

Probabilistic Precision and Recall

Towards Reliable Evaluation of Generative Models

Dogyun Park

Department of Computer Science, Korea University
Seoul, South Korea
gg933@korea.ac.kr

Suhyun Kim*

Korea Institute of Science and Technology
Seoul, South Korea
dr.suhyun.kim@gmail.com

Abstract

Assessing the fidelity and diversity of the generative model is a difficult but important issue for technological advancement. So, recent papers have introduced k -Nearest Neighbor (k NN) based precision-recall metrics to break down the statistical distance into fidelity and diversity. While they provide an intuitive method, we thoroughly analyze these metrics and identify oversimplified assumptions and undesirable properties of k NN that result in unreliable evaluation, such as susceptibility to outliers and insensitivity to distributional changes. Thus, we propose novel metrics, P -precision and P -recall (PP&PR), based on a probabilistic approach that address the problems. Through extensive investigations on toy experiments and state-of-the-art generative models, we show that our PP&PR provide more reliable estimates for comparing fidelity and diversity than the existing metrics. The codes are available at https://github.com/kdst-team/Probablistic_precision_recall.

1. Introduction

Modern state-of-the-art generative models such as Variational Autoencoders (VAEs) [18, 24], Generative Adversarial Networks (GANs) [8, 15, 13, 14], and Score-based Diffusion Models (SDMs) [12, 17, 29, 6, 32, 25, 30] have shown outstanding achievements that surpass human evaluative capabilities. Consequently, reliable evaluation metrics for comparing the performance of generative models are critical for further technological advancement. The widely used Fréchet Inception Distance (FID) evaluation metric [11] has been praised for its robustness and consistency with human perceptual evaluation. However, FID only provides a single score and does not diagnose whether the model lacks fidelity or diversity, two essential aspects for improving generative models. Recent papers [26, 28, 20, 22, 1] have introduced various two-value metrics to measure fidelity

and diversity, respectively. Especially, k -nearest neighbor (k NN) based metrics such as Improved Precision and Recall (IP&IR) [20] and Density and Coverage (D&C) [22] have been used to assess fidelity and diversity on various generative models [6, 35, 2, 24, 13, 4, 23]. The pipeline of these metrics [20, 22] starts with estimating distribution support of real and fake distribution by constructing sample-specific hyperspheres around each sample using k NN. Then, they measure the expectation of the likelihood of fake (real) samples with respect to the estimated real (fake) distribution support to represent fidelity (diversity).

Unfortunately, both IP&IR and D&C metrics present certain challenges that compromise their practical reliability. Specifically, IP&IR operate under the assumption of constant-density within the distribution support and overlook the overestimation linked to k NN, leading to an inflated distribution support with constant density. This scenario becomes especially problematic when the distribution houses outliers, resulting in metric overestimation (See Fig. 1a). Density metric [22] attempts to address these limitations by accumulating density over overlapped k NN hyperspheres. Yet, our findings suggest that it still remains unreliable to outliers. Furthermore, it tends to exhibit high variance, due to its extremely wide bound, which further limits the reliability of evaluations. On the other hand, Coverage metric [22] presents an alternative method that measures robust diversity even with the presence of outliers in the fake distribution. However, we show that it suffers from a conceptual limitation that makes it insensitive to changes in fake distribution. An in-depth discussion of these constraints is provided in Sec. 2. In the end, these issues underscores the need for careful revision of k NN-based metrics to provide more reliable evaluations.

We, therefore, propose a novel set of metrics, P -precision and P -recall (PP&PR), which aim to provide more reliable estimations for the fidelity and diversity of generative models. In particular, we derive our metrics based on a *probabilistic approach* that relaxes the oversimplified assumptions and undesirable properties of k NN, addressing the limitations of

*corresponding author

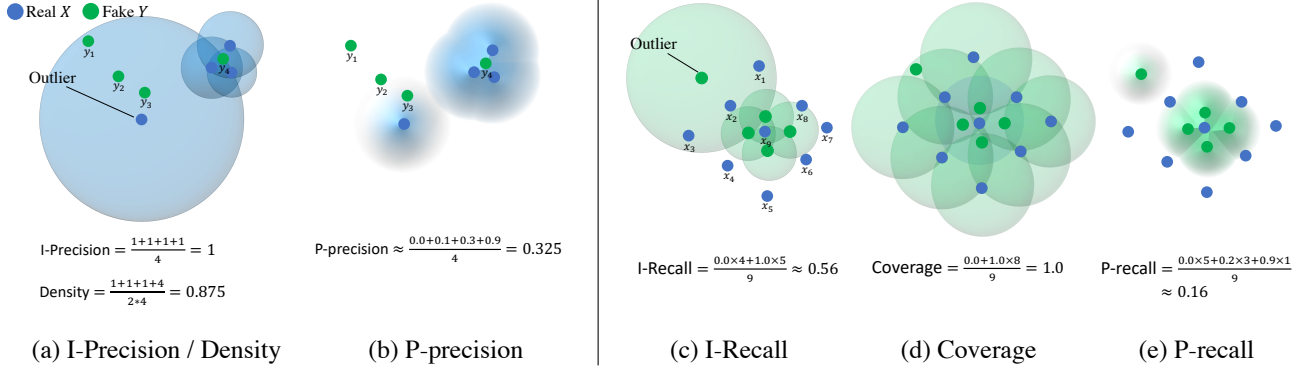


Figure 1: **Examples of IP&IR, D&C, and PP&PR** (better zoom-in). For simplicity, we use $k = 2$ for k -Nearest Neighbor in all metrics. (a), (c): Due to the overestimation of k NN by the outlier and constant-density assumption within hyperspheres, I-precision, I-recall, and Density denote the same values for different y in an overly large space (See y_1, y_2 , and y_3 in (a)), leading to unreliably-overestimated values. (b), (e): On the other hand, our P-precision and P-recall assign different scores to different y based on a probabilistic approach and address overestimation of k NN (See Sec. 4 for details). (d): This illustrates a case when Coverage exhibits its conceptual limitation. Even if the fake samples have lower relative diversity compared to real samples, they can still be included in multiple hyperspheres of real samples, leading to high Coverage.

existing metrics. Through extensive experiments on both toy and real-world datasets, we systematically demonstrate the superiority of our proposed metrics over IP&IR and D&C. To summarize, our contributions are as follows:

- We propose a set of two-value metrics, P-precision and P-recall, that provide reliable fidelity and diversity measures for generative models.
- We thoroughly analyze the limitations of existing metrics (IP&IR and D&C), such as vulnerability to outliers and insensitivity to diversity, and introduce a novel method based on the probabilistic approach that remedies the limitations.
- Through a series of experiments on toy and real-world datasets with state-of-the-art generative models, we empirically verify the limitations of current metrics and demonstrate the efficacy of our proposed metrics.

2. Preliminary

2.1. Feature embedding for evaluation

Statistical comparison of high dimensional data in the raw input space \mathcal{X} is difficult. Therefore, the usual evaluation starts with converting real data $X \sim P$ and fake data $Y \sim Q$ into embeddings using a CNN feature Φ ($\Phi(X)$ and $\Phi(Y)$) [11, 20, 27, 22]. This embedded feature space is believed to have more “meaningful” information to compare two samples [7, 33]. For simplicity, we use X and Y in place of $\Phi(X)$ and $\Phi(Y)$ for the remainder of the paper.

2.2. k NN-based fidelity and diversity measures

Now, we briefly introduce current k NN-based fidelity and diversity measures and discuss their limitations in Sec. 3.

Improved Precision and Recall (IP&IR). To represent fidelity, Precision is defined as the expected likelihood of a fake sample belonging to the support of the real distribution, which refers to the set of possible real samples whose probability density function is non-zero. Likewise, Recall is defined as the expected likelihood of a real sample belonging to the support of fake distribution to measure diversity.

$$\begin{aligned} \text{Precision} &:= \mathbb{E}_{Y \sim Q} [Y \in \text{supp}(P)] \\ \text{Recall} &:= \mathbb{E}_{X \sim P} [X \in \text{supp}(Q)], \end{aligned} \quad (1)$$

where $\text{supp}(\cdot)$ indicates the ground truth support of distribution. Then, IP&IR [20] approximates them using the empirical observations ($\{x_i\}_{i=1}^N$ and $\{y_j\}_{j=1}^M$) as

$$\begin{aligned} \text{Precision} &\approx \text{I-Precision} \\ &:= \frac{1}{M} \sum_{j=1}^M \text{BSR}_P(y_j) \\ \text{Recall} &\approx \text{I-Recall} \end{aligned} \quad (2)$$

$$:= \frac{1}{N} \sum_{i=1}^N \text{BSR}_Q(x_i).$$

The scoring rule of IP&IR, which we name as **Binary Scoring Rule (BSR)** for comparison with other scoring rules, is

$$\text{BSR}_P(y_j) = \begin{cases} 1, & \text{if } y_j \in \cup_{i=1}^N B(x_i, \text{NND}_k(x_i)) \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

where $B(x, r)$ is the hypersphere in \mathbb{R}^D around x with radius r , and $\text{NND}_k(x_i)$ denotes the Euclidean distance from x_i to the k^{th} nearest neighbor among $\{x_i\}_{i=1}^N$ excluding itself. In other words, $\text{BSR}_P(y_j)$ assigns a binary score of 1 or 0 to y_j based on whether it falls inside the union of hyperspheres. It is worth noting that BSR employs a *deterministic approach* that determines the support as the union of k NN hyperspheres without accounting for any uncertainty regarding the true support. This restricts the scoring rule to binary assignments, leading to a constant density within the approximated distribution support.

I-recall is defined using $\text{BSR}_Q(x_i)$, which is defined symmetrically with $\text{BSR}_P(y_j)$. It builds hyperspheres around fake samples and tests real samples.

Density and Coverage (D&C). D&C use different approaches to define fidelity and diversity in comparison to IP&IR. Specifically, Density [22] tackles the issue of constant-density over distribution support in IP&IR by quantifying the degree of overlap between k NN hyperspheres. To this end, Density defines a scoring rule that estimates the density function within distribution support by the superposition of hyperspheres, as opposed to the union in IP&IR. We denote this scoring rule as **Density Scoring Rule (DSR)**.

$$\text{DSR}_P(y_j) = \frac{1}{k} \sum_{i=1}^N \mathbb{1}_{y_j \in B(x_i, \text{NND}_k(x_i))} \quad (4)$$

where k is for the k -nearest neighbors. DSR employs a kind of Parzen window density estimation [10] with a uniform kernel whose width is adjusted automatically, except that the integral of the kernel does not sum up to 1. Now, Density is defined as the expectation of DSR over observed samples as

$$\text{Density} := \frac{1}{M} \sum_{j=1}^M \text{DSR}_P(y_j) \quad (5)$$

For diversity, the authors employ another scoring rule, **Coverage Scoring Rule (CSR)**, which is not symmetric to DSR. It builds k NN hyperspheres around the real samples and counts the real samples whose hyperspheres are occupied by the fake samples. Coverage has demonstrated effectiveness in handling outliers in fake distribution, as the hyperspheres of real samples are not influenced by fake outliers.

$$\begin{aligned} \text{CSR}_Q(x_i) &= \mathbb{1}_{\exists j \text{ s.t. } y_j \in B(x_i, \text{NND}_k(x_i))} \\ \text{Coverage} &:= \frac{1}{N} \sum_{i=1}^N \text{CSR}_Q(x_i) \end{aligned} \quad (6)$$

3. Limitations with scoring rules of IP&IR and D&C

Binary scoring rule. BSR can only assign constant density within an approximated distribution support, which can

hinder reliable estimates when comparing distributions with complex density functions. Moreover, the use of k NN in BSR makes it more susceptible to outliers as it unreliably inflates the size of the hypersphere due to the presence of such outliers. This results in overestimated value as illustrated in Fig. 1a and 1c. Both I-precision and I-recall face similar issues since they symmetrically use BSR.

Density scoring rule. We contend that DSR remains vulnerable to outliers as it *accumulates* density over the *overestimated constant-density* hyperspheres. For instance, if there are s outliers, all hyperspheres of outliers can overlap due to the overestimated size of hyperspheres. This causes the overlapped density to accumulate to s/k (note that DSR scales the overlapped density by a factor of $1/k$ in Eq. 4), resulting in a rapid increase with just a few outliers. We demonstrate this quantitatively with the experiment in Sec. 5. Additionally, reducing the density by a factor of $1/k$ makes the metric highly dependent on the choice of k , which is not desirable as we demonstrate in Fig. 3.

Lastly, DSR is bounded within $[0, N/k]$ (e.g., $N/k = 50000/5 = 10000$), which can result in samples with abnormally high values of DSR, compromising stable evaluations with high standard deviations as demonstrated in Fig. 2c.

Coverage scoring rule. While CSR can effectively handle outliers in fake distribution, it exhibits a conceptual limitation: a lack of sensitivity to the comparative diversity between fake and real distributions. This arises because CSR doesn't factor in distances between fake samples. Instead, its focus is solely on the count of real sample hyperspheres occupied by fake samples. This leads to a scenario where if real samples are more dispersed compared to fake ones, meaning the fake samples exhibit less relative diversity, the fake samples could still be included in multiple real hyperspheres. This, in turn, can produce a falsely elevated Coverage metric (a conceptual illustration is provided in Fig. 1d). Conversely, Coverage can be low if fake samples are more widely distributed compared to real samples. We elucidate this phenomenon quantitatively in Fig. 5b.

4. Method

To address the aforementioned limitations, we propose a novel scoring rule based on a probabilistic approach. Our motivation arises from the recognition that accurately estimating the ground truth support is infeasible with empirical observations alone. As a result, it becomes difficult to determine whether a sample belongs to the ground truth support with certainty (0 or 1). Therefore, we adopt a probabilistic perspective and define the probability of a sample belonging to the support, taking into account the uncertainty surrounding the ground truth support. Then, our scoring rule can

assign lower values to samples that are less likely to belong to the support such as outliers, effectively addressing outlier concerns. We also define the scoring rules for fidelity and diversity symmetrically, similar to IP&IR, to avoid the conceptual limitation of CSR. Although we mainly focus on describing the details of fidelity in the paper, diversity is defined symmetrically.

4.1. P-precision and P-recall (PP&PR)

Given the observations $\{x_i\}_{i=1}^N$, we assume that there are myriad possible supports of P , denoted as S_P . The collection of these supports is represented by \mathbb{S}_P . Subsequently, we introduce **Probabilistic-Precision** that characterizes precision as the probability of y_j being an element of support randomly chosen from \mathbb{S}_P . This is articulated as

$$\begin{aligned} \text{P-precision} &:= \mathbb{E}_{Y \sim Q, S_P \in \mathbb{S}_P} [Y \in S_P] \\ &= \mathbb{E}_{Y \sim Q} [\mathbb{E}_{S_P \in \mathbb{S}_P} [Y \in S_P]] \\ &= \mathbb{E}_{Y \sim Q} [\Pr(Y \in S_P)] \end{aligned} \quad (7)$$

Then, we approximate it with given observations $\{y_j\}_{j=1}^M$ as

$$\text{P-precision} \approx \frac{1}{M} \sum_{j=1}^M \Pr(y_j \in S_P). \quad (8)$$

In a similar manner, **Probabilistic-recall** is approximated as

$$\text{P-recall} \approx \frac{1}{N} \sum_{i=1}^N \Pr(x_i \in S_Q). \quad (9)$$

4.2. Probabilistic scoring rule

To estimate $\Pr(y_j \in S_P)$, we first divide S_P into subsets around each observation as

$$S_P := \bigcup_{i=1}^N s_P(x_i), \quad (10)$$

where $s_P(x_i)$, denoted as the subsupport around x_i , represents a subset of S_P that comprises elements for which the nearest reference among $\{x_i\}_{i=1}^N$ is x_i . Then, we can further elucidate $\Pr(y_j \in S_P)$ as

$$\begin{aligned} \Pr(y_j \in S_P) &= 1 - \Pr(y_j \notin S_P) \\ &= 1 - \Pr(y_j \notin \bigcup_{i=1}^N s_P(x_i)) \\ &= 1 - \Pr(y_j \in \bigcap_{i=1}^N s_P^c(x_i)) \\ &= 1 - \prod_{i=1}^N \Pr(y_j \in s_P^c(x_i)) \\ &= 1 - \prod_{i=1}^N (1 - \Pr(y_j \in s_P(x_i))) \end{aligned} \quad (11)$$

using Eq. 10 and complement rule, where $s_P^c(\cdot)$ denotes the complement set of $s_P(\cdot)$. The equality in the fourth line of Eq. 11 comes from assuming the independence of events $\{y_j \in s_P^c(x_i) | i = 1, 2, \dots, N\}$.

Now, we define our new scoring rule termed Probabilistic Scoring Rules (PSR) from the last equation in Eq. 11:

$$\begin{aligned} \text{PSR}_P(y_j) &:= 1 - \prod_{i=1}^N (1 - \Pr(y_j \in s_P(x_i))) \\ &= \Pr(y_j \in S_P) \end{aligned} \quad (12)$$

In contrast with deterministic methods such as BSR, which make certain assumptions about the definitive shape of subsupport, PSR estimates the probability of y_j to belong to the *uncertain* support by combining the probability to belong to the subsupport around each observation. Utilizing this probabilistic approach allows us to depart from defining the specific subsupport—a task that, in practice, can be ambiguous. Rather, we infer the likelihood of a sample being part of the subsupport based on an intuitive premise: samples in closer proximity to x_i are more likely to be an element of the subsupport compared to those further away. Accordingly, we approximate the probability $\Pr(y_j \in s_P(x_i))$ as

$$\Pr(y_j \in s_P(x_i)) \approx \begin{cases} 1 - \frac{d}{R}, & \text{if } d \leq R \\ 0, & \text{otherwise} \end{cases} \quad (13)$$

where d is l_2 distance between x_i and y_j . From Eq. 13, the probability is one when $d = 0$, given that we know for certain that x_i belongs to the subsupport. The probability monotonically decreases with an increase in d until it hits the threshold distance R , where the probability converges to zero. By Eq. 12 and Eq. 13, our proposed PSR can assign varying values to samples based on their likelihood of belonging to the support. Moreover, it is bounded within $[0,1]$ (see supplement for the proof) unlike DSR which has an extremely wide bound. Note that while other distance metrics d or monotonically decreasing functions could serve this purpose, we opted to use l_2 distance and a linear function since it was sufficiently sensitive based on our empirical results. We leave it for future work.

It is also important to select R in Eq. 13 since it controls the robustness of the scoring rule. Instead of sample-specific $R(x_i)$ based on k NN as used in IP&IR and D&C, we propose to use identical R for all samples x_i . We set R as a multiple of the average of k NN distances among $\{x_i\}_{i=1}^N$, since it can be used as an effective normalization factor in high-dimensional space. This is similar to using a fixed-variance Gaussian kernel for Kernel Density Estimation and offers several advantages over sample-specific k NN: 1) it suppresses the overestimation of k NN by the outlier as shown in Fig. 1b, and 2) makes metric more robust to the choice of k as shown in Fig. 3.




Figure 2: (a) Behavior of fidelity metrics between two Gaussian distributions $X \sim N(0, I)$ and $Y \sim N(u\mathbf{1}, I)$ as u moves between $[-3, 3]$ with outlier $x_o \sim N(-2, I)$ added to X . (b) Behavior of diversity metrics between two Gaussian distributions $X \sim N(u\mathbf{1}, I)$ and $Y \sim N(0, I)$ as u moves between $[-3, 3]$ with outlier $y_o \sim N(2, I)$ added to Y . (c) Estimated bias from the presumed true value between two identical Gaussian distributions for different numbers of datasets N . The line shows the means and the shaded area denotes standard deviations across 50 runs.

Finally, putting it all together, we propose P-precision as

$$\text{P-precision} = \frac{1}{M} \sum_{j=1}^M \text{PSR}_P(y_j) \quad (14)$$

where

$$\text{PSR}_P(y_j) = 1 - \prod_{i=1}^N (1 - \Pr(y_j \in s_P(x_i))),$$

$$\Pr(y_j \in s_P(x_i)) = \begin{cases} 1 - \frac{\|x_i - y_j\|_2}{R}, & \text{if } \|x_i - y_j\|_2 \leq R \\ 0, & \text{otherwise} \end{cases}$$

and

$$R = \frac{a}{N} \sum_{i=1}^N \text{NND}_k(x_i).$$

a is a hyperparameter that controls the magnitude of R , and k is for the k -nearest neighbors. In a similar manner, we propose P-recall as

$$\text{P-recall} = \frac{1}{N} \sum_{i=1}^N \text{PSR}_Q(x_i) \quad (15)$$

where

$$\text{PSR}_Q(x_i) = 1 - \prod_{j=1}^M (1 - \Pr(x_i \in s_Q(y_j))),$$

$$\Pr(x_i \in s_Q(y_j)) = \begin{cases} 1 - \frac{\|x_i - y_j\|_2}{R}, & \text{if } \|x_i - y_j\|_2 \leq R \\ 0, & \text{otherwise} \end{cases}$$

and

$$R = \frac{a}{M} \sum_{j=1}^M \text{NND}_k(y_j).$$

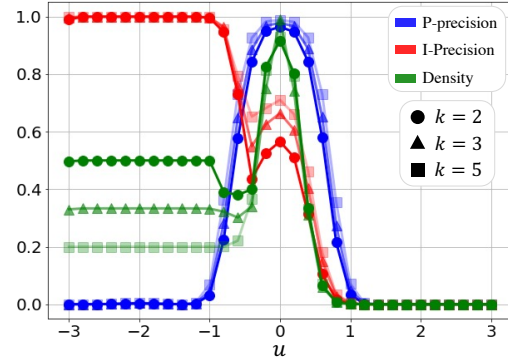


Figure 3: Ablation over k . We measure each metric for different k between $X \sim N(0, I)$ and $Y \sim N(u, I)$ as $u \in [-3.0, 3.0]$ with outlier $x_o \sim N(-2, I)$ added to X .

5. Experiments

In this section, we systematically demonstrate the limitations of current metrics and the improvements of our metrics through toy experiments. We then show the experiments with state-of-the-art generative models, highlighting that our metrics are more reliable for evaluating generative models than IP&IR and D&C. We use the settings described below for the remainder of the experiments unless stated otherwise. We set $k = 4$ and $a = 1.2$ for our metrics in all experiments. For IP&IR and D&C, we follow the recommended value of $k = 3$ and $k = 5$, respectively. In the toy experiments, we generate Gaussian distributions in \mathbb{R}^{64} and set $M = N = 10000$. For the experiments evaluating generative models, we use features extracted from the ImageNet [19] pre-trained Inception model [31] after the second fully connected layer. We compute metrics between



Figure 4: Qualitative examples sorted according to L . We used StyleGAN [15] trained on FFHQ [15] and BigGAN [3] trained on CIFAR-10 [19]. The top two rows are images with the highest L meaning they have high PSR but low DSR. Conversely, the bottom two rows are images with the lowest L , indicating low PSR and high DSR.




Figure 5: Behavior of metrics between $X \sim N(0, I)$ and $Y \sim N(0, vI)$ as v changes between $[0.2, 1.5]$. Because Density goes over 1 (up to nearly 1000), the y -axis for Density is on the right side of the plot for better visualization.

all available real samples (up to 50K) and 50K generated images following [6].

5.1. Toy experiments

5.1.1 Robustness to outliers

Gaussian distribution We first investigate how the metrics behave to outliers through a simple toy experiment with Gaussian distributions, following [22]. We set real $X \sim N(0, I)$ and fake $Y \sim N(u\mathbf{1}, I)$ where $\mathbf{1}$ is a vector of ones and I is the identity matrix, and measure the metrics between X and Y as u varies in $[-3, 3]$. To introduce an outlier, we add a sample $x_o \sim N(-2, I)$ to the set of real samples. A reliable metric should consistently decrease as u moves away from 0, regardless of the outlier. Fig. 2a shows the result for fidelity: I-precision gives overestimated values when

$u \leq 0$ due to the enlarged constant-density hypersphere by the outlier. Although Density alleviates the problem by reducing the density inside the hypersphere by $1/k$, it does not reflect the change for a wide range of u (See when $u \leq -1$). On the other hand, P-precision denotes reliable estimates even when the fake distribution comes closer to the outlier (PP = 0.006 when $u = -2$), and also shows a robust decrease as u moves away from 0. We observe a similar result on the diversity measures except for Coverage (See Fig. 2b).

In addition, we perform an ablation over k to investigate whether the improvement of our metric over the outlier is due to a properly chosen k . The same experiment described above is repeated with different k and the metric values are reported in Fig. 3. Our P-precision shows strong consistency over k , due to the employment of identical R for all real samples. However, I-precision continues to suffer from outliers, regardless of the value of k . Moreover, Density exhibits high variation over different values of k as it reduces the density by $1/k$, demonstrating that Density needs properly chosen k to achieve adequate robustness to outliers.

These results suggest that our metrics are more reliable to estimate fidelity and diversity between distributions, even in the presence of outliers.

5.1.2 Stability of metrics

We also test the stability of metrics by measuring the empirical bias of the metric from the presumed true value, which we approximate by averaging across 50 runs with 50K samples, between two identical Gaussian distributions. Fig. 2c shows the bias from the true value with respect to the number of datasets used for estimation, along with the average and

standard deviation across 50 runs. Our PP&PR show a small bias and standard deviation, even for small N , while IP&IR show a large bias. In addition, despite the theoretical value when two distributions are identical [22], Density does not behave nicely: at $N \leq 10000$, Density shows quite a large bias and standard deviation due to the extremely wide bound. This demonstrates that our PP&PR provides more reliable estimates than other metrics even for small datasets.

5.1.3 Reflecting fidelity and diversity

The main motive of the two-value metric is to capture both fidelity and diversity between two distributions. To assess whether these metrics are capable of reflecting fidelity and diversity, we design a scenario where the two values are at odds. Specifically, we generate samples from $X \sim N(0, I)$ and $Y \sim N(0, vI)$, and vary the value of v from 0.2 to 1.5. As v increases, we expect fidelity to decrease and diversity to increase. As shown in Fig. 5, while Density behaves as expected, Coverage does not; it remains constant at 1 as v increases from 0.2 to 1, despite the increasing relative diversity of fake distribution. We argue that this counter-intuitive behavior is due to the conceptual limitation of Coverage as discussed in Sec. 3. When v increases beyond 1, Coverage becomes 0, as fake samples are no longer included in the hyperspheres of real samples. In contrast, PP&PR and IP&IR nicely reflect the expected behavior.

5.2. Evaluating generative models

5.2.1 Comparing scoring rules

Here, we compare the scoring rules of fidelity metrics to demonstrate how nicely the estimated density reflects the quality of the fake sample. We compare PSR with DSR, as BSR only denotes a binary score. For effective comparison, we focus on cases where two scoring rules have conflicting results. Specifically, we compute $L = \{\text{PSR}_P(y_j) - \text{DSR}_P(y_j) \mid \forall_j\}$ using generative models $Q(Y)$ (StyleGAN [15] and BigGAN [3]) trained on real-world datasets $P(X)$ (CIFAR-10 [19] and FFHQ [15]). A high value of L indicates a high PSR but low DSR, while a low value of L indicates the opposite. For a fair comparison, we normalize DSR by its maximum value so that it ranges between 0 and 1. We show images of the highest L and lowest L in Fig. 4 (without cherry-picking). In general, images with a high value of L show more realistic and clear objects, while images with a low L are often distorted with artifacts (See supplement for more examples). Based on the assumption that realistic images should correspond to the high density of real distribution, our results in Fig. 4 qualitatively demonstrate the superiority of PSR over DSR in reflecting the semantic quality of generated images. We also conduct a user study selecting preferences between the

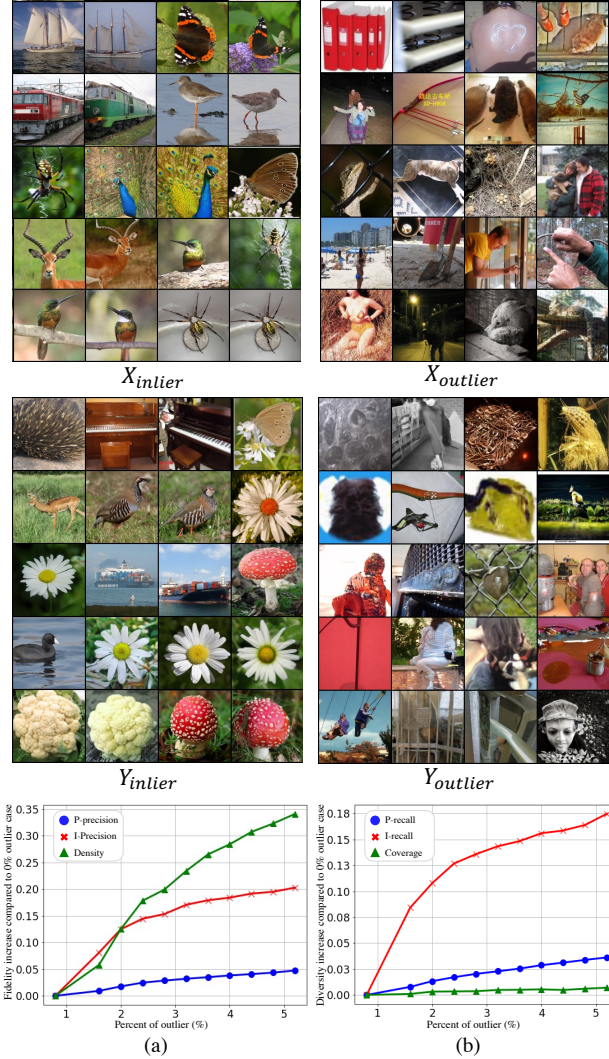


Figure 6: Behavior of metrics with respect to the number of outliers when evaluating the generative model Y on a real-world dataset X . (a) Fidelity metrics. (b) Diversity metrics.

highest L and the lowest L and it also shows a similar result to our observation (See supplement).

5.2.2 Robustness to outliers

To investigate the efficacy of our proposed metrics in handling outliers with generative models on real-world datasets, we conduct experiments on the state-of-the-art diffusion model (e.g., ADM [6]), which currently has gained attention for its powerful performance, trained on ImageNet [19] dataset (128^2 resolution). Then, we adopt the outlier selection criterion used in [22] which uses the distance to the k^{th} nearest neighbor. We apply this criterion to select inliers and outliers (at a ratio of 0.95:0.05) for both real and fake samples generated by ADM, resulting in the sets

Table 1: Quantitative result of various generative models on real-world datasets. The reported values are obtained by measuring each metric five times and taking the average. Bolded values indicate the top scores evaluated by the corresponding metric.

Model	FID↓	PP↑	PR↑	F ₁ ↑	IP↑	IR↑	F ₁ ↑	D↑	C↑	F ₁ ↑
ImageNet 256x256										
ADM [6]	4.95	0.538	0.732	0.621	0.681	0.688	0.684	1.52	0.876	1.112
ADM-G [6]	4.58	0.699	0.587	0.638	0.818	0.606	0.696	2.071	0.956	1.307
BigGAN [3]	8.12	0.751	0.465	0.574	0.874	0.403	0.551	2.481	0.945	1.368
AFHQv2 512x512										
StyleGAN2 [16]	4.62	0.568	0.689	0.623	0.716	0.494	0.584	1.886	0.790	1.113
StyleGAN3-R [14]	4.40	0.564	0.725	0.634	0.685	0.591	0.635	1.576	0.770	1.034
StyleGAN3-T [14]	4.04	0.567	0.727	0.637	0.699	0.578	0.632	1.624	0.792	1.065
LSUN Bedroom 256x256										
DDPM [16]	4.88	0.799	0.749	0.773	0.606	0.444	0.512	1.701	0.977	1.241
ADM [6]	1.91	0.839	0.731	0.781	0.659	0.494	0.565	1.929	0.993	1.311
StyleGAN2 [16]	2.35	0.801	0.753	0.776	0.591	0.501	0.543	1.732	0.986	1.257

$X = X_{inlier} \cup X_{outlier}$ and $Y = Y_{inlier} \cup Y_{outlier}$. We provide examples of the outliers and inliers in Fig. 6. Note that the inliers tend to exhibit more recognizable objects with clear semantic features compared to the outliers.

We first observe how fidelity metrics between X_{inlier} and Y respond when a portion of X_{inlier} is gradually substituted with the subset of $X_{outlier}$. To visualize this, we plot the metric increments compared to the case when no outliers are included in X_{inlier} , as shown in Fig. 6a. The results show that I-precision rapidly increases (0.2 increments) when only 5% of inliers are gradually replaced by outliers, confirming its vulnerability to outliers. Although Density is less affected by outliers initially, it eventually surpasses the increment of I-precision as it accumulates the density over hyperspheres, supporting our claim that Density is still vulnerable to outliers. On the other hand, our P-precision shows the least increment (0.05 increments for 5% outliers), validating its effectiveness in handling outliers.

Likewise, we also test the robustness of diversity measures between X and Y_{inlier} by gradually replacing a portion of Y_{inlier} with the subset of $Y_{outlier}$. In Fig. 6b, we observe that I-recall exhibits a similar vulnerability to outliers as I-precision, while our P-recall is less affected by the presence of outliers. Besides, Coverage demonstrates strong robustness to outliers due to its alternative diversity measure. However, this comes at the cost of losing accuracy, as demonstrated in the experiments in Fig. 5b and 7b.

5.2.3 Reflecting fidelity and diversity

We now verify whether the metrics accurately reflect the change in fidelity and diversity of fake samples generated

from generative models. In order to do that, we employ classifier guidance [6] during the image sampling from ADM. This technique uses gradients of a pre-trained classifier to perform conditional sampling from pre-trained unconditional diffusion models. By adjusting the scale of gradients from the classifier, we can explicitly control the balance between fidelity and diversity of generated images. For instance, a larger gradient scale results in the diffusion model sampling more from the modes of the pre-trained classifier, yielding higher fidelity but lower diversity. For more details, please refer to [6]. Therefore, we sample images from ADM, the same model as in Sec. 5.2.2, using gradient scales ranging from 0.5 to 10.0, and measure metrics (IP&IR, D&C and PP&PR) between real and generated samples for each scale. Fig. 7 shows the results; our PP&PR and IP&IR effectively capture the trade-off between fidelity and diversity whereas D&C does not exhibit a clear relationship with the gradient scale. This is due to the conceptual limitation of Coverage in accurately measuring the diversity in generated samples.

5.2.4 Image generation benchmark

To evaluate the practical usefulness of our proposed metrics in assessing the state-of-the-art generative models, we present the quantitative results on real-world datasets (LSUN bedroom [34], AFHQ-v2 [5], and ImageNet) for various types of generative models, along with FID. The results, presented in Tab. 1, are obtained by measuring each metric five times and taking their average. For clarity in the presentation, the variance associated with these measurements can be found in the supplementary materials. While FID serves as a holistic reference point for distinguishing superior




Figure 7: Fidelity versus diversity when increasing the gradient scale from 0.5 to 10.0 for sampling diffusion models with classifier guidance [6]. A larger gradient scale tends to result in higher fidelity and lower diversity.




Figure 8: Ablation study on PSR: (A): probabilistic estimation for subsupport. (B): identical R .

model performance, our metrics delve deeper, elaborating on aspects of fidelity and diversity. For instance, our metrics explain that the improvement of ADM over StyleGAN [15] on LSUN bedroom in terms of FID is due to an increase in fidelity at the cost of diversity. An analysis between ADM and ADM-G further reveals that the classifier guidance in ADM-G augments fidelity, yet curtails diversity. This demonstrates the usefulness of our metrics in elaborating on the effect of the specific component in models.

Moreover, to observe whether the metrics harmonize with the widely adopted FID metric, we analyze the consistency between FID and two-value metrics. This is achieved by computing the F_1 -score [9], which is a harmonic mean of fidelity and diversity. The results demonstrate that F_1 -scores of our metric mirror the trends set by FID (as underscored by the bold scores for each dataset in Tab. 1).

6. Ablation study

Here, we conduct an ablation study on our proposed scoring rule, PSR, to demonstrate the effectiveness of our approach: (A) probabilistic estimation for subsupport (Eq. 13) and (B) identical R . We repeat the outlier toy experiment

with Gaussian distributions in Sec. 5.1.1 with different configurations and the result for fidelity metric is shown in Fig. 8. In both cases where the feature (A) and (B) are individually removed, the metric becomes more susceptible to outlier than the case where both are used ((A)+(B)) but less affected than the case where both are all removed (see when $u \in [-3, -1]$). This demonstrates that our probabilistic approach for subsupport and addressing overestimation of sample-specific k NN both contribute to estimating accurate density and combining them together maximizes it.

7. Conclusions and limitations

In this paper, we introduced a set of two-value metrics, P-precision and P-recall, based on the probabilistic approach that addresses the limitations of existing metrics (IP&IR and D&C). We have conducted a systematic investigation of the limitations and demonstrated the effectiveness of our proposed metrics by evaluating toy experiments and state-of-the-art generative models. Nevertheless, we recognize there is still room for improvement. While we utilized commonly-used ImageNet embeddings [11, 20, 22] for feature extraction, recent studies such as [1, 21], have proposed domain-agnostic embeddings to enhance the feature quality compared to previous embeddings that rely on dataset-specific statistics. Thus, future research focusing on the quality of feature embedding combined with our method may further enhance the reliability and usefulness of the metric.

Acknowledgements This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-00456, Development of Ultra-high Speech Quality Technology for remote Multi-speaker Conference System and No.RS-2023-00225630, Development of Artificial Intelligence for Text-based 3D Movie Generation).

References

- [1] Ahmed Alaa, Boris Van Breugel, Evgeny S Saveliev, and Mihaela van der Schaar. How faithful is your synthetic data? sample-level metrics for evaluating and auditing generative models. In *International Conference on Machine Learning*, pages 290–306. PMLR, 2022. 1, 9
- [2] Ivan Anokhin, Kirill Demochkin, Taras Khakhulin, Gleb Sterkin, Victor Lempitsky, and Denis Korzhenkov. Image generators with conditionally-independent pixel synthesis. In *Conference on Computer Vision and Pattern Recognition*, pages 14278–14287, 2021. 1
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *International Conference on Learning Representations*, 2019. 6, 7, 8
- [4] Eric R Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. pi-gan: Periodic implicit generative

- adversarial networks for 3d-aware image synthesis. In *Conference on Computer Vision and Pattern Recognition*, pages 5799–5809, 2021. [1](#)
- [5] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *Conference on Computer Vision and Pattern Recognition*, pages 8188–8197, 2020. [8](#)
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [1](#), [6](#), [7](#), [8](#), [9](#)
- [7] David A Forsyth and Jean Ponce. *Computer vision: a modern approach*. Pearson, 2012. [2](#)
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63, 2020. [1](#)
- [9] Cyril Goutte and Eric Gaussier. A probabilistic interpretation of precision, recall and f-score, with implication for evaluation. In *European conference on information retrieval*, pages 345–359. Springer, 2005. [9](#)
- [10] Peter E Hart, David G Stork, and Richard O Duda. *Pattern classification*. Wiley Hoboken, 2000. [3](#)
- [11] 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. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [9](#)
- [12] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33, 2020. [1](#)
- [13] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020. [1](#)
- [14] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in neural information processing systems*, 34, 2021. [1](#), [8](#)
- [15] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. [1](#), [6](#), [7](#), [9](#)
- [16] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. [8](#)
- [17] Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. *Advances in neural information processing systems*, 34, 2021. [1](#)
- [18] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014. [1](#)
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. [5](#), [6](#), [7](#)
- [20] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *Advances in neural information processing systems*, 32, 2019. [1](#), [2](#), [9](#)
- [21] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fréchet inception distance. In *International Conference on Learning Representations*, 2023. [9](#)
- [22] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *International Conference on Machine Learning*, pages 7176–7185. PMLR, 2020. [1](#), [2](#), [3](#), [6](#), [7](#), [9](#)
- [23] Evangelos Ntavelis, Mohamad Shahbazi, Iason Kastanis, Radu Timofte, Martin Danelljan, and Luc Van Gool. Arbitrary-scale image synthesis. In *Conference on Computer Vision and Pattern Recognition*, pages 11533–11542, 2022. [1](#)
- [24] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2. In *Advances in neural information processing systems*, 2019. [1](#)
- [25] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [1](#)
- [26] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. *Advances in neural information processing systems*, 31, 2018. [1](#)
- [27] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. *Advances in neural information processing systems*, 29, 2016. [2](#)
- [28] Loic Simon, Ryan Webster, and Julien Rabin. Revisiting precision recall definition for generative modeling. In *International Conference on Machine Learning*, pages 5799–5808. PMLR, 2019. [1](#)
- [29] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-based diffusion models. *Advances in neural information processing systems*, 34, 2021. [1](#)
- [30] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021. [1](#)
- [31] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. [5](#)
- [32] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34, 2021. [1](#)
- [33] 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. [2](#)

- [34] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 8
- [35] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. The spatially-correlative loss for various image translation tasks. In *Conference on Computer Vision and Pattern Recognition*, pages 16407–16417, 2021. 1