

Anomaly Score: Evaluating Generative Models and Individual Generated Images based on Complexity and Vulnerability

Jaehui Hwang*, Junghyuk Lee*, Jong-Seok Lee

School of Integrated Technology / BK21 Graduate Program in Intelligent Semiconductor Technology
 Yonsei University, Korea * Equal contribution

{jaehui.hwang, junghyuklee, jong-seok.lee}@yonsei.ac.kr

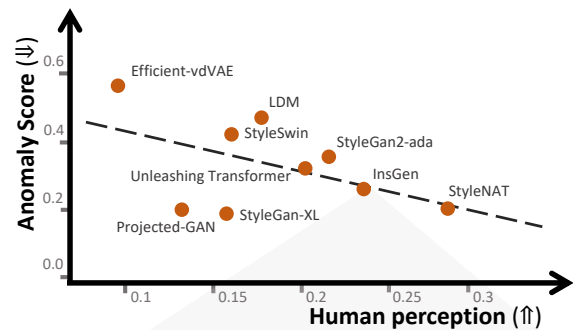
Abstract

With the advancement of generative models, the assessment of generated images becomes increasingly more important. Previous methods measure distances between features of reference and generated images from trained vision models. In this paper, we conduct an extensive investigation into the relationship between the representation space and input space around generated images. We first propose two measures related to the presence of unnatural elements within images: complexity, which indicates how non-linear the representation space is, and vulnerability, which is related to how easily the extracted feature changes by adversarial input changes. Based on these, we introduce a new metric to evaluating image-generative models called anomaly score (AS). Moreover, we propose AS-i (anomaly score for individual images) that can effectively evaluate generated images individually. Experimental results demonstrate the validity of the proposed approach.

1. Introduction

The advancement of deep learning has significantly fostered the development of generative AI, particularly in the domain of image generation. Initially, the focus was primarily on generative adversarial networks (GANs) [19, 28, 30, 52], which employed a generator and a discriminator. Recently, various generative models have been suggested, including autoencoder-based models [22, 32] and diffusion-based models [5, 16, 25, 53]. Simultaneously, evaluating the performance of generative models has become more critical.

The performance of generative models can be evaluated in various ways [7]. Assessing how similar generated images are to real images (i.e., fidelity) is one of the most important and challenging aspects of evaluating the performance. An accurate approach is conducting subjective tests, where human subjects are asked to judge the naturalness of generated images, but this is resource-intensive.



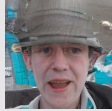



				
AS-i (ψ)	1471.68	906.51	300.92	71.26
Rarity (ψ)	24.52	12.48	Outlier	17.51
Realism (\hat{f})	1.10	1.06	0.97	1.02
Human (\hat{f})	0%	14.29%	78.57%	92.85%

Figure 1. **Proposed AS for evaluating generative models and AS-i for individual images.** The graph on the top shows the proposed AS aligns well with the human perception of evaluating various generative models trained on the FFHQ dataset. On the bottom, several generated images are shown with our AS-i score, rarity score [21], realism score [34], and human evaluation. A value of the human evaluation indicates the proportion of participants who assess that the image is a natural image in our subjective test. The best score for each metric, indicating an image to be the most natural image is highlighted in blue. In terms of naturalness, AS-i shows the best alignment with human evaluation. On the other hand, the rarity score prefers the second image, which is unnatural, as the most common in real images. The realism score also overestimates the leftmost image to be the most realistic.

To efficiently evaluate generative models in various aspects, several objective metrics have been proposed [2, 4, 12, 21, 24, 34, 40, 49, 50, 59]. Most existing metrics involve comparing sample statistics between the sets of real and generated images after extracting features from pre-trained vision

models. For instance, the Fréchet inception distance (FID) utilizes the Inception model [55] to extract Inception features and measures the 2-Wasserstein distance of the Inception features between the real and generated datasets.

However, it is argued that such metrics are often misaligned with the human judgment on naturalness. In [35], it is shown that FID has a null space where the score is unaffected by the change of naturalness of generated images. Furthermore, [54] shows that the score often focuses on image parts unrelated to the naturalness. To sum up, existing metrics are subject to inconsistency to a certain extent in assessing the naturalness of generated images.

We argue that relying only on feature distances is insufficient to assess the naturalness of generated images. Consider a pair of real images having a certain distance in a representation space (i.e., feature space). We can modify one of the images (e.g., by adding noise) so that the modified one is at the same certain distance from the original image. In this case, while the feature distance between the two real images is the same as the distance between the chosen real image and its modified version, the modified image has significantly different contents in terms of naturalness. This also applies to generated images, having certain distances to real images may not be able to accurately represent whether the generated images are natural or not.

In order to address this limitation, in this paper, we propose two novel metrics: anomaly score (AS) for evaluating generative models and anomaly score for individual images (AS-i) for evaluating individual generated images. Instead of simply measuring distances between features, AS and AS-i capture the relationship between the input space and the representation space based on two new perspectives, *complexity* and *vulnerability*. We demonstrate that both metrics show significant correlations with the human-perceived naturalness of generated images (Fig. 1).

We define *complexity* as the amount of variations in the direction of feature changes with respect to linear input changes. A trained neural network model typically implements a non-linear function, and the degree of non-linearity (which we refer to as *complexity*) depends on the feature location in the representation space [31]. We observe that *complexity* tends to become larger for real images compared to unnatural generated images.

Vulnerability reflects how easily the extracted feature of an image is changed due to adversarial input changes. We apply the concept of adversarial attacks [13, 20, 41], which point out weaknesses of deep learning models and are also utilized for tasks capturing characteristics of images, such as out-of-distribution detection [38]. We observe that *vulnerability* tends to be smaller for real images compared to generated unnatural images.

Our contributions are summarized as follows.

1. We introduce *complexity* and *vulnerability*, to examine

the characteristics of the representation space. *Complexity* measures how non-linear the representation space around an image is with respect to the linear input changes. And, *vulnerability* captures how easily an extracted feature is changed by adversarial input changes. We demonstrate that *complexity* and *vulnerability* of generated images are significantly different compared to those of real images.

2. We propose a novel metric called anomaly score (AS) to evaluate generative models in terms of naturalness based on *complexity* and *vulnerability*. AS is the difference of joint distributions of *complexity* and *vulnerability* between the sets of reference real images and generated images, which is quantified by 2D Kolmogorov-Smirnov (KS) statistics. Our method aligns better human judgment about the unnaturalness of generated images compared to the existing method.
3. We suggest the anomaly score for individual images (AS-i) to assess generated images individually. By conducting subjective tests, we demonstrate that AS-i outperforms existing methods for image evaluation.

2. Related works

2.1. Generative models

Generative models have garnered significant attention for their ability to generate realistic data samples. GANs [19, 28, 30, 52] leverage a game-theoretic approach, employing a generator and discriminator in a competitive setting. Variational autoencoders (VAEs) [22, 32], on the other hand, model the distribution of the training data with a likelihood function, learning latent variable representations to generate data that closely matches the observed distribution. Recently, diffusion models [5, 16, 25, 53] have emerged as a powerful approach in generating high-quality images and capturing complex data distributions.

2.2. Evaluation of generative models

Evaluation of generative models mostly involves comparing sample statistics between the generated data and the real target data. Existing metrics can be categorized into three groups based on the way of evaluation: summarizing overall performance of generative models in a single score [4, 24, 50], evaluating different aspects of performance (e.g., fidelity and diversity) of models separately [34, 43, 49], and assessing individual generated images [21, 34].

In the first category, FID [24] is one of the most widely used metrics, which measures how well a generative model can reproduce the target data distribution. It employs the trained Inception model [55] to extract features from the generated and the real images and measures the 2-Wasserstein distance between the two feature distributions. However, it often fails to model the density of the feature




Figure 2. **Tendency of linearity around real data.** We compute the change in the linearity of representation spaces developed by ConvNeXt-tiny and DINO-V2 when random noise is added to real images from the ImageNet dataset.

distributions [37] and to align with human perception [35].

For the second category, Precision and Recall [34, 49] measure fidelity and diversity of samples from generative models, respectively, by extending the classical precision and recall metrics for machine learning. They construct manifolds of the real samples and the generated samples in a certain representation space, then count the proportions of generated and real samples that belong to the real and generated manifolds, respectively. While such twofold metrics are effective for a diagnostic purpose, they are often vulnerable to outliers [43] and suffer from high computational costs for measuring pairwise distances between samples.

Few studies have addressed evaluation of individual samples [21, 34]. The realism score [34] examines the relationship between a generated image and the real manifold. Still, it is based on the manifold that is vulnerable to outliers, thus may become deviated from human evaluation (as will be shown in our experiments). The rarity score [21] focuses on how rare or uncommon a generated image is in order to consider the performance of generative models in terms of creativity. Thus, the rarity score does not provide accurate information about the naturalness of an image.

3. Analyzing representation space around generated images

In this section, we demonstrate that the representation space, which is the space of features extracted by trained models for vision tasks, around generated images exhibits distinct properties in comparison to that around real images.

3.1. Complexity

Motivation. When an input of a trained model changes linearly, it is expected that its feature representation (i.e., output of the model before the softmax function) does not change linearly but instead exhibits curvature due to the non-linearity inherent in deep neural networks. In [31], it is observed that the regions around the features of training images in the representation space appear curved (i.e., complex) after training, i.e., a linear movement in the input space yields a curved trajectory in the representation space;

on the other hand, the representation space near modified images with random noise is less complex.

To verify this, we compute the changes of the average angle between $M(x^k) - M(x^{k-1})$ and $M(x^{k+1}) - M(x^k)$ with varying k , where $M(\cdot)$ indicates the model used for extracting features, which is referred to as feature model for simplicity, and x^k indicates a changed input x by adding a constant noise k times. In Fig. 2, we can observe that the angle decreases as k increases. This phenomenon suggests that the regions far from training images in the representation space are less complex when compared to those around the training images themselves. In a similar context, we hypothesize that the representation space around generated images is less complex than that surrounding reference real images.

Definition. To assess the *complexity* of the representation space around an image, we gradually add random noises to the image and quantify the angular variations in the corresponding feature movements. Let x and N denote an image and a Gaussian random noise vector having a unit length, respectively. x is corrupted gradually by ϵN , i.e., the noised image at step k is computed as $x^k = x^0 + k\epsilon N$, where ϵ is the parameter controlling the magnitude of noise and $x^0 = x$. Then, we calculate the angle between the changes of the output feature within two successive steps (i.e., $k-1$ and k). For instance, when the features for x^{k-2} , x^{k-1} , and x^k are on a straight line, the angle is zero. We compute the *complexity* $C(\cdot)$ by averaging the angles across the changes over multiple steps, which is formulated as follows.

$$C(x) = \frac{1}{K-1} \sum_{k=1}^{K-1} \left(\cos^{-1} \frac{(M(x^k) - M(x^{k-1})) \cdot (M(x^{k+1}) - M(x^k))}{\|M(x^k) - M(x^{k-1})\| \|M(x^{k+1}) - M(x^k)\|} \right), \quad (1)$$

where $M(\cdot)$ denotes the feature model and K represents the total number of steps of adding random noise. Note that the feature model can be chosen among various models trained for vision tasks, including ImageNet classification models and models trained by self-supervised learning methods.

Experimental setup. We conduct an experiment to compare the *complexity* defined above for real images and generated images. We use six pre-trained models for the feature model ($M(\cdot)$): three supervised ImageNet classification models, ResNet50 [23], ViT-S [17], and ConvNeXt-tiny [39], and three self-supervised models, DINO [9], DINO-V2 [46], and CLIP [48]. We utilize three reference datasets, CIFAR10 [33], ImageNet [14], and FFHQ [28]. We employ generated datasets produced by PFGM++ [60] trained with CIFAR10, RQ Transformer [36] trained with ImageNet, and InsGen [61] and StyleNAT [57] trained with FFHQ from dgm-eval [54]. And we utilize 10000 generated images from each dataset. We set $\epsilon = 0.01$ and $K = 10$.

Results. Tab. 1 shows the average values of the *complexity* of the reference dataset and the generated dataset. Overall, the *complexity* of the generated images is smaller than that of the reference images, which is confirmed by statisti-

		ViT	ConvNeXt	DINO-V2
CIFAR10	Reference	0.1046	0.0986	0.0578
	Generated	0.1018	0.0975	0.0573
	<i>p</i> -value	<0.0001*	<0.0001*	<0.0005*
ImageNet	Reference	0.0519	0.0485	0.0337
	Generated	0.0410	0.0287	0.0102
	<i>p</i> -value	<0.01*	<0.0001*	<0.0001*
FFHQ	Reference	0.0643	0.0627	0.0311
	Generated	0.0638	0.0525	0.0302
	<i>p</i> -value	0.2495	<0.0001*	<0.0001*

Table 1. **Complexity of various datasets.** We compare the average value of *complexity* for various feature models, ViT-S [17], ConvNeXt-tiny [39], and DINO-V2 [46]. ‘Reference’ indicates the original dataset, such as CIFAR10 [33], ImageNet [14], and FFHQ [28]. ‘Generated’ denotes the *complexity* obtained from datasets generated by PFGM++ [60], RQ Transformer [36], and InsGen [61] trained with the respective reference datasets. ‘*p*-value’ denotes the *p*-value of the one-tailed *t*-test under the null hypothesis that *complexity* of the generated dataset is equal to that of the reference dataset. The cases with statistical significance are marked with ‘*’.




Figure 3. **Distribution of complexity.** The cumulative distribution function (CDF) of *complexity* is depicted for various feature models and generative models trained on FFHQ. Each row shows distributions for each feature model: ViT-S (left), ConvNeXt-tiny (mid), and DINO-V2 (right). Each column indicates a different type of generative model: InsGen [61] (top) and StyleNAT [57] (bottom). Note that InsGen is assessed as a low-performance model compared to StyleNAT by the human evaluation [54].

cal tests, indicating that as expected, generated images are located in less complex regions in the representation space than the reference (real) images. Fig. 3 presents the cumulative distribution function (CDF) of *complexity* for various feature models and datasets. In most cases, the distribution of the generated dataset differs from that of the reference dataset. Furthermore, the datasets generated by the InsGen model (left column), which generates relatively low-quality images, have more deviated distributions of *complexity* compared to the datasets generated by the StyleNAT model, which generates higher-quality images. The difference between the original dataset and the generated one is

especially prominent in the case of using the ConvNeXt-tiny feature model.

3.2. Vulnerability

Motivation. We exploit the idea of adversarial attack [13, 20, 41] as another tool for examining the relationship between the input space and the representation space. In particular, we generate adversarial perturbations that cause large changes in the representation space with small changes in the input space. As a result, we observe that unnatural components (regions) of images tend to cause large changes. The left column of Fig. 4 shows examples of unnatural components of images, specifically the chin of the girl (upper panel) and the right side of the man (lower panel). For each image, we employ the SLIC algorithm [1] to divide the image into 20 super-pixels. Then, we randomly select 3 to 6 super-pixels among 20 super-pixels, add adversarial perturbations determined by PGD [41] to the selected super-pixels, and obtain the feature of the attacked image from a feature model. We repeat this process 20 times. We apply linear regression between the binary variables indicating whether each super-pixel is attacked and the amount of feature change, and the obtained coefficient of each variable is considered as the contribution of each super-pixel to the feature change (see the appendix for more details). The second to fourth columns of Fig. 4 show the contribution of each super-pixel to the feature changes for different feature models with the color coding. Across all models, unnatural super-pixels are consistently highlighted in red, i.e., the feature of the image is largely changed when we add adversarial perturbations to the unnatural super-pixels. Based on this result, we consider that examining the feature change of an image under adversarial attack can be an effective way to identify unnaturalness of the image.

Definition. The PGD attack [41] is a widely used method for altering prediction results of a model through iterative perturbations applied to images. While the original PGD attack targeting classification models aims to maximize the cross-entropy loss, we maximize the L_2 loss between the features of the original and attacked images. Starting from $x^0 = x + \delta N$, the image is iteratively perturbed as follows. Let P^j and x^j denote the adversarial perturbation and the attacked image at the j -th step of the attack, respectively. The modified PGD update rule is given by

$$\tilde{P}^j = \nabla_{x^j} L_2(M(x), M(x^j)), \quad (2)$$

$$P^j = \alpha \tilde{P}^j / \|\tilde{P}^j\|, \quad (3)$$

$$x^{j+1} = \text{Clip}_{0,255}(x^j + P^j), \quad (4)$$

where α is the attack size at each step, $M(\cdot)$ indicates the feature model, and $L_2(M(x), M(x^j))$ is the L_2 loss between the features of the original and modified inputs. $\text{Clip}_{0,255}(\cdot)$ is a clipping function that limits values to the

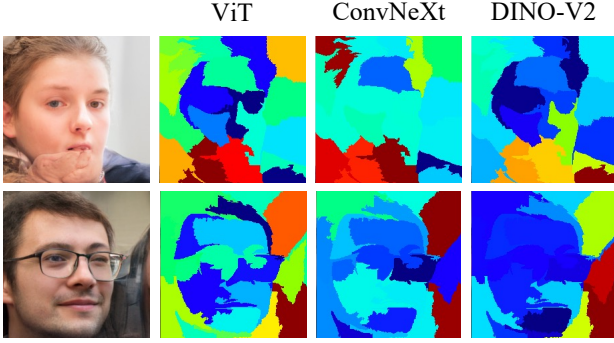


Figure 4. **Image components causing large changes by adversarial attack.** We partition images into super-pixels and assess their contribution to feature changes by the attack. Starting from the left: the original image, and the level of contributions on the changes in the feature extracted by ViT-S [17], ConvNeXt-tiny [39], and DINO-V2 [46], respectively. Red denotes a high level of impact on the changes, while blue indicates a low level of influence on the changes.

		ViT	ConvNeXt	DINO-V2
CIFAR10	Reference	24.56	14.57	29.41
	Generated	24.98	15.24	30.32
	p -value	<0.0001*	<0.0001*	<0.0001*
ImageNet	Reference	11.73	8.69	8.06
	Generated	15.80	12.45	12.85
	p -value	<0.0001*	<0.0001*	<0.0001*
FFHQ	Reference	18.30	14.57	12.90
	Generated	19.22	17.21	16.34
	p -value	<0.0001*	<0.0001*	<0.0001*

Table 2. **Vulnerability of various datasets.** We compare the average value of *vulnerability* for various feature models, ViT-S [17], ConvNeXt-tiny [39], and DINO-V2 [46]. ‘Reference’ indicates the original dataset, such as CIFAR10 [33], ImageNet [14], and FFHQ [28]. ‘Generated’ denotes the *vulnerability* obtained from datasets generated by PFGM++ [60], RQ Transformer [36], and InsGen [61] trained with the respective reference datasets. ‘ p -value’ denotes the p -value of the one-tailed t -test under the null hypothesis that *vulnerability* of the generated dataset is equal to that of the reference dataset. The cases with statistical significance are marked with ‘*’.

range between 0 and 255. Then, we define *vulnerability* of image x , $V(x)$, as follows:

$$V(x) = \text{dist}(M(x), M(x^J)), \quad (5)$$

where J is the total number of steps of the attack and $\text{dist}(A, B)$ indicates the L_2 distance between two vectors A and B .

Experimental setup. We conduct an experiment to examine the validity of the *vulnerability* for evaluating unnaturalness of generated images. We set $\alpha = 0.01$, $\delta = 10^{-6}$, and

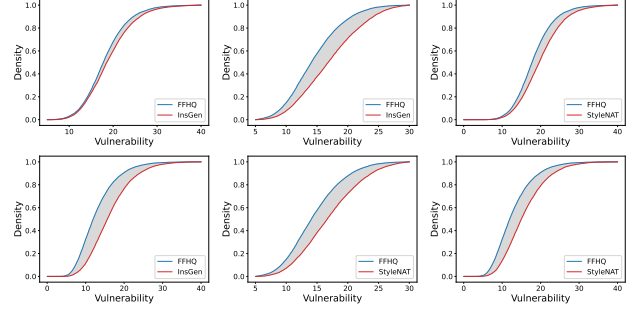


Figure 5. **Distribution of vulnerability.** The cumulative distribution function (CDF) of *vulnerability* is depicted for various feature models and generative models trained on FFHQ. Each row shows distributions for each feature model: ViT-S (left), ConvNeXt-tiny (mid), and DINO-V2 (right). Each column indicates a different type of generative model: InsGen [61] (top) and StyleNAT [57] (bottom). Note that InsGen is assessed as a low-performance model compared to StyleNAT by the human evaluation [54].




Figure 6. **Examples of the MVT dataset.** Images having low MVT levels are clear for people to recognize. On the other hand, a high MVT level means that people need a long time to recognize the image contents due to weak naturalness.

$J = 10$ in this experiment, while keeping all other settings, such as feature models and generative models, the same as those used in Sec. 3.1.

Results. Tab. 2 shows the average *vulnerability* of images from various datasets on ViT-S, ConvNeXt-tiny, and DINO-V2 as feature models. Additional results for other models are shown in the appendix. It can be seen that the *vulnerability* of generated images is larger than that of reference images. Fig. 5 also shows the significant disparity in the *vulnerability* distribution between real and generated images. Furthermore, the high-quality generated dataset by the StyleNAT model has a more similar distribution to that of the reference dataset compared to the relatively low-quality generated dataset by the InsGen model in most cases, which is consistent with the results observed in the comparison of *complexity*.

Vulnerability and naturalness. We further demonstrate the relationship between *vulnerability* and naturalness. We employ the MVT dataset [42], which contains information on the difficulty for humans to recognize objects in each image by measuring the minimum viewing time required for recognition. In the dataset, the viewing time is separated into several levels called MVT levels (i.e., 17 ms, 50 ms,




Figure 7. **Vulnerability vs. MVT levels.** The average value of *vulnerability* of images having each MVT level. Note that the value of *vulnerability* is normalized by the maximum and the minimum values in each feature model.

100 ms, 150 ms, 250 ms, and 10,000 ms). Each MVT level means the time within which over 50% of participants can correctly classify an object. It is linked to the level of unnaturalness of the images. In other words, when an image contains clear content (Fig. 6a), humans swiftly identify the depicted object within 50 ms. Conversely, an image with unnatural components (Fig. 6b) poses a challenge for human recognition, requiring over 10,000 ms.

Fig. 7 shows the *vulnerability* with respect to the MVT level utilizing the MVT dataset. Here, “Err” indicates that over 50% of participants misclassify an image given sufficient time. The MVT level is highly correlated with the *vulnerability* on all feature models, i.e., low naturalness is related to high *vulnerability*.

4. Evaluating generative models

In the previous section, we explored the distinct properties in the representation space around generated images in comparison to those of real images. The representation space around generated images is more complex (*complexity*) and contains a certain path that is vulnerable to adversarial changes (*vulnerability*). Thus, we propose a novel metric for evaluating generative models by capturing anomalies based on both *complexity* and *vulnerability*, called anomaly score.

4.1. Anomaly score for generative models

We define anomaly score (AS) for evaluating generative models as the difference of the bivariate distributions of *complexity* and *vulnerability* between the reference and generated datasets. We denote the set of anomaly vectors for the dataset X as $A(X) = \{[C(x), V(x)]\}_{x \in X}$.

To compare the distributions of the generated dataset and the reference dataset, we employ the Kolmogorov-Smirnov (KS) statistic that measures a non-parametric statistical difference between two distributions. Our anomaly score, AS, utilizing the 2D KS statistic, is defined as

$$AS = \text{Sup} \left| \text{CDF}(A(X^{real})) - \text{CDF}(A(X^{generated})) \right|, \quad (6)$$




Figure 8. **Performances of our method and FID using DINO-V2 for overall datasets.** Each dot represents a distinct dataset generated by a generative model. A high human error rate indicates a high-quality dataset, while a high AS score means a low-quality dataset.

where X^{real} is the reference dataset for the generated dataset, $X^{generated}$, and $\text{CDF}(\cdot)$ is the cumulative distribution function of input vectors. We compute 2D KS statistics by referencing github.com/syrte/ndtest. AS has the minimum value of 0 when the two distributions are identical and the maximum value of 1 when the two distributions are significantly different, i.e., their CDFs are completely non-overlapped.

4.2. Experiments

Setup. To evaluate the effectiveness of our metric, we examine whether our scores align with the subjective scores reported in [54]. In this subjective test, human viewers responded whether a given image appeared fake (generated) or real. Note that a high human error rate for a generated dataset indicates that many viewers cannot identify the images as fake ones, implying that the images in the dataset have high quality and are realistic. We employ generated datasets produced by various generative models including GANs [3, 8, 10, 18, 26, 27, 29, 30, 45, 51, 52, 57, 58, 61, 62], VAEs [32, 36], a flow-based generative model [11], and diffusion models [5, 6, 15, 16, 25, 44, 47, 56, 60] from dgm-eval [54] and utilize 10000 generated images from each dataset. We set K of Eq. (1) and J of Eq. (5) to 10. During feature extraction, our method needs K -times inferences for computing *complexity* and J -times inferences for computing *vulnerability*, while feature extraction FID needs one inference. In our setting, we need 20-times inferences.

Results. Fig. 8 shows the comprehensive performance of our method with DINO-V2 as a feature model. It contains evaluation results on all generative models targeting one of the CIFAR10, ImageNet, and FFHQ datasets. For comparison, the conventional FID with DINO-V2 as a feature model is also evaluated. We can observe that our method shows better performance than FID in terms of evaluating naturalness of generated images. Our method has a relatively high correlation (-0.81 pearson correlation coefficient (PCC)) with human perception, while FID has a lower correlation (-0.54 PCC).

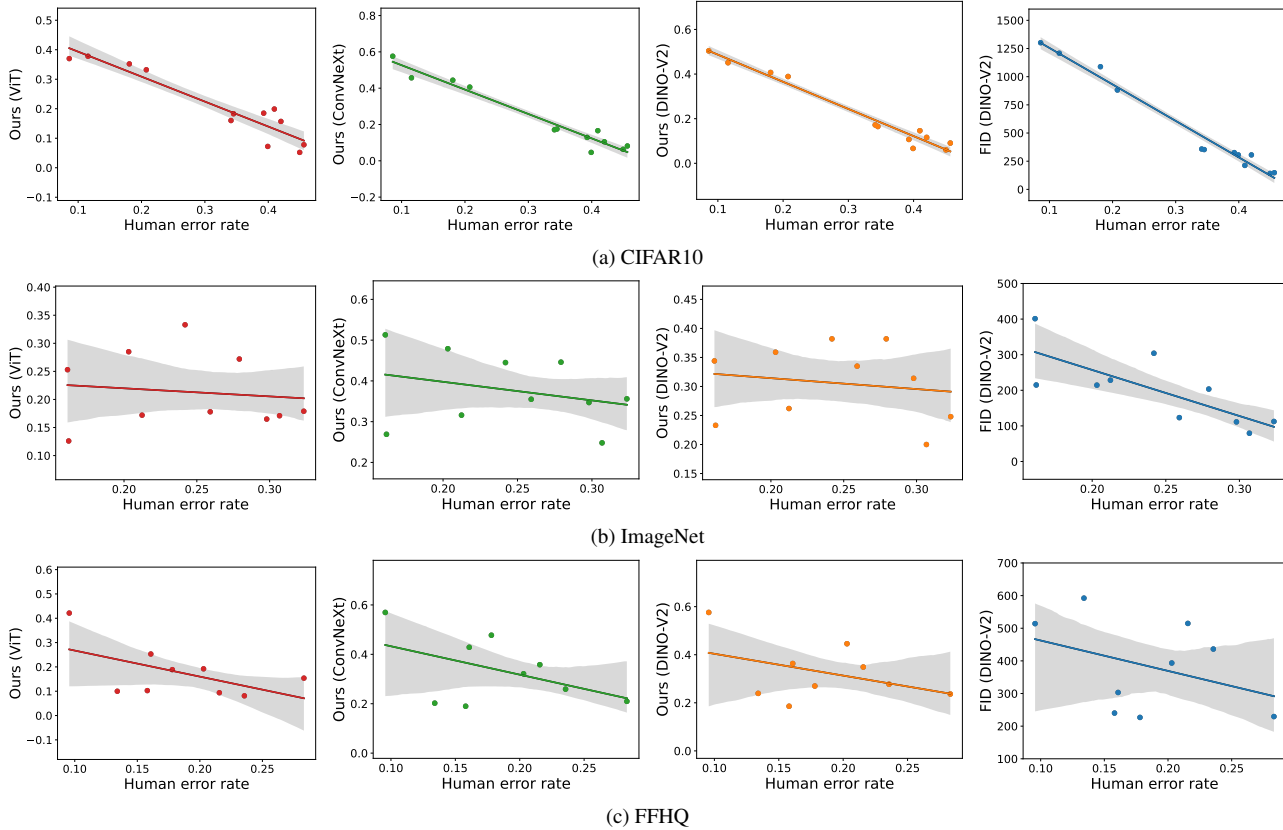


Figure 9. **Overall results of evaluating generative models on each dataset type.** Each dot represents a distinct dataset generated by a generative model. A high human error rate indicates a high-quality dataset, while a high AS score means a low-quality dataset. The first three columns show AS with different feature models: DINO-V2, ConvNeXt-tiny, and ViT-S, respectively. The last column is the result of FID [24] with the DINO-V2 model.

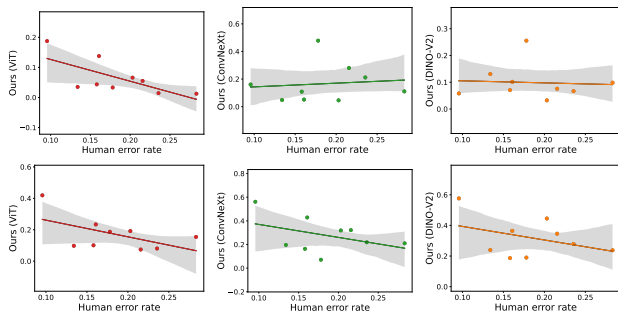


Figure 10. **Anomaly score using one measure.** Each dot represents a distinct generated dataset produced by a generative model for the FFHQ dataset. A high human error rate means a high-quality dataset, while a high AS implies a low-quality dataset. Each column corresponds to a different feature model, and each row represents the results obtained by using the distribution of *complexity* (top) or *vulnerability* (bottom).

Fig. 9 shows AS and FID with various feature models for evaluating generative models trained on different datasets. Our method outperforms FID (the last column) on the FFHQ dataset. AS has a relatively high correlation (-0.56 PCC) with human perception, while FID has a

lower correlation (-0.38 PCC). We can observe that AS with DINO-V2 (-0.98 PCC), ConvNeXt-tiny (-0.30 PCC), and ViT-S (-0.56 PCC) are well aligned with the human error rate on CIFAR10, ImageNet, and FFHQ, respectively. Additional results regarding the evaluation of generative models using other feature models are presented in the appendix.

1D test. In our method, we utilize both *complexity* and *vulnerability* of the images. We conduct a comparative analysis using only one of the two in the anomaly vector, where the 1D KS statistic is used instead of the 2D KS statistic. Fig. 10 presents the results on the FFHQ dataset. *Complexity* performs well with ViT-S (-0.69 PCC) but not with ConvNeXt-tiny and DINO-V2 (0.11 and -0.07 PCCs, respectively). *Vulnerability* shows decent performance over all feature models (-0.55, -0.40, and -0.39 PCCs, respectively), but is outperformed by our method employing both *complexity* and *vulnerability* (-0.56, -0.48, and -0.42 PCCs, respectively). The results on the overall dataset are shown in Tab. 3. We can observe that the 2D-metric shows better performance than the 1D-metrics as well as FID when we use DINO-V2 as a feature model. These results demonstrate the benefit of employing both *complexity* and *vulnerability*

	PCC	SRCC
1D-Complexity	-0.376	-0.388
1D-Vulnerability	-0.790	-0.722
2D (Ours)	<u>-0.806</u>	<u>-0.738</u>
FID	-0.540	-0.525

Table 3. **Evaluation results of generated datasets.** We compare Pearson correlation coefficient (PCC) and Spearman rank correlation coefficient (SRCC) between the human error rate and the result of each metric by using DINO-V2 as a feature model.

in our AS in terms of both performance and consistency.

5. Evaluating individual generated images

In this section, we propose a method for evaluating generated images individually based on *complexity* and *vulnerability* defined in Sec. 3.

5.1. Anomaly score for individual generated images

We adopt a simple and effective formula to capture the properties of *complexity* and *vulnerability* of an image in a single score. The proposed anomaly score for individual images (AS-i) is defined as follows:

$$\text{AS-i}(x) = \frac{V(x)}{C(x)}, \quad (7)$$

where x represents an individual generated image. When the image is natural, *complexity* around the image increases (large $C(x)$) and *vulnerability* of the image decreases (small $V(x)$), hence AS-i becomes small. On the other hand, when the image is unnatural, AS-i becomes large.

5.2. Subjective test

Experimental settings. To verify that AS-i captures human judgement well in terms of the naturalness of images, we conduct a subjective test using the images generated by InsGen [61]. We set five levels of AS-i (highest in the dataset, 900, 600, 300, and lowest in the dataset) using the ConvNeXt-tiny feature model. The highest level of AS-i ranges from 1434 to 1746 with an average of 1561 and the lowest level ranges from 53 to 80 with an average of 71. Then, for each level, we form an image subset by sampling 20 images having AS-i close to the level. Fig. 11 shows several examples of images that belong to the highest, medium (600), and the lowest levels. 14 participants are asked to judge if each image appears natural. We consider an image to be natural if over 50% of the participants respond that the image is natural as follows the rule used in [42]. Then, we obtain the proportion of the number of images that are identified as natural in each subset.

Results and comparison. Tab. 4 presents the results of the subjective test. It is observed that as AS-i level increases,




Figure 11. **Examples having various levels of AS-i.**

AS-i	Human	Rarity [21]	Realism [34]
Low	0.75	21.82	1.039
300	0.50	18.47	1.036
600	0.40	21.51	1.030
900	0.30	21.54	1.032
high	0.25	20.16	1.010

Table 4. **Evaluation of AS-i.** For each level of AS-i, the proportion of images judged as natural by participants, the average rarity score, and the average realism score are shown.

fewer images in a subset are judged as natural in a consistent manner. On the other hand, the average rarity score [21] and realism score [34] are not aligned well with human perception. The PCCs for AS-i, rarity score, and realism score are -0.88, -0.49, and 0.63, respectively.

6. Conclusion

We have proposed new metrics, AS and AS-i, for evaluation of generative models and individual generated images, respectively. Both are based on *complexity* and *vulnerability*, which examine the representation space around the images. *Complexity* captures the curvedness of the representation space, while *vulnerability* tests the changes in the representation space under adversarial attack. We demonstrated that the proposed metrics accord well with human judgments and outperform existing metrics.

Acknowledgement This work is supported by the Korea Agency for Infrastructure Technology Advancement(KAIA) grant funded by the Ministry of Land, Infrastructure and Transport (Grant KA163379).

References

- [1] Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *IEEE TPAMI*, 34(11):2274–2282, 2012. 4
- [2] Motasem Alfarrar, Juan C Pérez, Anna Frühstück, Philip HS Torr, Peter Wonka, and Bernard Ghanem. On the robustness of quality measures for GANs. In *ECCV*, 2022. 1
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017. 6
- [4] Mikołaj Bińkowski, Danica J. Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *ICLR*, 2018. 1, 2
- [5] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023. 1, 2, 6
- [6] Sam Bond-Taylor, Peter Hesse, Hiroshi Sasaki, Toby P Breckon, and Chris G Willcocks. Unleashing transformers: Parallel token prediction with discrete absorbing diffusion for fast high-resolution image generation from vector-quantized codes. In *ECCV*, 2022. 6
- [7] Ali Borji. Pros and cons of gan evaluation measures. *Computer Vision and Image Understanding*, 179:41–65, 2019. 1
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *ICLR*, 2018. 6
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 3
- [10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *CVPR*, 2022. 6
- [11] Ricky TQ Chen, Jens Behrmann, David K Duvenaud, and Jörn-Henrik Jacobsen. Residual flows for invertible generative modeling. *NeurIPS*, 2019. 6
- [12] Min Jin Chong and David Forsyth. Effectively unbiased FID and Inception score and where to find them. In *CVPR*, 2020. 1
- [13] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *ICML*, 2020. 2, 4
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 3, 4, 5
- [15] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 2021. 6
- [16] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *NeurIPS*, 2021. 1, 2, 6
- [17] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3, 4, 5
- [18] Kangning Du, Huaqiang Zhou, Lin Cao, Yanan Guo, and Tao Wang. Mhgan: Multi-hierarchies generative adversarial network for high-quality face sketch synthesis. *IEEE Access*, 8:212995–213011, 2020. 6
- [19] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1, 2
- [20] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 2, 4
- [21] Jiyeon Han, Hwanil Choi, Yunjey Choi, Junho Kim, Jung-Woo Ha, and Jaesik Choi. Rarity score : A new metric to evaluate the uncommonness of synthesized images. In *ICLR*, 2023. 1, 2, 3, 8
- [22] Louay Hazami, Rayhane Mama, and Ragavan Thuraiatnam. Efficient-VDVAE: Less is more. *arXiv preprint arXiv:2203.13751*, 2022. 1, 2
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [24] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *NeurIPS*, 2017. 1, 2, 7
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 2020. 1, 2, 6
- [26] Minguk Kang, Woohyeon Shim, Minsu Cho, and Jaesik Park. Rebooting aggan: Auxiliary classifier gans with stable training. *NeurIPS*, 2021. 6
- [27] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up gans for text-to-image synthesis. In *CVPR*, 2023. 6
- [28] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1, 2, 3, 4, 5
- [29] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 6
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *CVPR*, 2020. 1, 2, 6
- [31] Juyeop Kim, Junha Park, Songkuk Kim, and Jong-Seok Lee. Curved representation space of vision transformers. In *AAAI*, 2024. 2, 3
- [32] Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *ICLR*, 2014. 1, 2, 6
- [33] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. Technical report, 2009. 3, 4, 5
- [34] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *NeurIPS*, 2019. 1, 2, 3, 8
- [35] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of ImageNet classes in fréchet Inception distance. In *ICLR*, 2023. 2, 3

- [36] Doyup Lee, Chihon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. Autoregressive image generation using residual quantization. In *CVPR*, 2022. 3, 4, 5, 6
- [37] Junghyuk Lee and Jong-Seok Lee. TREND: Truncated generalized normal density estimation of Inception embeddings for GAN evaluation. In *ECCV*, 2022. 3
- [38] Shiyu Liang, Yixuan Li, and Rayadurgam Srikant. Enhancing the reliability of out-of-distribution image detection in neural networks. In *ICLR*, 2018. 2
- [39] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 3, 4, 5
- [40] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. Are GANs created equal? a large-scale study. In *NeurIPS*, 2018. 1
- [41] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 2, 4
- [42] David Mayo, Jesse Cummings, Xinyu Lin, Dan Gutfreund, Boris Katz, and Andrei Barbu. How hard are computer vision datasets? calibrating dataset difficulty to viewing time. In *NIPS Datasets and Benchmarks Track*, 2023. 5, 8
- [43] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *NeurIPS*, 2020. 2, 3
- [44] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 6
- [45] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *ICML*, 2017. 6
- [46] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 3, 4, 5
- [47] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 6
- [48] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*. PMLR, 2021. 3
- [49] Mehdi S. M. Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *NeurIPS*, 2018. 1, 2, 3
- [50] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training GANs. In *NeurIPS*, 2016. 1, 2
- [51] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. *NeurIPS*, 2021. 6
- [52] Axel Sauer, Katja Schwarz, and Andreas Geiger. StyleGAN-XL: Scaling StyleGAN to large diverse datasets. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Proceedings*, pages 1–10, 2022. 1, 2, 6
- [53] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*. PMLR, 2015. 1, 2
- [54] George Stein, Jesse C. Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Leigh Ross, Valentin Vilecroze, Zhaoyan Liu, Anthony L. Caterini, J. Eric T. Taylor, and Gabriel Loaizaganem. Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models. *arXiv preprint arXiv:2306.04675*, 2023. 2, 3, 4, 5, 6
- [55] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception architecture for computer vision. In *CVPR*, 2016. 2
- [56] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *NeurIPS*, 2021. 6
- [57] Steven Walton, Ali Hassani, Xingqian Xu, Zhangyang Wang, and Humphrey Shi. StyleNAT: Giving each head a new perspective. *arXiv preprint arXiv:2211.05770*, 2022. 3, 4, 5, 6
- [58] Yan Wu, Jeff Donahue, David Balduzzi, Karen Simonyan, and Timothy Lillicrap. Logan: Latent optimization for generative adversarial networks. *arXiv preprint arXiv:1912.00953*, 2019. 6
- [59] Qiantong Xu, Gao Huang, Yang Yuan, Chuan Guo, Yu Sun, Felix Wu, and Kilian Weinberger. An empirical study on evaluation metrics of generative adversarial networks. *arXiv preprint arXiv:1806.07755*, 2018. 1
- [60] Yilun Xu, Ziming Liu, Yonglong Tian, Shangyuan Tong, Max Tegmark, and Tommi Jaakkola. PFGM++: Unlocking the potential of physics-inspired generative models. In *ICML*, 2023. 3, 4, 5, 6
- [61] Ceyuan Yang, Yujun Shen, Yinghao Xu, and Bolei Zhou. Data-efficient instance generation from instance discrimination. *NeurIPS*, 2021. 3, 4, 5, 6, 8
- [62] Bowen Zhang, Shuyang Gu, Bo Zhang, Jianmin Bao, Dong Chen, Fang Wen, Yong Wang, and Baining Guo. Styleswin: Transformer-based gan for high-resolution image generation. In *CVPR*, 2022. 6