

Exploring and Exploiting Hubness Priors for High-Quality GAN Latent Sampling

Yuanbang Liang¹ Jing Wu¹ Yu-Kun Lai¹ Yipeng Qin¹

Abstract

Despite the extensive studies on Generative Adversarial Networks (GANs), how to reliably sample high-quality images from their latent spaces remains an under-explored topic. In this paper, we propose a novel GAN latent sampling method by exploring and exploiting the *hubness priors* of GAN latent distributions. Our key insight is that the high dimensionality of the GAN latent space will inevitably lead to the emergence of *hub* latents that usually have much larger sampling densities than other latents in the latent space. As a result, these *hub* latents are better trained and thus contribute more to the synthesis of high-quality images. Unlike the a posteriori “cherry-picking”, our method is highly efficient as it is an a priori method that identifies high-quality latents before the synthesis of images. Furthermore, we show that the well-known but purely empirical truncation trick is a naive approximation of the central clustering effect of *hub* latents, which not only uncovers the rationale of the truncation trick, but also indicates the superiority and fundamentality of our method. Extensive experimental results demonstrate the effectiveness of the proposed method. Our code is available at: <https://github.com/Byronliang8/HubnessGANSampling>.

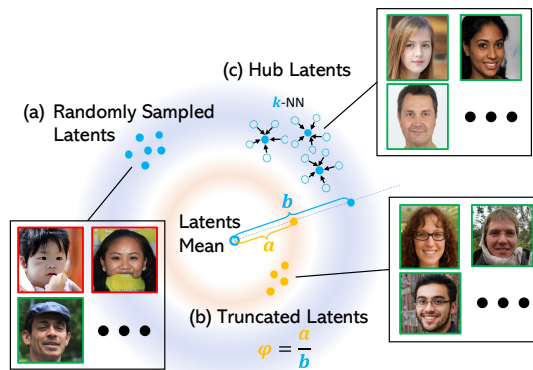


Figure 1. Our method vs. random latent sampling and the truncation trick (Marchesi, 2017; Brock et al., 2019; Karras et al., 2019). All images are generated using StyleGAN2 (Karras et al., 2020). (a) Random latent sampling yields both high-quality (green box) and low-quality (red box) images; (b) The truncation trick improves the quality of synthesized images by empirically truncating randomly sampled latents according to a scaling parameter ψ (e.g. $\psi = 0.7$), which is a naive approximation of the “central clustering effect” of our *hub* latents; (c) Our method identifies high-quality latents as the *hub* latents that are more likely to be among the k -nearest neighbors of other latents (Radovanovic et al., 2010). The blue and orange rings illustrate the high-dimensional Gaussian (latent) distribution (Menon et al., 2020) and their truncated version respectively.

1. Introduction

Generative adversarial networks (GANs) are a type of deep generative models that have revolutionized a variety of applications in computer vision and computer graphics, e.g. image synthesis (Karras et al., 2019; Park et al., 2019; Zhu et al., 2020), image editing (Abdal et al., 2019; 2020; Tov et al., 2021), image-to-image translation (Isola et al., 2017;

¹School of Computer Science and Informatics, Cardiff University, Cardiff, CF24 4AG, UK. Correspondence to: Yipeng Qin <qiny16@cardiff.ac.uk>.

Zhu et al., 2017; Richardson et al., 2021). Among them, novel image synthesis via random latent sampling is the most fundamental. It not only generates novel instances from the data distribution, but also measures how close the learned distribution is to the data distribution. Through the lens of the quality of synthesized images, we have witnessed significant progress in GANs over the past several years. Specifically, starting from the groundbreaking vanilla GAN (Goodfellow et al., 2014), DCGAN (Radford et al., 2015) laid the foundation for GAN architectures as deep convolutional neural networks; ProGAN (Karras et al., 2018) showed that GANs can generate high-quality images at high resolutions; BigGAN (Brock et al., 2019) addressed the problem of class-conditional image synthesis; the StyleGAN series (Karras et al., 2019; 2020; 2021) further boosted

the quality and controllability of synthesized images with their style-based generator architectures and several novel techniques.

Nevertheless, with such improvements, the quality variance among images generated by randomly sampled latents has become increasingly striking (Fig. 1). Without curation, the quality of GAN synthesized images can occasionally be very low, which hinders the deployment of GANs in real-world applications. As a naive solution, “cherry-picking” is commonly used to select high-quality images from those synthesized with randomly sampled latents in an *a posteriori* manner. However, in the absence of reliable quantitative measures of the quality of a single GAN-synthesized image¹, existing “cherry-picking” methods are barely manual, thereby being tedious and unscalable. Addressing this issue, the well-known “truncation trick” (Marchesi, 2017; Brock et al., 2019; Karras et al., 2019) was proposed, which “truncates” randomly sampled GAN latents towards their mean based on the observation that the images synthesized from close-to-mean latents are usually of higher quality. Although effective, the truncation trick is a purely empirical “trick” that brings few new insights to the community.

In this paper, we propose a novel latent sampling method for GANs by exploring and exploiting the *hubness* phenomenon (Radovanovic et al., 2010) in their latent spaces, which facilitates their synthesis of high-quality images in an *a priori* manner. Specifically, our key insights include: i) the high dimensionality of the GAN latent space will inevitably lead to the emergence of *hub* latents that are much more likely to be among the nearest neighbors of other latents in the latent space, *i.e.* the hubness phenomenon; ii) in general, the quality of a GAN synthesized image is positively correlated with the *hub value* of its corresponding latent, *i.e.* the number of times a latent becomes a k -nearest neighbor (k -NN) of other latents in a given latent sample set. We believe that this positive correlation originates from the well-known close relationship between k -NN and density estimation. In other words, a higher *hub value* usually indicates a higher sampling density, which has a positive effect on the training and thus the quality of synthesized images. Therefore, we formulate the above insights as the proposed *hubness priors* and propose a corresponding method to sample high-quality GAN latents that yield high-quality synthesized images. Compared to “cherry-picking”, our method is highly efficient as it is *a priori* (*i.e.* our high-quality latents are determined before the synthesis of images) and automatic (*i.e.* with little human-intervention). Furthermore, we show that the well-known truncation trick is a naive approximation of the “central clustering effect” of our *hub* latents (Radovanovic et al., 2010). This not only uncovers

¹Existing quantitative measures like FID and Inception scores are all statistical ones that are only applicable to distributions.

the rationale of the truncation trick, but also indicates that our method is superior and more fundamental. Extensive experimental results demonstrate the effectiveness of the proposed method.

In summary, our contributions include:

- We uncover the existence of *hubness* phenomenon in the GAN latent space, which has a significant correlation with the quality of GAN synthesized images, *i.e.* the proposed *hubness priors*.
- We propose a novel GAN latent sampling algorithm that identifies high-quality *hub* latents based on our *hubness priors*, which allows efficient and high-quality image synthesis for GANs.
- We show that the well-known truncation trick is a naive approximation of the “central clustering effect” of our *hub* latents. This not only uncovers the rationale of the truncation trick, but also indicates that our method is superior and more fundamental.

2. Related Work

2.1. Generative Adversarial Network (GAN)

Since the seminal work of Goodfellow et al. (2014), Generative Adversarial Networks (GANs) have become a major type of deep generative models and have been extensively studied in recent years. Existing works mostly focus on the choices of architectures (Radford et al., 2015; He et al., 2016; Jiang et al., 2021), loss functions (Arjovsky et al., 2017; Mao et al., 2017), regularization and normalization techniques (Gulrajani et al., 2017; Mescheder et al., 2018; Miyato et al., 2018; Qin et al., 2020), aiming to stabilize the training of GANs and improve the quality of synthesized images. To date, the best-performing GANs include the ProGAN (Karras et al., 2018) and the StyleGAN series (Karras et al., 2019; 2020; 2021) developed by Nvidia for unconditional image synthesis, and the BigGAN (Brock et al., 2019) developed by DeepMind for conditional image synthesis. Despite the success of these methods, they all have the long-standing problem that a large proportion of the images synthesized from randomly sampled latents are low-quality ones with artifacts, which hinders their applications in visual content generation.

GAN Latent Sampling. To sidestep the above-mentioned problem and obtain high-quality synthesized images, two workaround solutions were proposed: i) “cherry-picking” and ii) the truncation trick (Marchesi, 2017; Brock et al., 2019; Karras et al., 2019). Between them, the first approach is a naive solution as one can always “cherry-pick” high-quality ones from a set of synthesized images in an *a*

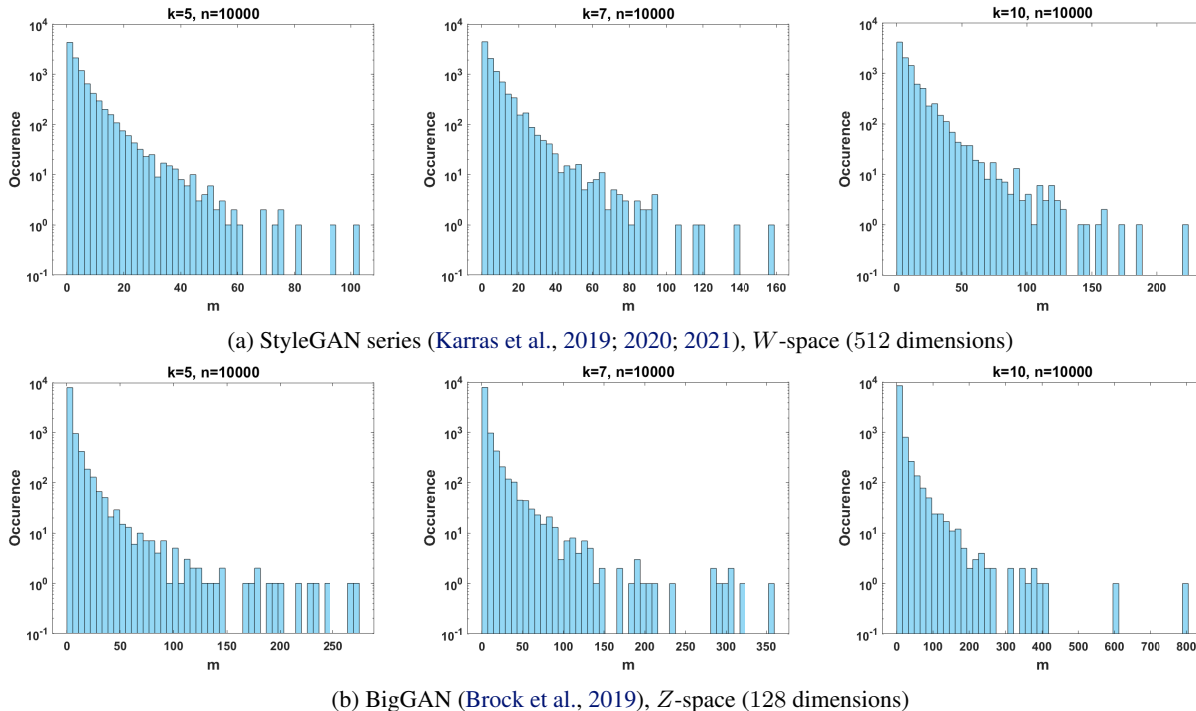


Figure 2. Distributions of m -hub latents for state-of-the-art GANs, $k = 5, 7, 10$ (the k -NN algorithm) and $n = 10000$ (size of latent sample set S). All distributions are highly tailed to the right, which shows the existence of hubness phenomenon (Radovanovic et al., 2010) in GAN latent spaces. Similar phenomena hold for StyleGAN Z -space and ProGAN (Appendix A). Note that y -axis is in log-scale.

posteriori manner with visual inspection. Obviously, this method is inefficient as it requires intensive human labor and is not applicable for large-scale image synthesis tasks. Unlike “cherry-picking”, the truncation trick is an automatic method that can synthesize high-quality images by normalizing sampled latents to be close to their mean. However, it is a purely *empirical* method with few insights. In this paper, we propose a novel latent sampling method for GANs based on the observation of *hubness* phenomenon in their high dimensional latent spaces, which is efficient with solid theoretical insights. We also show that the truncation trick is a naive approximation of our method due to the “central clustering effect” of *hub* latents.

2.2. Hubness Phenomenon

Hubness is a well-known phenomenon that describes the impact of the notorious “curse of dimensionality” on nearest neighbors (Radovanovic et al., 2010). In a nutshell, the hubness phenomenon is proved to be an inherent property of data distributions in high-dimensional space (Newman et al., 1983; Newman & Rinott, 1985; Radovanovic et al., 2010), which uncovers an interesting but counter-intuitive fact: high dimensionality leads to the emergence of “popular” nearest neighbors (*a.k.a.* the *hub* points). In other words, the hub points are those that are much more likely to be among the k -nearest neighbours of other points in

a sample set. This fact poses challenges for algorithms that rely on nearest neighbor search. Addressing such challenges, hubness-aware methods were proposed and applied in various areas, *e.g.* gene expression classification (Buza, 2016a;b), time-series classification (Tomašev et al., 2015) and electroencephalograph classification (Buza & Koller, 2016). Meanwhile, hubness-aware k -nearest neighbor (k -NN) methods were also proposed, *e.g.* hubness-weighted k -NN (Radovanovic et al., 2010), hubness-fuzzy k -NN (Tomašev et al., 2014), hubness-information k -NN (Tomašev & Mladenec, 2011), Naive Hubness-Bayesian k -NN (Tomašev et al., 2011), and Augmented Naive Hubness-Bayesian k -NN (Tomašev & Mladenec, 2013).

In this paper, in contrast to previous methods that treat hubness as an undesirable phenomenon and aim to mitigate it, we show that the hubness phenomenon can be effectively used as priors for the sampling of high-quality GAN latents that produce high-quality synthesized images.

3. Hubness Priors for GAN Latent Sampling

In this section, we first explore the hubness of GAN latents (Section 3.1) and then exploit the insights obtained as priors to develop a novel algorithm for the sampling of high-quality latents for GANs (Section 3.2).



(a) Baseline (random latent sampling) (b) Our method (W -space hubness priors) (c) Low-quality latent sampling (Alg. 2)

Figure 3. (a) and (b): Effectiveness of our method (hubness priors) against the baseline (random latent sampling). We use $n = 10000$, $k = 5$ and hub value threshold $t = 50$ in our method. The StyleGAN2 (Karras et al., 2020) images generated using our method are almost always of high quality while those generated using the baseline contain both high-quality and low-quality (red boxes) results. (c): Low-quality StyleGAN2 (Karras et al., 2020) images generated using the **reversed** version of our method, *i.e.* Algorithm 2 (Appendix B), where $n = 10000$, $k = 5$ and hub value threshold $t_{lq} = 1$. Almost all images are of low quality.

3.1. Exploring Hubness of GAN Latents

Inspired by previous studies on the hubness phenomenon of data distributions in high dimensional space (Radovanovic et al., 2010), let $Z \in \mathbb{R}^d$ be a d -dimensional GAN latent space, $S = \{z_1, z_2, \dots, z_n\}$, $z_i \in Z$ be a set of latents sampled from a d -dimensional standard normal distribution $\mathcal{N}(0, I)$, k be the parameter of the k -nearest neighbor algorithm, we define m -hub latents as:

Definition 3.1. Latent code z_i ($1 \leq i \leq n$) is an m -hub latent if z_i is among the k -nearest neighbors of m ($m < n$) sampled latents in S , where m is the *hub value* of z_i .

With the above definition, we explore the hubness of GAN latents by investigating the distributions of m -hub latents in the latent spaces of state-of-the-art GANs (Brock et al., 2019; Karras et al., 2018; 2019; 2020; 2021). As Fig. 2 shows, it can be observed that the distributions of m -hub latents are highly tailed to the right. Thus, we argue that the GAN latents are not uniformly distributed and that a small portion of them are much more likely to be close to other latents in the latent space, *i.e.*, with large m . Therefore, these latents tend to have larger sampling densities and are thus better trained than other latents during GAN training. Based on the heuristics that well-trained latents are more likely to yield high-quality images, we conjecture that the hubness phenomenon can be used as priors to identify GAN latents that generate high-quality results:

Conjecture 3.2. (Hubness Priors) The quality of GAN synthesized images and the hub values m of their corresponding latents are positively correlated.

Please see Section 4.2 for an empirical justification of our conjecture.

Algorithm 1 GAN Latent Sampling with Hubness Priors

Input: a set of GAN latents $S = \{z_1, z_2, \dots, z_n\}$ randomly sampled from a standard normal distribution $\mathcal{N}(0, I)$, a hyper-parameter k , a threshold t

Output: S_{hq}

Step 1

$m_{1,2,\dots,n} \leftarrow 0$

for $i \leftarrow 1$ to n **do**

$\{\text{idx}_1, \text{idx}_2, \dots, \text{idx}_k\} \leftarrow k\text{-NN}(z_i)$

for $j \leftarrow 1$ to k **do**

$m_{\text{idx}_j} \leftarrow m_{\text{idx}_j} + 1$

end for

end for

Step 2

$S_{hq} \leftarrow \emptyset$

for $i \leftarrow 1$ to n **do**

if $m_i \geq t$ **then**

$S_{hq} \leftarrow S_{hq} \cup z_i$

end if

end for

Remark on Random Latent Sampling Previously, it was widely believed that GAN latents are *unbiased* as they are sampled from a simple but well-behaved noise distribution, *i.e.*, the standard normal distribution $\mathcal{N}(0, I)$, which is isotropic and has most of its density on a hypersphere surface in the high-dimensional GAN latent space (Menon et al., 2020). This implies that all sampled latents should be approximately uniformly distributed and of similar norms², thereby contributing to the sampling in a similar manner. While in this paper, we counter this popular belief by show-

²In latest implementations (Karras et al., 2019; 2020; 2021), the latents are explicitly normalized to be of the same norm.



(a) ProGAN-HQ


(b) ProGAN-LQ




(c) BigGAN-HQ



(d) BigGAN-LQ



(e) StyleGAN3-HQ




(f) StyleGAN3-LQ

Figure 4. Performance of our method on ProGAN (Karras et al., 2018), BigGAN (Brock et al., 2019) and StyleGAN3 (Karras et al., 2021). It can be observed that our method works well on all GAN architectures. (a) and (b), (c) and (d), (e) and (f) are images synthesized using high-quality (HQ) and low-quality (LQ) latents obtained by our method with ProGAN, BigGAN and StyleGAN3 respectively. We use Algorithm 1 to obtain HQ latents and Algorithm 2 (Appendix B) to obtain LQ latents respectively. We use $n = 10000$, $k = 5$, $t = 50$ and $t_{lq} = 1$.


ing that GAN latents are actually *biased* from the observation of hubness phenomenon in GAN latent spaces. Among all latents, the *hub* ones tend to have higher sampling densities and are thus better trained by GANs, thereby generating higher quality images.

3.2. Exploiting Hubness Priors for High-quality GAN Latent Sampling


As Conjecture 3.2 states, the identification of high-quality GAN latents relies on their hub values m . Thus, given a set of GAN latents $S = \{z_1, z_2, \dots, z_n\}$ randomly sampled from a standard normal distribution $\mathcal{N}(0, I)$, a hyperparameter k , and a threshold t , we utilize the proposed hubness priors and design a simple two-step GAN latent sampling algorithm: First, we compute the hub value m_i for




(a) StyleGAN-Car-HQ



(b) StyleGAN-Car-LQ



(c) StyleGAN-Cat-HQ



(d) StyleGAN-Cat-LQ

Figure 5. Performance of our method on StyleGAN2 pretrained on different image domains. It can be observed that our method works well on all domains. (a) and (b), (c) and (d) are images synthesized from high-quality (HQ) and low-quality (LQ) latents obtained by our method using a StyleGAN2 pretrained on the cars domain and the cats domain, respectively. We use Algorithm 1 to obtain HQ latents and Algorithm 2 (Appendix B) to obtain LQ latents respectively. We use $n = 10000$, $k = 5$, $t = 50$ and $t_{lq} = 1$. Please see Appendix A for results on the horse domain.

each latent $z_i \in S$ using a standard k -NN (k -nearest neighbor) algorithm; Second, we identify z_i as a high-quality latent if m_i is larger than a user-defined threshold t , and add z_i into a set S_{hq} . The set S_{hq} is the output of our algorithm, which contains all the high-quality latents identified. Algorithm 1 shows the pseudocode of our algorithm. Note that our algorithm is fundamental and widely applicable to different types of GANs as long as they sample latents from a standard normal distribution, e.g. conditional GANs (Brock et al., 2019).

Relationship to Truncation Trick. To our knowledge, the truncation trick (Marchesi, 2017; Brock et al., 2019; Karras et al., 2019) is the only a priori method to sample high-quality GAN latents before our work, which is based on a heuristic that high-quality latents are those close to their mean. However, such a heuristic is purely empirical with few insights. Surprisingly, the proposed *hubness priors* have revealed the rationale of the truncation trick: the *hub* latents obtained by our method tend to cluster towards their mean (Radovanovic et al., 2010). Thus, we argue that the well-known truncation trick is a naive approximation of our method as it only captures near-mean *hub* latents but overlooks those that are relatively far from the mean. Please see Section 4.5 for an empirical justification of our claims.

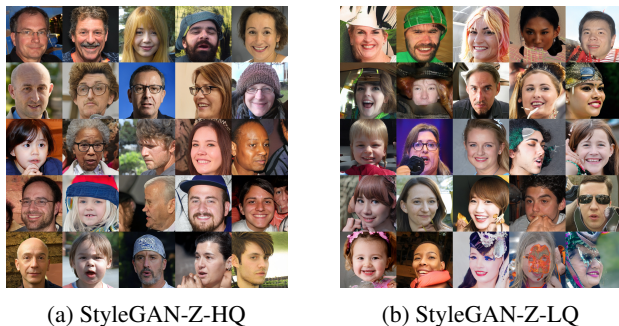


Figure 6. Performance of our method on StyleGAN2’s Z -space. We use $n = 10000$, $k = 5$, $t = 50$ and $t_{lq} = 1$.

4. Experimental Results

4.1. Experimental Setup

Due to its a priori nature, our method allows for the sampling of high-quality GAN latents before the synthesis of images. Thus, for the sampling of StyleGAN’s Z -space and other GANs’ latents, we use an Intel(R) Core(TM) i7-10875H CPU; for the sampling of StyleGAN’s W -space latents, we use a GeForce RTX 2080 Ti GPU as the computation involves passing Z -space latents through a fully-connected mapping network (Karras et al., 2019; 2020; 2021). For the synthesis of high-quality images, we use publicly-released Github codes of StyleGANs³ (Karras et al., 2019; 2020; 2021), BigGAN⁴ (Brock et al., 2019), ProGAN⁵ (Karras et al., 2018) with a GeForce RTX 2080 Ti GPU. Unless specified, all results are generated with the W -space of StyleGAN2 (Karras et al., 2020). All quantitative results are averaged over three runs. **Note that JPEG is applied to compress the synthesized images to meet the size limit. Please see the supplementary materials for uncompressed images.**

4.2. Effectiveness of Hubs Priors

As Figs. 3 (a) and (b) show, we compare the images generated by StyleGAN2 (Karras et al., 2020) using our method with those generated using the baseline, *i.e.* random latent sampling. It can be observed that our method consistently yields high-quality images while the baseline generates both high-quality and low-quality images, which demonstrates the effectiveness of the proposed *hubness priors*. Quantitatively, we observed better FID scores of images generated using our method than those by the baseline (Table 2).

Low-quality Latents. As Conjecture 3.2 implies, the proposed *hubness priors* can also be used to identify low-quality latents that yield unrealistic synthesized images. Thus, as

³StyleGAN2,3: <https://github.com/NVlabs/stylegan2>, <https://github.com/NVlabs/stylegan3>.

⁴<https://github.com/ajbrock/BigGAN-PyTorch>

⁵https://github.com/tkarras/progressive_growing_of_gans

a complement to high-quality latent sampling, we implement low-quality GAN latent sampling by reversing the thresholding scheme in Algorithm 1 to $m_i \leq t_{lq}$ and have Algorithm 2 (Appendix B). As Fig. 3 (c) shows, almost all synthesized images are of low quality, which justifies the effectiveness of the proposed *hubness priors*.

In fact, our *hubness priors* can be used to sort all sampled latents into a *hubness spectrum* according to their hub values m (Appendix C), where the quality of images changes from high to low from left to right with decreasing m .

4.3. Versatility

To demonstrate the versatility of our method, we show that it generalizes across different GAN architectures, different image domains and different latent spaces of the StyleGAN series (Karras et al., 2019; 2020; 2021).

Different GAN Architectures. As Fig. 4 shows, to justify that our method works across different GAN architectures, we show that our method also works on three other state-of-the-art GAN architectures, *i.e.* ProGAN (Karras et al., 2018), BigGAN (Brock et al., 2019), and the recent StyleGAN3 (Karras et al., 2021).

Different Image Domains. As Fig. 5 shows, to justify that our method works across different image domains, we show that our method also works on StyleGAN2 models pretrained on other images domains⁶: cars, cats and horses.

StyleGAN’s Z -space. As Fig. 6 shows, our method also works for the Z -space of StyleGAN2 (Karras et al., 2020). However, we observed that the quality variance of synthesized images is slightly lower when using the W -space. Thus, we propose to use the W -space for StyleGAN2.

4.4. Justification of Algorithmic Choices

Threshold t . In our method, given a fixed latent sample set S , the threshold t determines the trade-off⁷ between image quality and number of output latents: the larger t , the higher image quality, but the fewer output images. However, as Fig. 7 and Table 1 show, we observed that the image quality remains high for various choices of t . Since the image quality is not sensitive to the choice of t in a relatively large range, we suggest using $t = 50$ as the default value for the case when $n = 10000$, $k = 5$. Note that we can easily extend our algorithm to output a user-specified number of images (denoted as n') by using a revised scheme: if there

⁶All pre-trained networks are available at: <https://github.com/NVlabs/stylegan2>.

⁷Note that this trade-off only applies to a fixed S . Our method can generate an infinite number of high-quality samples by simply using multiple latent sets S_1, S_2, \dots, S_N or a larger S .



Figure 7. Performance of our method with different choices of threshold $t = 60, 50, 40, 35$. We use $n = 10000, k = 5$.




Figure 8. Performance of our method with different choices of hyper-parameter $k = 3, 5, 7, 10$. We use $t = 50, n = 10000$.

are enough images in S , we first sort all images in the descending order of hub value m , and keep the top n' latents; otherwise, we successively draw more latent sets S_i and keep all $m > t$ images from them until we get n' images.

Hyper-parameter k . We tested the performance of our algorithm with various choices of $k = 1, 3, 5, 7, 10$ in the k -NN algorithm. Apart from the case when no hub latents can be found ($k = 1$), we show the results of $k = 3, 5, 7, 10$ in Fig. 8 and Table 1. It can be observed that the image quality is not sensitive to the choice of k . Nevertheless, we noticed that using a larger k yields more output hub latents for a given latent set S and threshold t , but at the cost of slightly longer computation (Appendix E). To achieve a balance, we suggest using $k = 5$ as a default value when $n = 10000, t = 50$.

Size of Latent Sample Set n . As Table 1 shows, we also test the performance of the proposed method against various sizes $n = 10000, 20000, 30000, 40000$ of latent sample set S . Please see Appendix A for qualitative results. Similar to above, we observed that (i) although the FID scores get slightly better with increasing n , the image quality is not sensitive to the choice of n ; (ii) using a larger n yields more output hub latents but at the cost of longer computation (Appendix E). To achieve a balance, we suggest using $n = 10000$ as a default value when $k = 5, t = 50$.

Table 1. FID scores of StyleGAN2 images synthesized using our method with different choices of k, t and n , whose default values are $k = 5, t = 50$ and $n = 10000$. We sample 2,000 images to compute the FIDs, whose rationale is discussed in Sec. 4.5.

k	FID \downarrow	t	FID \downarrow	n	FID \downarrow
3	22.793	60	20.749	10000	22.782
5	22.782	50	22.782	20000	22.021
7	22.720	40	24.517	30000	21.679
10	22.560	35	25.412	40000	19.124

4.5. Relationship with Truncation Trick

The truncation trick (Marchesi, 2017; Brock et al., 2019; Karras et al., 2019) has been widely used in state-of-the-art GANs. Specifically, it truncates randomly sampled latents \mathbf{w} to $\mathbf{w}' = \bar{\mathbf{w}} + \psi(\mathbf{w} - \bar{\mathbf{w}})$ to obtain high-quality latents that yield high-quality synthesized images, where $\bar{\mathbf{w}}$ is the mean of a large number of randomly sampled latents, ψ is a scaling parameter. As discussed in Section 3.2, we argue that it is a naive approximation of our method.

Distance to the Means of Hub and All Latents. To justify our claim, we first investigate the distances of our *hub latents* to their mean and their distances to the mean of all sampled latents. As Fig. 9 shows, it can be observed that: i) Our hub latents are closer to both the hub mean and the all latent mean than randomly sampled latents, which justifies the ‘‘central clustering effect’’ of our hub

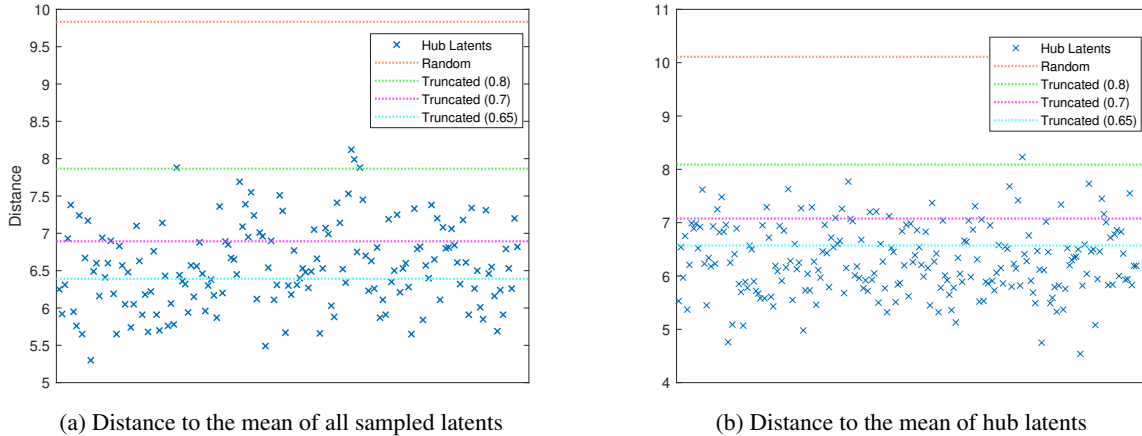


Figure 9. The distances of our hub latents to (a) the mean of all sampled latents and (b) the mean of hub latents. Random: the average distance of randomly sampled latents; Truncated (ψ_0): the average distance of latents after truncation trick ($\psi = \psi_0$).

Table 2. Comparison of FID scores of StyleGAN2 synthesized images using our method and the truncation trick. FFHQ-1 and FFHQ-2: real images sampled from the FFHQ dataset (Karras et al., 2019); Hubs (50): our method with $t = 50$; Truncated (0.7): truncation trick with $\psi = 0.7$; Random: random sampling. We sample 2000 latents/images for all methods compared. The FID scores between i) FFHQ-1 and FFHQ-2; and ii) Random and FFHQ-1, FFHQ-2 are used as baselines. Dist2Mean: distances of sampled latents to the all latent mean.

Methods	FID↓		Dist2Mean
	FFHQ-1	FFHQ-2	
FFHQ-2	16.505	—	—
Hubs (50)	21.955	23.609	6.247
Truncated (0.7)	25.097	25.127	6.893
Random	35.455	35.598	9.847

latents (Radovanovic et al., 2010). ii) Surprisingly, the distances of most hub latents are around 6.0 to 7.0 for both cases, which is roughly the same as the distances of randomly sampled latents truncated with a parameter $\psi = 0.7$, *i.e.* the StyleGAN-recommended (Karras et al., 2019) parameter value for the truncation trick. However, StyleGAN obtained the value $\psi = 0.7$ empirically via try-and-error while we obtain it as a byproduct of our method, which justifies the superiority and fundamentality of our approach. iii) A small portion of our hub latents are of larger distances (*e.g.* around 7.5 and 8.0) to the means, which will be overlooked by the truncation trick with $\psi = 0.7$. In addition, applying the truncation trick with $\psi = 0.8$ are prone to get low-quality latents that yield low-quality images while our “distant” hub latents are still of high quality (Fig. 11). This further justifies the superiority of our method against the truncation trick.

FID Scores. As Table 2 shows, we also justify the superiority of our method by comparing the FID scores (Heusel

et al., 2017) of images generated by StyleGAN2 using both the truncation trick, $\psi = 0.7$ (Karras et al., 2019) and our method. Specifically, we compute the FID scores between images generated by i) real images sampled from the FFHQ datasets, *i.e.* FFHQ-1 and FFHQ-2 in Table 2; ii) our hub latents and FFHQ-1, FFHQ-2; iii) truncated latents ($\psi = 0.7$) and FFHQ-1, FFHQ-2; iv) randomly sampled latents and FFHQ-1, FFHQ-2. It can be observed that i) both our method and the truncation trick outperform random sampling; ii) our method achieves better FID scores than the truncation trick. Note that we intentionally used a small number of images (*i.e.* 2,000) to compute FID to avoid covering the entire distribution and thus suffer less from the restriction of latent spaces. In comparison with the results in (Jung & Keuper, 2021) and the bias-free FID_∞ (Chong & Forsyth, 2020) computed with 10K images (Table 3), our FID scores of “Truncated (0.7)” images are better than “Random”, which is consistent with human perception. Note that our method outperforms Truncated (0.7) in both cases. Examples of StyleGAN2 synthesized images after the truncation trick ($\psi = 0.7$) are shown in Appendix D. Nevertheless, even using a small number of images, FID may still not be a good evaluation metric for our task. Therefore, we resort to the precision and recall metrics (Kynkäänniemi et al., 2019) that make more sense.

Table 3. FID_∞ scores (Chong & Forsyth, 2020) computed with 10K images, which are ineffective as they capture the entire distribution and thus suffer from the restriction of latent spaces. **Red**: random sampling has the best score, which contradicts human perception as the images sampled with it are of the lowest quality.

Method	Hubs (50)	Truncated (0.7)	Random
$FID_\infty \downarrow$	15.398	15.761	2.923

Precision and Recall (Kynkäänniemi et al., 2019). As Table 4 shows, our method achieves a high precision com-




Figure 10. The class distributions of the StyleGAN2 model pretrained on the CIFAR10 dataset with (a) our hubness-based sampling ($t = 50$), (b) the truncation trick ($\psi = 0.7$), and (c) the random sampling methods. WD: Wasserstein distance.




Figure 11. StyleGAN2 images synthesized from (a) distant hub latents far from their mean; (b) truncated latents ($\psi = 0.8$).

parable to Truncated (0.3) which sacrifices the synthesis diversity (*i.e.* low recall) while retaining a very high recall comparable to Random which includes many low-quality results (*i.e.* low precision). This further justifies the superiority of our method.

Table 4. Comparison of precision and recall (Kynkäänniemi et al., 2019) of StyleGAN2 synthesized images using our method and the truncation trick.

Method	Precision \uparrow	Recall \uparrow
Hubs (50)	<u>0.890</u>	<u>0.324</u>
Truncated (0.3)	0.892	0.015
Truncated (0.7)	0.811	0.223
Random	0.720	0.393

4.6. Impact on Class Balance

We further investigate how our method affects the class balance of unconditional GANs pre-trained on multi-class datasets. As Fig. 10 shows, we evaluate the class balance of a StyleGAN2 model pretrained on the CIFAR10 dataset with i) random sampling⁸ (*i.e.* Random), ii) truncation trick ($\psi = 0.7$) and iii) our hubness-based sampling method. Specifically, we sample 50,000 images each and use a pre-trained CIFAR10 classifier⁹ to estimate their class distributions. Note that although a “larger” difference can be

observed visually, our method actually preserves the class balance better as it has a smaller Wasserstein distance to the distribution of Random than the truncation trick. In addition, as Table 5 shows, our method achieves a better Inception Score (Salimans et al., 2016) that favours balanced and high-confidence classifications, which further justifies the superiority of our method in preserving class balance.

Table 5. Evaluation of class balance with Inception Scores (IS) (Salimans et al., 2016) of StyleGAN2 pretrained on the CIFAR10 dataset using our hubness-based sampling ($t = 50$), the truncation trick ($\psi = 0.7$), and the random sampling methods.

Method	Hubs (50)	Truncated (0.7)	Random
IS	6.212	6.059	7.080

5. Conclusions

In this paper, we address the quality variance of GAN synthesized images by investigating the sampling of GAN latents. Specifically, we first show that GAN latents are not uniformly distributed in the latent space due to the *hubness* phenomenon of data distributions in high dimensional space. In addition, there exist *hub* latents that are much more likely to be nearest neighbors of others and contribute more to the synthesis of high-quality images. Then, we formulate the above as the *hubness priors* and propose a novel GAN latent sampling algorithm, which allows for efficient and high-quality image synthesis for GANs. Furthermore, we show that the well-known truncation trick is a naive approximation of our method that utilizes the “central clustering effect” of *hub* latents, which not only uncovers the rationale of the truncation trick, but also indicates that our method is superior and more fundamental.

Acknowledgements

We appreciate the reviewers’ constructive comments in improving the paper. This research was partially funded by the UK Engineering and Physical Sciences Research Council (EPSRC) through the Doctoral Training Partnerships (DTP) with No. EP/T517951/1 (2599521).

⁸<https://github.com/POSTECH-CVLab/PyTorch-StudioGAN>

⁹<https://github.com/open-mmlab/mmlclassification>, ResNet50

References

- Abdal, R., Qin, Y., and Wonka, P. Image2StyleGAN: How to embed images into the styleGAN latent space? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4432–4441, 2019.
- Abdal, R., Qin, Y., and Wonka, P. Image2StyleGAN++: How to edit the embedded images? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8296–8305, 2020.
- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *International conference on machine learning*, pp. 214–223. PMLR, 2017.
- Brock, A., Donahue, J., and Simonyan, K. Large scale GAN training for high fidelity natural image synthesis. In *International Conference on Learning Representations*, 2019.
- Buza, K. Classification of gene expression data: a hubness-aware semi-supervised approach. *Computer methods and programs in biomedicine*, 127:105–113, 2016a.
- Buza, K. Semi-supervised naive hubness Bayesian k-nearest neighbor for gene expression data. In *Proceedings of the 9th International Conference on Computer Recognition Systems CORES 2015*, pp. 101–110. Springer, 2016b.
- Buza, K. A. and Koller, J. Classification of electroencephalograph data: A hubness-aware approach. *Acta Polytechnica Hungarica*, 13(2):27–46, 2016.
- Chong, M. J. and Forsyth, D. Effectively unbiased FID and inception score and where to find them. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6070–6079, 2020.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of Wasserstein GANs. *Advances in neural information processing systems*, 30, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- 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. *Advances in neural information processing systems*, 30, 2017.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Jiang, Y., Chang, S., and Wang, Z. TransGAN: Two pure transformers can make one strong GAN, and that can scale up. *Advances in Neural Information Processing Systems*, 34, 2021.
- Jung, S. and Keuper, M. Internalized biases in fréchet inception distance. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018.
- Karras, T., Laine, S., and Aila, T. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4401–4410, 2019.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8110–8119, 2020.
- Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., and Aila, T. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., and Aila, T. Improved precision and recall metric for assessing generative models. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mao, X., Li, Q., Xie, H., Lau, R. Y., Wang, Z., and Paul Smolley, S. Least squares generative adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2794–2802, 2017.
- Marchesi, M. Megapixel size image creation using generative adversarial networks. *arXiv preprint arXiv:1706.00082*, 2017.
- Menon, S., Damian, A., Hu, S., Ravi, N., and Rudin, C. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2437–2445, 2020.
- Mescheder, L., Geiger, A., and Nowozin, S. Which training methods for GANs do actually converge? In *International conference on machine learning*, pp. 3481–3490. PMLR, 2018.

- Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks. In *International Conference on Learning Representations*, 2018.
- Newman, C. M. and Rinott, Y. Nearest neighbors and Voronoi volumes in high-dimensional point processes with various distance functions. *Advances in Applied Probability*, 17(4):794–809, 1985.
- Newman, C. M., Rinott, Y., and Tversky, A. Nearest neighbors and Voronoi regions in certain point processes. *Advances in Applied Probability*, 15(4):726–751, 1983.
- Park, T., Liu, M.-Y., Wang, T.-C., and Zhu, J.-Y. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2337–2346, 2019.
- Qin, Y., Mitra, N., and Wonka, P. How does Lipschitz regularization influence GAN training? In *European Conference on Computer Vision*, pp. 310–326. Springer, 2020.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Radovanovic, M., Nanopoulos, A., and Ivanovic, M. Hubs in space: Popular nearest neighbors in high-dimensional data. *Journal of Machine Learning Research*, 11(sept): 2487–2531, 2010.
- Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., and Cohen-Or, D. Encoding in style: a StyleGAN encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2287–2296, 2021.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training GANs. *Advances in neural information processing systems*, 29, 2016.
- Tomasev, N. and Mladenic, D. Nearest neighbor voting in high-dimensional data: Learning from past occurrences. In *2011 IEEE 11th International Conference on Data Mining Workshops*, pp. 1215–1218. IEEE, 2011.
- Tomašev, N. and Mladenić, D. Hub co-occurrence modeling for robust high-dimensional kNN classification. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 643–659. Springer, 2013.
- Tomašev, N., Radovanovic, M., Mladenic, D., and Ivanovic, M. A probabilistic approach to nearest-neighbor classification: Naive hubness Bayesian kNN. In *Proc. 20th ACM Int. Conf. on Information and Knowledge Management (CIKM)*, pp. 2173–2176, 2011.
- Tomašev, N., Radovanović, M., Mladenić, D., and Ivanović, M. Hubness-based fuzzy measures for high-dimensional k-nearest neighbor classification. *International Journal of Machine Learning and Cybernetics*, 5(3):445–458, 2014.
- Tomašev, N., Buza, K., Marussy, K., and Kis, P. B. Hubness-aware classification, instance selection and feature construction: Survey and extensions to time-series. In *Feature selection for data and pattern recognition*, pp. 231–262. Springer, 2015.
- Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., and Cohen-Or, D. Designing an encoder for StyleGAN image manipulation. *ACM Transactions on Graphics (TOG)*, 40(4):1–14, 2021.
- Zhu, J.-Y., Park, T., Isola, P., and Efros, A. A. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pp. 2223–2232, 2017.
- Zhu, P., Abdal, R., Qin, Y., and Wonka, P. SEAN: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5104–5113, 2020.

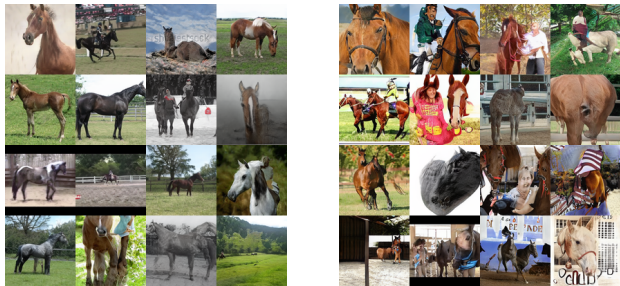
A. Additional Experimental Results

Different Image Domains. As Fig. 12 shows, as complement to Fig. 5 in the main paper, to justify that our method works across different image domains, we show that our method also works on a StyleGAN2 model pretrained on the horse domain.

Distributions of Hub Latents. As a complement to Fig. 2, Fig. 13 shows additional results on the distributions of m -hub latents in the latent spaces of state-of-the-art GANs.

Choice of size of Latent Sample Set n . Fig. 14 shows the qualitative results of our method with $n = 20000, 30000, 40000$ (see Fig. 3 for results when $n = 10000$). It can be observed that the visual quality remains similar across different n , which indicates that the image quality is not sensitive to the choice of n .

Quantitative results with BigGAN (Brock et al., 2019). As Table 6 shows, our method outperforms Truncated (0.7) with the BigGAN architecture pretrained on the 1000-class ImageNet ILSVRC 2012 dataset on precision and



(a) StyleGAN-Horse-HQ

(b) StyleGAN-Horse-LQ

Figure 12. Performance of our method on StyleGAN2 pretrained on the horse domain. It can be observed that our method works well. (a) and (b) are images synthesized from high-quality (HQ) and low-quality (LQ) latents obtained by Algorithm 1 and Algorithm 2 (Appendix B), respectively. We use $n = 10000$, $k = 5$, $t = 50$ and $t_{lq} = 1$.

recall (Kynkäänniemi et al., 2019), which further justifies the superiority of our method.

Table 6. Quantitative results with BigGAN (ImageNet).

Method	Precision \uparrow	Recall \uparrow
Hubs (50)	0.147	0.311
Truncated (0.7)	0.131	0.264

B. Pseudocode of Low-quality GAN Latent Sampling

The pseudocode of our low-quality GAN latent sampling algorithm is shown in Algorithm 2. It is a simple inverse of Algorithm 1, using a different thresholding scheme.

C. Hubness Spectrum

Fig. 15 shows the *hubness spectrum* obtained by our hubness priors. It can be observed that the quality of images changes from high to low from left to right with decreasing m .

D. Example Images with Truncation Trick

We show examples of StyleGAN2 synthesized images after the truncation trick ($\psi = 0.7$) in Fig. 16.

E. Running Time

Table 7 shows the running time of our method with different choices of k and n . It can be observed that the running time increases mildly with k but significantly with n .

F. Limitation and Future Work

Although our method allows for the sampling of high-quality latents, the quality of synthesized images is bounded by the performance of the pre-trained GANs used to synthesize

Algorithm 2 Low-quality GAN Latent Sampling with Hubness Priors

Input: a set of GAN latents $\mathcal{S} = \{z_1, z_2, \dots, z_n\}$ sampled from a standard normal distribution $\mathcal{N}(0, I)$, a hyperparameter k , a threshold t_{lq}

Output: S_{lq}

Step 1

$m_{1,2,\dots,n} \leftarrow 0$

for $i \leftarrow 1$ to n **do**

$\{\text{id}x_1, \text{id}x_2, \dots, \text{id}x_k\} \leftarrow k\text{-NN}(z_i)$

for $j \leftarrow 1$ to k **do**

$m_{\text{id}x_j} \leftarrow m_{\text{id}x_j} + 1$

end for

end for

Step 2

$S_{lq} \leftarrow \emptyset$

for $i \leftarrow 1$ to n **do**

if $m_i \leq t_{lq}$ **then**

$S_{lq} \leftarrow S_{lq} \cup z_i$

end if

end for

Table 7. Running time of our method using the StyleGAN2 W -space with different choices of k and n . The default parameter values are $k = 5$, $t = 50$ and $n = 10000$.

k	Time(s)	n	Time(s)
3	163s	10000	167s
5	167s	20000	647s
7	176s	30000	1272s
10	185s	40000	2554s

them. Also, we observed that the proposed *hubness priors* may overlook some relatively high-quality images with small *hub* values m (Fig. 17). We conjecture that the reason might be that the limited sizes of latent sample sets (e.g. $n = 10000, 20000, \dots$) cannot capture all *hub* latents. This is partially verified by our experiment on the choice of n . However, it is difficult to test larger n due to the $O(n^2)$ time complexity to compute the hub values m for all points in a latent sample set. We hope to investigate this issue in future work. We also hope to apply our insights on the hubness phenomenon in GAN latent space to improve the training of GANs and make GANs unbiased for all latents. The acceleration of our algorithm is also a very interesting direction for future work.




Figure 13. Distributions of m -hub latents for state-of-the-art GANs, $k = 5, 7, 10$ (the k -NN algorithm) and $n = 10000$ (size of latent sample set S). *: Although both are 512-dimensional, the ProGAN (Karras et al., 2018) latents are sampled directly from $\mathcal{N}(0, I)$ while the StyleGAN latents further normalized the sampled latents to be of the same norm (Karras et al., 2019). All distributions are highly tailed to the right, which demonstrates the existence of hubness phenomenon (Radovanovic et al., 2010) in GAN latent spaces.

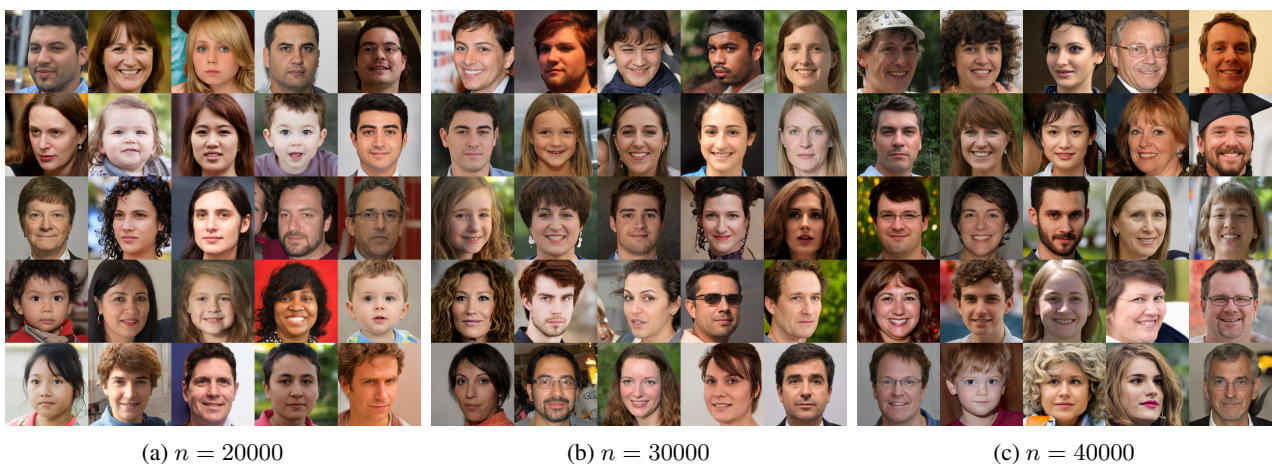


Figure 14. Performance of our method with different sizes $n = 20000, 30000, 40000$ of sample set S . We use $k = 5, t = 50$.

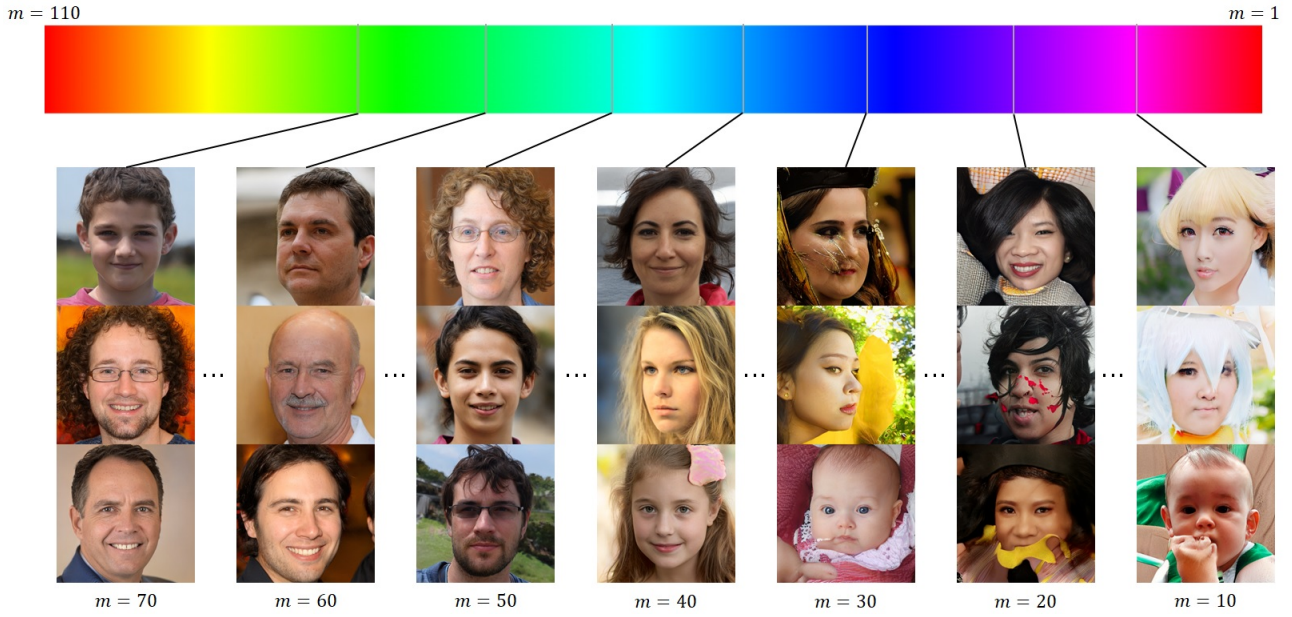


Figure 15. The *hubness spectrum* of StyleGAN2 (Karras et al., 2020) synthesized images ranked according to their *hub* values m . We use $n = 10000$, $k = 5$. Note that the spectrum is highly tailed to the left and thus there are few images in the range $m = (70, 110)$.



Figure 16. Examples of StyleGAN2 synthesized images after the truncation trick ($\psi = 0.7$).



Figure 17. Relatively high-quality StyleGAN2 (Karras et al., 2020) synthesized images with small hub values m . However, there are still small artifacts in these images (*e.g.* background and facial details).