

Demystifying Randomly Initialized Networks for Evaluating Generative Models

Junhyuk Lee, Jun-Hyuk Kim, Jong-Seok Lee

Yonsei University, South Korea
 {junhyuklee, junhyuk.kim, jong-seok.lee}@yonsei.ac.kr

Abstract

Evaluation of generative models is mostly based on the comparison between the estimated distribution and the ground truth distribution in a certain feature space. To embed samples into informative features, previous works often use convolutional neural networks optimized for classification, which is criticized by recent studies. Therefore, various feature spaces have been explored to discover alternatives. Among them, a surprising approach is to use a randomly initialized neural network for feature embedding. However, the fundamental basis to employ the random features has not been sufficiently justified. In this paper, we rigorously investigate the feature space of models with random weights in comparison to that of trained models. Furthermore, we provide an empirical evidence to choose networks for random features to obtain consistent and reliable results. Our results indicate that the features from random networks can evaluate generative models well similarly to those from trained networks, and furthermore, the two types of features can be used together in a complementary way.

Introduction

In recent years, generative models have flourished in deep learning literature, which can generate realistic and high-quality image samples. The main objectives of generative models are to estimate the distribution of the training data and to generate samples that are likely to belong to the distribution. In the image domain, generated images are highly realistic thanks to state-of-the-art deep generative models including variational autoencoders (VAEs) (Hazami, Mama, and Thuraiatnam 2022), generative adversarial networks (GANs) (Brock, Donahue, and Simonyan 2019), and diffusion models (Dhariwal and Nichol 2021).

Evaluation of generative models is one of the most crucial parts to develop the models. A classical approach is to measure the likelihood of the real data under the estimated distribution. However, for intractable models such as GANs, the likelihood cannot be directly computed and estimation of the likelihood is impractical when generated samples are high-dimensional data. Moreover, Theis, van den Oord, and Bethge (2016) argued that the likelihood is often unrelated to the quality of generated samples.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

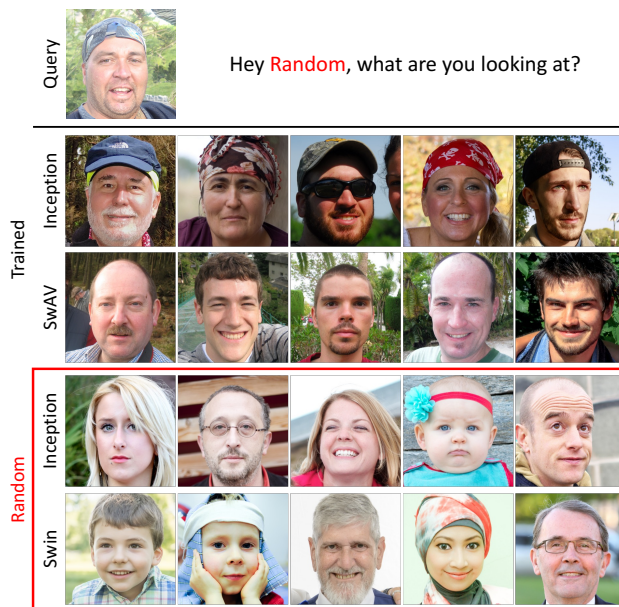


Figure 1: Five nearest neighbors among the ground truth real images by each model for a generated query image. We argue that the random features have different perspectives from the trained ones.

Various evaluation metrics have been proposed so far other than the likelihood measure. In general, they are based on the comparison between the learned distribution and the ground truth distribution in certain feature spaces. Most early metrics use classification models to embed images into informative feature spaces (Salimans et al. 2016; Heusel et al. 2017; Sajjadi et al. 2018; Bińkowski et al. 2018). For instance, the Fréchet Inception distance (FID) (Heusel et al. 2017), which is the most ubiquitously used metric to this day, uses the Inception model (Szegedy et al. 2016) trained on the ImageNet dataset (Deng et al. 2009). These metrics take advantage of the classification models that are regarded as effective to capture high-level semantic information.

However, recent studies have argued that features optimized for classification may be limited for evaluating generative models. In (Naem et al. 2020), it was shown that the evaluation results can be underrated in domains differ-

ent from that of natural images. It was also observed that the Inception feature yields unreliable evaluation results when it comes to non-ImageNet datasets (Morozov, Voynov, and Babenko 2021). Kynkäänniemi et al. (2022) showed that the Inception feature has the null space where the FID can be improved without enhancing the quality of generated images, which is due to the bias of the classification task.

One approach to address these issues is to use the models that are more sophisticated than the classification models. Morozov, Voynov, and Babenko (2021) proposed to use features from self-supervised learning-based models such as SwAV (Caron et al. 2020) to deal with the images that do not belong to the ImageNet classes such as face, bedroom, and church. Kynkäänniemi et al. (2022) showed that the evaluation results agree with human assessment when the CLIP model (Radford et al. 2021) trained on the vision-language modeling task is employed.

Curiously, randomly initialized networks have been considered as another approach. In (Naeem et al. 2020), features from convolutional neural networks (CNNs) with random weights yielded sensible evaluation results for the generative models in domains other than natural images such as hand-written digits and speech data. Sauer, Schwarz, and Geiger (2022) also employed a randomly initialized CNN to evaluate generative models. Although features from randomly initialized networks (referred to as “random features” in this paper) are occasionally considered in the literature, we notice that justification of using the random features has been insufficiently addressed especially for the domains where the sophisticated features already exist (e.g., natural images in Figure 1). In particular, the following questions remain unanswered:

1. Do random features even capture any useful information?
2. What are the advantages of random features than those from trained networks?
3. Under what circumstances do random features work?

In this paper, we demystify randomly initialized networks for evaluating generative models via extensive analysis. To answer the above research questions, first, we examine the feature spaces established by randomly initialized networks by inspecting both sample-based and distribution-based distances. Then, we perform in-depth analysis of faulty images created by generative models. We also demonstrate that random features have distinct perspectives compared to the features from trained networks (referred to as “trained features”). Finally, we present comprehensive evaluation results of various generative models including GANs, VAEs, and diffusion models using random features.

Our main contributions are summarized as follows:

- We provide various analyses on the feature spaces of randomly initialized networks in the viewpoint of evaluating generative models. Based on the analyses, we establish a fundamental basis of the random features to be employed to evaluate generative models.
- We show the cases where the random features have different characteristics compared to the trained ones. We examine the faulty images from generative models to understand how the images appear in the feature spaces. It

is shown that random features have different criteria to judge the faulty images from the trained features.

- We rigorously evaluate generative models with features from various networks with random weights. As a result, we demonstrate that the random features are useful for evaluation of generative models and, furthermore, they can provide information complementary to that by the trained features regarding the characteristics of the generated images.

Related Work

Generative Models

Generative models aim to estimate the distribution of the target training data. Recently, the state-of-the-art models can generate highly realistic images after they are trained on challenging real-world image datasets (Deng et al. 2009; Karras, Laine, and Aila 2019). As one of the likelihood-based generative models, VAEs estimate the distribution by mapping images to the learned latent space (Hazami, Mama, and Thuraiatnam 2022), but often generate blurry images (Huang et al. 2018). GANs optimize the adversarial loss, which play a min-max game between the generator and discriminator (Goodfellow et al. 2014). GANs significantly improve the visual quality of the generated images, while the estimated densities are intractable. Diffusion models learn to revert the progressive diffusion process that adds noise to an image. Recently, they yield comparable results to GANs with tractable estimated densities (Dhariwal and Nichol 2021).

Evaluation of Generative Models

For likelihood-based generative models, the likelihood measure can be directly adopted for evaluation. However, it is not the case for GANs since they are based on likelihood-free training. For GANs, kernel density estimation using Parzen window can be employed as an indirect likelihood measure, but a significantly large number of samples are required for high-resolution images (Theis, van den Oord, and Bethge 2016). It was also argued that the likelihood measure is sometimes irrelevant to the quality of samples.

To properly consider the sample quality in an efficient manner, the Inception score was proposed (Salimans et al. 2016), which utilizes the trained classification model, i.e., Inception (Szegedy et al. 2016). However, since the target distribution is not considered, the evaluation results often fail to detect mode collapse and favor highly classifiable images rather than high quality images.

FID is the Fréchet distance (also known as the 2-Wasserstein distance and the earth mover’s distance) between the trained and target distributions in the feature space of the Inception model. By assuming that the features follow normal distributions, it is computed as follows,

$$\text{FID} = \|\mu^g - \mu^r\|_2^2 + \text{Trace} \left(\Sigma^g + \Sigma^r - 2(\Sigma^g \Sigma^r)^{\frac{1}{2}} \right), \quad (1)$$

where (μ^g, Σ^g) and (μ^r, Σ^r) are the sample mean and the covariance of the generated and real data, respectively, in

the Inception feature space. Thanks to its simplicity and reasonable accuracy, FID is ubiquitously used for evaluation of various generative models. However, limitations of FID have been also pointed out: FID is based on an over-simplified normality assumption (Bińkowski et al. 2018; Lee and Lee 2021), has a bias to the number of samples (Chong and Forsyth 2020), and is unexpectedly sensitive to subtle pre-processing methods (Parmar, Zhang, and Zhu 2022).

Another approach is to separately consider two different aspects of performance (i.e., fidelity and diversity) rather than a single value such as FID (Sajjadi et al. 2018; Kynkäänniemi et al. 2019; Naeem et al. 2020). Such twofold metrics are particularly useful for diagnostic purposes. There also exist attempts to evaluate the quality of individual samples (Kynkäänniemi et al. 2019; Han et al. 2022). They are usually based on the distance of an image to its nearest neighbor among a set of real or generated images.

Most of the aforementioned metrics take advantage of classification models (e.g., the Inception model) to obtain meaningful features. However, it has been argued that the feature spaces are classification-oriented and may not be necessarily perceptual (Morozov, Voynov, and Babenko 2021; Kynkäänniemi et al. 2022). For alternative feature spaces, models trained on sophisticated tasks such as self-supervised learning and vision-language modeling are employed. It has been shown that features from the SwAV (Caron et al. 2020) and CLIP (Radford et al. 2021) models are effective for non-ImageNet datasets such as the FFHQ dataset (Karras, Laine, and Aila 2019).

Randomly Initialized Networks

In order to explain the expressive power of deep neural networks, models with random weights are often considered. Ulyanov, Vedaldi, and Lempitsky (2018) showed that the structure of CNNs itself has an implicit prior on the image domain, thus models can perform image super-resolution, in-painting, and de-noising tasks without a large number of training images. The lottery ticket hypothesis (Frankle and Carbin 2019; Ramanujan et al. 2020) demonstrates that a randomly initialized network contains subnetworks that perform comparably to the trained networks. Furthermore, features from randomly initialized networks can be used to predict class labels of images (Amid et al. 2022).

Naeem et al. (2020) argued that features from the VGG models (Simonyan and Zisserman 2015) with random weights are effective for evaluating generated samples in domains other than natural images such as hand-written digits and speech. In addition, they observed that the random features yield similar evaluation results to the trained features for natural images. In (Sauer, Schwarz, and Geiger 2022), it was also observed that the evaluation results using the feature of the Inception model with random weights are similar to those using the trained features. However, we argue that clear grounds beyond these observations are required to rely on random features for evaluating generative models.

Experimental Setup

For the real-world image datasets, we use the ImageNet (Deng et al. 2009) and FFHQ (Karras, Laine, and

Aila 2019) datasets. We also use a non-natural cartoon image dataset that consists of Pokémon images. For generated images, we use various generative models including E-VDVAE (Hazami, Mama, and Thuraiatnam 2022), BigGAN (Brock, Donahue, and Simonyan 2019), ProjectedGAN (Sauer et al. 2021), StyleGAN (Karras, Laine, and Aila 2019), StyleGAN2 (Karras et al. 2020), StyleGAN-XL (Sauer, Schwarz, and Geiger 2022), and ADM (Dhariwal and Nichol 2021). For the diffusion models, we use the classification-guidance model (ADM-C) and the up-scaling model (ADM-U).

We use 50,000 generated images for evaluating each generative model. We use the official codes of the models with pre-trained weights. For the ImageNet class-conditional models (i.e., BigGAN, ADM-C, and ADM-U), we generate 50 images for each of the 1,000 classes. BigGAN and ADM-U generate images at a resolution of 256×256 pixels, while E-VDVAE and ADM-C generate images at a resolution of 64×64 pixels. For FFHQ, we use StyleGAN and StyleGAN2 at 1024×1024 pixels, and E-VDVAE and ProjectedGAN at 256×256 pixels. We use ProjectedGAN and StyleGAN-XL trained on Pokémon to generate images at a resolution of 256×256 pixels.

For trained features, the Inception model (Szegedy et al. 2016) and the SwAV model with the ResNet50 structure (Caron et al. 2020) are used. For random features, we employ the Inception and VGG (Simonyan and Zisserman 2015) models that have been used in the previous studies (Sauer, Schwarz, and Geiger 2022; Naeem et al. 2020). We also consider Transformer models: Swin-T (Liu et al. 2021) and ViT-T (Dosovitskiy et al. 2020) with random weights. The Kaiming uniform (He et al. 2015) is used to initialize the weights. The evaluation results using the random features are averaged over five runs using different seeds for the random state initialization. Variations across the five runs are provided in the Supplementary Material¹, where we do not observe significant difference between runs.

For evaluation metrics, we employ the FID using various trained and random models for feature embedding. We provide further results using different metrics including twofold metrics in the Supplementary Material, which are consistent to the results using FID.

Understanding Random Feature Spaces

In general, the evaluation methods of generative models are based on the comparison between the learned distribution and the actual distribution in certain feature spaces. Therefore, understanding the feature spaces is crucial to properly interpret the evaluation results. In this section, we examine the feature spaces of randomly initialized networks through comparison with trained feature spaces: 1) which information they capture, and 2) how well they differentiate disturbances.

Nearest Neighbor Retrieval

First, we examine the characteristics of the images that are closely located in a feature space. We choose a query image

¹<https://arxiv.org/abs/2208.09218>

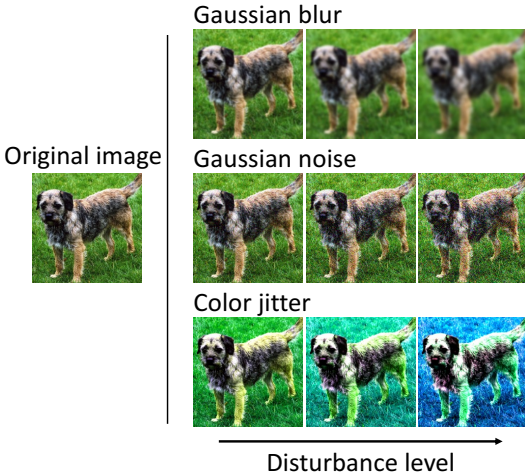


Figure 2: Example of disturbed images

generated from StyleGAN2 trained on FFHQ and retrieve its nearest real samples with respect to trained or random feature spaces. We use Inception and SwAV to extract trained features, and Inception and Swin with random weights to extract random ones.

Figure 1 shows a query image and its five nearest images with respect to each feature space. In the case of trained Inception, all retrieved samples contain adults wearing caps, which is one of the classes in the ImageNet dataset. SwAV retrieves images of adult males in front of green foliage as in the query image even if they do not wear caps. In the cases of both random features, the retrieved samples are overall in similar colors and brightness to the query image but highly diverse in terms of age, gender, background, etc. In summary, we observe that while the trained features mostly focus on high-level semantic information such as gender, age, backgrounds, and accessories, the random features mostly focus on low-level information such as color and brightness.

Disturbance

Next, we examine how well random features capture differences between distributions. We apply various low-level image disturbances such as blur, noise, and color jitter to a set of images. We also consider the disturbance in terms of high-level semantics, called class contamination, which replaces some images with those from different datasets. Then, we measure the difference between the original and disturbed sets using FID, but with randomly initialized Inception, VGG, Swin, and ViT as feature extractors. For comparison, the trained Inception and SwAV are also employed.

For the 50,000 images from the ImageNet validation dataset, three increasing levels of disturbances are applied (Figure 2): Gaussian blur with width $\sigma_{gb} \in \{1, 2, 3\}$, Gaussian noise with variance $\sigma_{gn}^2 \in \{0.05, 0.10, 0.15\}$, color jitter that changes all of the brightness, contrast, saturation, and hue with ratio $r_c \in \{0.1, 0.2, 0.3\}$, and class contamination with replacement ratio $r_r \in \{0.25, 0.5, 0.75\}$.

Feature	Level	Gaussian blur	Gaussian noise	Color jitter	Class contamination
Inception (trained)	1	7.81	7.95	17.29	19.64
	2	23.13	21.34	22.93	50.14
	3	40.67	41.09	28.16	89.63
SwAV (trained)	1	0.21	0.64	6.62	3.70
	2	1.45	1.92	7.02	9.50
	3	10.66	3.41	7.37	17.50
Inception (random)	1	29.88	0.10	469.89	2.22
	2	92.84	1.05	482.55	8.79
	3	148.37	4.00	496.70	20.86
VGG (random)	1	0.20	<0.01	4.14	0.03
	2	0.54	0.02	4.29	0.09
	3	0.84	0.07	4.46	0.20
Swin (random)	1	0.54	<0.01	436.09	3.66
	2	1.77	0.04	449.63	14.34
	3	3.05	0.16	459.22	32.77
ViT (random)	1	0.17	<0.01	220.25	2.32
	2	0.63	0.01	225.77	8.86
	3	1.20	0.04	229.07	19.90

Table 1: FID of disturbed images using features from trained and random models. The values for Inception (random) are scaled by $\times 10^{14}$.

Table 1 shows FID values between the original and disturbed images. Overall, all features correctly yield larger FID along increasing levels of disturbances. However, the trained and random features differently assess the distributions of the disturbed images with respect to the type of disturbances. For instance, both trained Inception and SwAV determine that the class contamination at level 3 causes the largest difference in the distribution. On the other hand, all random features give the largest score to the color jitter at level 3. These results are consistent with the observations in Figure 1, i.e., the trained features mostly focus on the high-level semantics such as class labels.

In terms of low-level disturbances, the trained Inception feature judges the degradations due to the noise, blur, and color jitter to be similar at level 2. SwAV also considers the blurred and noisy images to have similar scores at level 2. Considering that the noise is less perceptible for humans than the others (Figure 2), it can be argued that the trained Inception and SwAV features have excessive sensitivity to noise. On the contrary, all of the random features judge that the noise and color jitter cause the smallest and the largest difference, respectively, which is consistent with human perception for the low-level disturbances.

Therefore, the trained models and randomly initialized models have different criteria to judge how the distributions become different under disturbances. In most cases, the random features are more precise to the changes in low-level information such as color, blur, and noise. On the other hand, the trained features focus on high-level semantics such as class categories, and thus are often inconsistent with the human perception in terms of low-level information.

Disturbances of Generated Images

In the previous section, we discovered clues on how the different feature models regard differences in distributions using the disturbances applied to real images in a controlled manner. In this section, we investigate the disturbances that practically occur in generative models, in the viewpoint of how they appear in the trained or random feature spaces.

Can We Distinguish Disturbances in Generated Images?

In order to determine faulty images, we use a method similar to the outlier detection method (Naeem et al. 2020) with modifications in the criteria. We consider two feature spaces to separate different types of outliers. First, we use the pre-trained CLIP model that is designed to extract high-level semantic information for vision-language relationship modeling (Radford et al. 2021). We employ the visual encoder part of the CLIP model with a vision transformer structure. Second, we consider the lower stage of the CLIP model to develop a feature space that focuses on low-level information. Specifically, we use the output of the first convolutional layer of CLIP after global average pooling, which is also known as the convolutional stem.

To detect outliers, we compute the nearest neighbor distance d_m for a set of generated images, which is the distance between the feature of an image and that of its k^{th} nearest neighbor ($k = 5$), where the features are the outputs of a model $m \in \{\text{CLIP}, \text{STEM}\}$. We compute d_{CLIP} and d_{STEM} , which focus more on high-level and low-level information, respectively. Then, we obtain a set of outliers such that they have the largest d_m . In order to separate the high-level and low-level outliers, we consider the images that belong to only a single set of outliers and disregard outliers contained in both sets. More specifically, we take the images that have the top $\alpha\%$ largest d_{CLIP} and the bottom $(100 - \alpha)\%$ largest d_{STEM} for high-level outliers, and vice versa for low-level outliers. As a result, the outliers are far from the other samples, i.e., the corresponding feature space regards them as having distinct characteristics that are rare in the others.

Figure 3 shows the high-level and low-level outliers generated by StyleGAN2 trained on FFHQ ($\alpha=67\%$). In Figure 3a, the images have severe errors in high-level semantics such as unidentifiable faces, body parts, and accessories. On the other hand, in Figure 3b, most of the images show acceptable quality of faces but with vivid colors around them. Outliers for other datasets are presented in the Supplementary Material. The evaluation metrics would regard both outliers as deviated compared to the real FFHQ dataset. The high-level outliers are obviously failure cases for modeling the real distribution, whereas the low-level outliers are rare in the real distribution, which is also an abnormality in a different sense. Therefore, it would be desirable to correctly consider these outliers when evaluating generative models, which is examined below.



Figure 3: StyleGAN2 outliers identified using d_{CLIP} and d_{STEM}

How Do the Disturbances of Outliers Appear in Different Feature Spaces?

We analyze how a feature space regards changes in distributions by the different types of outliers. First, we build two subsets containing high- and low-level outlier images, respectively, from a set of images synthesized by a generative model trained on a real dataset. Then, we progressively replace ten images from the real dataset with the outliers from one of the subsets. For each step of replacement, we measure the FID between the original real dataset and the replaced set in various feature spaces. In this case, the distributions of the replaced sets can be regarded as distributions estimated by faulty generative models and the number of replaced images as the degree of the difference of the distributions.

In this experiment, we randomly sample 5,000 images from FFHQ as a real dataset. Among a set of 10,000 images generated by StyleGAN2, we compose two subsets consisting of 1,880 high-level and low-level outlier images, respectively. We only replace real images at each step, while keeping the previously added outliers. For feature spaces, we use two trained models (Inception and SwAV) and four randomly initialized models (Inception, VGG, Swin, and ViT). Figure 4 shows the FID with respect to the proportion of the outliers in different feature spaces.

In Figures 4a and 4b, replacing with high-level outliers yields larger FID than replacing with low-level outliers, indicating that both models judge the high-level outliers to incur larger difference. Still, both trained features are quite sensitive to the low-level outliers by assessing them to be comparable to the high-level outliers. On the contrary, all random features regard that the low-level outliers change the distribution much more than the high-level outliers (Figures 4c-4f). Note that the random features tend to fluctuate more along gradual changes in the distribution, which might be risky when distributions with subtle differences

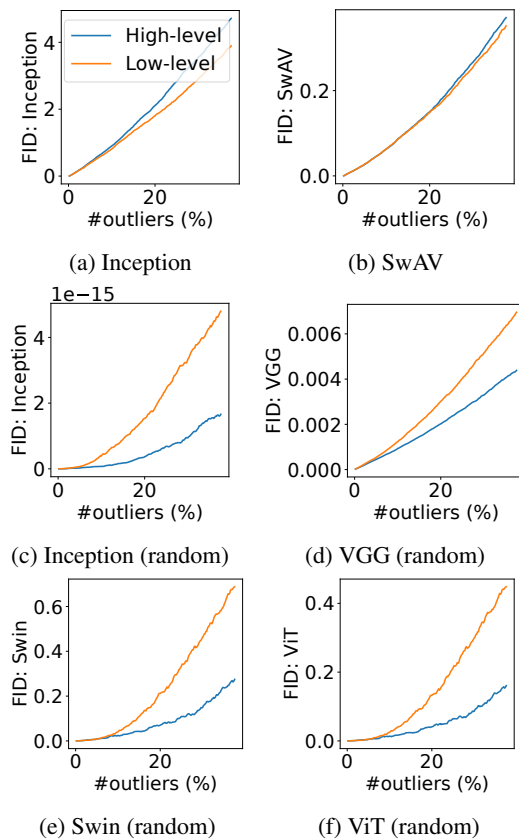


Figure 4: FID vs. the proportion of outliers for different features

are compared. Results of different datasets are shown in the Supplementary Material. These results indicate that sensitivity to different types of outliers is different depending on whether the feature is learned or random, which are consistent to those in Table 1, especially for class contamination and color jitter. The trained features are more sensitive to the high-level outliers than the low-level outliers, which also assess class contamination as severer degradation than color jitter. On the other hand, the random features are sensitive to the low-level outliers, which yield the largest FID for color jitter.

Do Random Features Correctly Evaluate Generative Models?

Comprehensive Evaluation

In this section, we conduct comprehensive evaluation of various generative models using random features. Table 2 shows the results using different features for FID. In most cases, in terms of determining the superiority of one against the other within a dataset, both trained and random features yield consistent results except for BigGAN and ADM-U. We observe that while BigGAN and ADM-U can generate decent samples, a considerable number of faulty samples with diverse disturbances are generated as shown in Figure 5. Thus, it would be necessary to utilize various feature spaces

Datasets	Trained		Random			
	Inception	SwAV	Inception	VGG	Swin	ViT
ImageNet 256²			$\times 10^{-15}$			
BigGAN	6.24	4.45	2.73	0.009	1.07	0.62
ADM-U	4.94	2.25	8.52	0.012	3.57	1.61
ImageNet 64²						
E-VDVAE	46.23	6.52	367.68	0.279	4.58	2.17
ADM-C	8.91	1.24	28.95	0.021	2.92	1.23
FFHQ 1024²						
StyleGAN	4.58	0.92	4.62	0.008	0.34	0.18
StyleGAN2	3.10	0.42	0.49	0.005	0.23	0.13
FFHQ 256²						
E-VDVAE	72.10	3.42	321.11	0.447	7.87	4.87
ProjectedGAN	4.18	1.32	18.01	0.028	0.68	0.35
Pokémon 256²						
ProjectedGAN	27.88	4.80	90.40	0.343	3.86	2.79
StyleGAN-XL	25.46	3.43	31.22	0.271	1.12	0.67

Table 2: Evaluation results in terms of FID using different feature spaces for various generative models



Figure 5: Generated images of BigGAN and ADM-U trained on ImageNet. The samples are randomly selected.

that could consider various aspects of faulty samples. We will discuss more later.

Do the Results Vary with Different Random Seeds for Initialization? We examine how much the evaluation results using random features vary along different random seeds for weight initialization in Figure 6. We can observe that all data points are below the dashed lines, which indicates that the results of determining superior models are not reversed depending on the seed. In Figures 6a and 6b, FIDs of the CNNs with random weights vary with random seeds. Surprisingly, FIDs using Swin and ViT are highly consistent over the five different seeds as shown in Figures 6c and 6d. We surmise that successive softmax operations in the Transformer blocks are helpful in this consistency because they normalize intermediate features. Consequently, in order to obtain reliable results, it would be better to employ a Transformer structure and to incorporate the results using different random seeds (e.g., by averaging).

Do the Results Vary with Different Model Sizes? In addition, we investigate whether the evaluation results are affected by the size of randomly initialized models. Based on the results in Figure 6, we consider Swin and ViT models

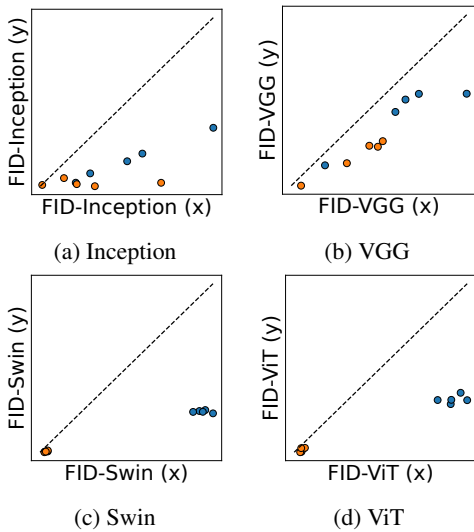


Figure 6: Evaluation results of generative models using random features with five different seeds for weight initialization of the feature networks. The blue circles are FIDs of ADM-U (x-axis) and BigGAN (y-axis), and the orange ones are FIDs of StyleGAN (x-axis) and StyleGAN2 (y-axis).

Models	Swin-T	Swin-S	Swin-B	ViT-T	ViT-B
#Parameters	28M	50M	88M	6M	86M
ImageNet 256²					
BigGAN	1.07	1.07	1.54	0.62	3.14
ADM-U	3.57	3.70	5.03	1.61	6.61

Table 3: FIDs of BigGAN and ADM-U using random features from Swin and ViT with different model sizes

that show consistent results over the different random seeds. We evaluate BigGAN and ADM-U (trained on ImageNet 256²) with random features from larger sizes of Swin and ViT (i.e., Swin-S, Swin-B, and ViT-B). Swin-S has more Transformer blocks in an intermediate layer compared to Swin-T. Swin-B produces higher dimensional outputs than the others (768→1024) with the same number of blocks to that of Swin-S. ViT-B has a deeper structure with a higher dimension of outputs (192→768) than ViT-T.

Table 3 shows the results. Swin-T and Swin-S show similar FIDs, and the superiority of BigGAN is also consistent in Swin-B although the scale of FID increases. ViT-B is also consistent to ViT-T in determining superiority of BigGAN to ADM-U except for the increase of the FID scale. Overall, changing the depth of Swin hardly affects the evaluation results and the dimension of the features of Swin and ViT is related to the scale of FID. Thus, any of the tested models can be used for evaluation, and smaller models like Swin-T and ViT-T have the advantage of computational efficiency.

Are Random Features Adoptable?

In the previous sections, evaluation results using random features are mostly similar to those using trained features. However, the random features sometimes yield different re-

sults from the trained ones (Tables 1 and 2). This is because the random and trained features for the generated images become different, especially for high- and low-level outliers (Figure 4). Thus, we propose to use both types of features in order to evaluate generative models in various aspects.

To examine when the complementarity of the two features can be expected, we measure the Spearman’s rank order correlation coefficient between d_{CLIP} and d_{STEM} for a set of generated images. If the correlation is low, the images have different ranks in terms of d_{CLIP} and d_{STEM} . In other words, high-level and low-level outliers are distinct. As shown in Figure 4, in this case, the evaluation results would be different depending on whether the used feature is trained or random.

For FFHQ, the correlation coefficients of ProjectedGAN, StyleGAN, and StyleGAN2 images are 0.276, 0.238, and 0.152, respectively. For ImageNet, the correlation coefficients of ADM-U and BigGAN are 0.191 and **0.127**, respectively. In the case of BigGAN, which shows the lowest correlation coefficient, random and trained features are especially likely to yield different evaluation results, which is actually observed (Table 2). As shown in Figure 5, it is quite challenging for generative models to learn the distribution of a thousand of classes of the ImageNet dataset. Thus, it would be more likely for generative models to produce various types of faulty images.

Thus, to consider different aspects of outliers, it would be desirable to use both trained and random features for evaluation. In addition, according to the results in Figure 6, we suggest to use a model with a Transformer structure (e.g., Swin or ViT) for random features in order to minimize the effect of randomness.

Conclusion

In this paper, we demystified randomly initialized networks for evaluating generative models. We investigated the random features in various aspects as follows.

- We performed in-depth analysis of random features by examining how the random features capture sample-based and distribution-based differences. We observed that random features focus more on low-level information (e.g., color and brightness) compared to trained ones that focus on high-level semantics (e.g., class labels).
- We investigated faulty images from generative models and separated them into high-level and low-level outliers. We also demonstrated how the differences in the distribution are assessed by random features by replacing real images with the outliers. We found that random features are sensitive to the low-level outliers, showing larger FID to the low-level outliers than the high-level outliers.
- We rigorously evaluated various state-of-the-art generative models using random features. We found that while the evaluation results are similar to those using trained features in most cases, the random features yield different evaluation results from the trained ones when generated images are distinct in terms of high-level and low-level outliers. Therefore, it can be beneficial to employ both trained and random features for thorough evaluation of generative models.

Acknowledgements

This work was supported in part by the Artificial Intelligence Graduate School Program, Yonsei University under Grant 2020-0-01361, and in part by the Ministry of Trade, Industry and Energy (MOTIE) under Grant P0014268.

References

- Amid, E.; Anil, R.; Kotłowski, W.; and Warmuth, M. K. 2022. Learning from randomly initialized neural network features. arXiv:2202.06438.
- Bifıkowski, M.; Sutherland, D. J.; Arbel, M.; and Gretton, A. 2018. Demystifying MMD GANs. In *Proceedings of the International Conference on Learning Representations*.
- Brock, A.; Donahue, J.; and Simonyan, K. 2019. Large scale GAN training for high fidelity natural image synthesis. In *Proceedings of the International Conference on Learning Representations*.
- Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. In *Advances in Neural Information Processing Systems*.
- Chong, M. J.; and Forsyth, D. 2020. Effectively unbiased FID and Inception score and where to find them. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6070–6079.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Kai Li; and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Dhariwal, P.; and Nichol, A. Q. 2021. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the International Conference on Learning Representations*.
- Frankle, J.; and Carbin, M. 2019. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *Proceedings of the International Conference on Learning Representations*.
- Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2672–2680.
- Han, J.; Choi, H.; Choi, Y.; Kim, J.; Ha, J.-W.; and Choi, J. 2022. Rarity score: A new metric to evaluate the uncommonness of synthesized images. arXiv:2206.08549.
- Hazami, L.; Mama, R.; and Thurairatnam, R. 2022. Efficient-VDVAE: Less is more. arXiv:2203.13751.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. In *Proceedings of the International Conference on Computer Vision*, 1026–1034.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, 6629–6640.
- Huang, H.; He, R.; Sun, Z.; Tan, T.; et al. 2018. IntroVAE: Introspective variational autoencoders for photographic image synthesis. In *Advances in Neural Information Processing Systems*.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020. Analyzing and improving the image quality of StyleGAN. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Kynkäänniemi, T.; Karras, T.; Laine, S.; Lehtinen, J.; and Aila, T. 2019. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, volume 32.
- Kynkäänniemi, T.; Karras, T.; Aittala, M.; Aila, T.; and Lehtinen, J. 2022. The role of ImageNet classes in Fréchet Inception distance. arXiv:2203.06026.
- Lee, J.; and Lee, J.-S. 2021. TREND: Truncated generalized normal density estimation of Inception embeddings for GAN evaluation. arXiv:2104.14767.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical vision Transformer using shifted windows. In *Proceedings of the International Conference on Computer Vision*, 10012–10022.
- Morozov, S.; Voynov, A.; and Babenko, A. 2021. On self-supervised image representations for GAN evaluation. In *Proceedings of the International Conference on Learning Representations*.
- Naeem, M. F.; Oh, S. J.; Uh, Y.; Choi, Y.; and Yoo, J. 2020. Reliable fidelity and diversity metrics for generative models. In *Advances in Neural Information Processing Systems*, volume 119, 7176–7185.
- Parmar, G.; Zhang, R.; and Zhu, J.-Y. 2022. On aliased resizing and surprising subtleties in GAN evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 11410–11420.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 8748–8763.
- Ramanujan, V.; Wortsman, M.; Kembhavi, A.; Farhadi, A.; and Rastegari, M. 2020. What’s hidden in a randomly weighted neural network? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Sajjadi, M. S. M.; Bachem, O.; Lucic, M.; Bousquet, O.; and Gelly, S. 2018. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems*, 5234–5243.

- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; Chen, X.; and Chen, X. 2016. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, 2234–2242.
- Sauer, A.; Chitta, K.; Müller, J.; and Geiger, A. 2021. Projected GANs converge faster. In *Advances in Neural Information Processing Systems*.
- Sauer, A.; Schwarz, K.; and Geiger, A. 2022. StyleGAN-XL: Scaling StyleGAN to large diverse datasets. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*, 1–10.
- Simonyan, K.; and Zisserman, A. 2015. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations*.
- Szegedy, C.; Vanhoucke, V.; Ioffe, S.; Shlens, J.; and Wojna, Z. 2016. Rethinking the Inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2818–2826.
- Theis, L.; van den Oord, A.; and Bethge, M. 2016. A note on the evaluation of generative models. In *Proceedings of the International Conference on Learning Representations*.
- Ulyanov, D.; Vedaldi, A.; and Lempitsky, V. 2018. Deep image prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 9446–9454.