, Catherine Olsson, Trevor Darrell, Ian Goodfellow, and Augustus Odena. Discriminator rejection sampling. In *International Conference on Learning Representations*, 2019. 1, 3, 4, 5, 6
- [4] Yogesh Balaji, Rama Chellappa, and Soheil Feizi. Robust optimal transport with applications in generative modeling and domain adaptation. *Advances in Neural Information Processing Systems*, 33, 2020. 2
- [5] Léon Bottou, Jonas Peters, Joaquin Quinonero Candela, Denis Xavier Charles, Max Chickering, Elon Portugaly, Dipankar Ray, Patrice Y Simard, and Ed Snelson. Counterfactual Reasoning and Learning Systems: the example of Computational Advertising. *Journal of Machine Learning Research*, 14(1):3207–3260, 2013. 4, 5
- [6] A. Brock, J. Donahue, and K. Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019. 1, 5
- [7] 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. In *Advances in Neural Information Processing Systems*, volume 33, pages 12275–12287. Curran Associates, Inc., 2020. 3
- [8] L. Faury, U. Tanielian, E. Dohmatob, E. Smirnova, and F. Vasile. Distributionally robust counterfactual risk minimization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*, pages 3850–3857. AAAI Press, 2020. 5
- [9] C. Fefferman, S. Mitter, and H. Narayanan. Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29:983–1049, 2016. 3
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and J. Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014. 1, 2
- [11] Aditya Grover, Jiaming Song, Ashish Kapoor, Kenneth Tran, Alekh Agarwal, Eric J Horvitz, and Stefano Ermon. Bias correction of learned generative models using likelihood-free importance weighting. In *Advances in Neural Information Processing Systems*, pages 11056–11068, 2019. 3, 4, 6
- [12] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A.C. Courville. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30*, pages 5767–5777. Curran Associates, Inc., 2017. 6
- [13] Swaminathan Gurumurthy, Ravi Kiran Sarvadevabhatla, and R Venkatesh Babu. Deligan: Generative adversarial networks for diverse and limited data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 166–174, 2017. 1, 2
- [14] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970. 3
- [15] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pages 6626–6637, 2017. 6
- [16] T. Karras, T. Aila, S. Laine, and J. Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. 5, 6
- [17] T. Karras, S. Laine, and T. Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2, 5, 7
- [18] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 5, 7
- [19] Mahyar Khayatkhoei, Maneesh K Singh, and Ahmed Elgammal. Disconnected manifold learning for generative adversarial networks. In *Advances in Neural Information Processing Systems*, pages 7343–7353, 2018. 1, 2, 6
- [20] T. Kynkäänniemi, T. Karras, S. Laine, J. Lehtinen, and T. Aila. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems 32*, pages 3927–3936. Curran Associates, Inc., 2019. 2, 5, 6
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998. 6

- [22] Jae Hyun Lim and Jong Chul Ye. Geometric gan. *arXiv preprint arXiv:1705.02894*, 2017. 6
- [23] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 6
- [24] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018. 6
- [25] Christian Robert and George Casella. *Monte Carlo statistical methods*. Springer Science & Business Media, 2013. 3, 5
- [26] Donald B Rubin. Using the sir algorithm to simulate posterior distributions. *Bayesian statistics*, 3:395–402, 1988. 3
- [27] M.S.M. Sajjadi, O. Bachem, M. Lucic, O. Bousquet, and S. Gelly. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems 31*, pages 5228–5237. Curran Associates, Inc., 2018. 2, 6
- [28] Adith Swaminathan and Thorsten Joachims. Batch Learning from Logged Bandit Feedback through Counterfactual Risk Minimization. *Journal of Machine Learning Research*, 16(1):1731–1755, 2015. 4
- [29] A. Swaminathan and T. Joachims. The self-normalized estimator for counterfactual learning. In *Advances in Neural Information Processing Systems 28*, pages 3231–3239. Curran Associates, Inc., 2015. 4
- [30] Akinori Tanaka. Discriminator optimal transport. In *Advances in Neural Information Processing Systems*, pages 6813–6823, 2019. 3, 5, 6
- [31] U. Tanielian, T. Issenhuth, E. Dohmatob, and J. Mary. Learning disconnected manifolds: a no gan’s land. In *International Conference on Machine Learning*, 2020. 1, 2, 3
- [32] Ryan Turner, Jane Hung, Eric Frank, Yunus Saatchi, and Jason Yosinski. Metropolis-hastings generative adversarial networks. In *International Conference on Machine Learning*, pages 6345–6353, 2019. 1, 3, 6
- [33] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 7