



**HAL**  
open science

# A New Perspective on Precision and Recall for Generative Models

Benjamin Sykes, Loïc Simon, Julien Rabin, Jalal Fadili

► **To cite this version:**

Benjamin Sykes, Loïc Simon, Julien Rabin, Jalal Fadili. A New Perspective on Precision and Recall for Generative Models. 2026. <hal-05344404v2>

**HAL Id: hal-05344404**

**<https://hal.science/hal-05344404v2>**

Preprint submitted on 12 May 2026

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# A New Perspective on Precision and Recall for Generative Models

**Benjamin Sykes**

BENJAMIN.SYKES@UNICAEN.FR

*NORMANDIE UNIV, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 CAEN, FRANCE*

**Loïc Simon**

LOIC.SIMON@ENSICAEN.FR

*NORMANDIE UNIV, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 CAEN, FRANCE*

**Julien Rabin**

JULIEN.RABIN@ENSICAEN.FR

*NORMANDIE UNIV, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 CAEN, FRANCE*

**Jalal Fadili**

JALAL.FADILI@ENSICAEN.FR

*NORMANDIE UNIV, UNICAEN, ENSICAEN, CNRS, GREYC, 14000 CAEN, FRANCE*

## Abstract

With the recent success of generative models in image and text, the question of their evaluation has recently gained a lot of attention. While most methods from the state of the art rely on scalar metrics, the introduction of Precision and Recall (PR) for generative model has opened up a new avenue of research. The associated PR curve allows for a richer analysis, but their estimation poses several challenges. In this paper, we present a new framework for estimating entire PR curves based on a binary classification standpoint. We conduct a thorough statistical analysis of the proposed estimates. As a byproduct, we obtain a minimax upper bound on the PR estimation risk. We also show that our framework extends several landmark PR metrics of the literature which by design are restrained to the extreme values of the curve. Finally, we study the different behaviors of the curves obtained experimentally in various settings.

**Keywords:** machine-learning, generative-modeling, statistics, evaluation metric

## 1 Introduction

### 1.1 Problem Statement

In this paper, we consider metrics designed to evaluate the faithfulness of a generative model to the distribution it is trained to capture. The problem intrinsically consists in evaluating the “closeness” of a target distribution, hereafter denoted by  $P$  and a generated one, denoted by  $Q$ . This problem in itself is a challenging one as generative models have recently been able to create ever more complex data in high dimension. Additionally, the notion of quality in generative models is inherently subjective — unlike in supervised learning, there is no unique correct output or ground-truth for one generated data. The notion of “closeness” stated above still remains to be precisely defined and several methods have been proposed to deal with this question, ranging from human feedback to the use of pre-trained encoder as a learned metric. We will first briefly see how different methods have emerged with the




Figure 1: Left: two illustrative distributions  $P$  and  $Q$  (example borrowed from Kynkäänniemi et al. (2019)) — Right: the PR-curve is the frontier of the shaded area composed of all admissible PR pairs  $(\beta, \alpha)$ . In essence, these pairs represent the mass of  $P$  and  $Q$  that one can recover by selecting a subset of the common support (gray area on the left). More precisely, by selecting regions of high likelihood of  $P$ , one trades Precision ( $\alpha$ ) in favor of Recall ( $\beta$ ). The extreme values  $\beta_0(P, Q)$  and  $\alpha_\infty(P, Q)$  embody the respective masses of the entire common support.

will to tackle generative model evaluation and which refreshing point of view we present in this field.

## 1.2 Context

In the early days of training deep generative models for still images, several scalar metrics were defined to objectively capture a notion of quality on generated images. While these metrics were appealing at first to compare two generative models, they were omitting a large quantity of information about the distributions under study. To alleviate these drawbacks, Sajjadi et al. (2018) defined a new and formal framework to evaluate generative models. This framework allows to capture the quality of a generative model in terms of **fidelity**: “*how close the generated data are from the target data?*” and **diversity**: “*how much of the target space do the generated images cover?*”. In this framework, both information are captured through a Precision and Recall curve which is inherently different from those of binary classification. The intuitive description of the metric is described in Figure 1.

While PR-curves are strongly theoretically grounded, practitioners have preferred scalar metrics of Precision and Recall over the entire curve. Whereas the initial formulation by Sajjadi et al. (2018) defined clusters based on samples using K-Means, Kynkäänniemi et al. (2019) defined scalar Precision and Recall metrics based on kNN support estimation. Naeem et al. (2020) proposed other metrics to evaluate fidelity and diversity, still based on kNN yet in a different manner. One key contribution was that these metrics were less sensitive to outliers. In Khayatkhoei and AbdAlmageed (2023), the authors proposed a new metric constructed in order to be robust in high dimension. This metric summarizes as the minimum between the two previous metrics. Finally, Park and Kim (2023) created novel Precision and Recall metrics estimating not simply the support of the distribution under

study but estimating the probability density in a Parzen like framework. A more detailed overview of the state of the art will be provided later on in Section 4.

On the one hand, this effort is justified by practical arguments which entail simplicity as well as an intuitive understanding of the extreme values. On the other hand, restricting to extreme values only allows for a very shallow comparison. For instance, extreme Precision and Recall scores may both attain their saturation level of 1 for entirely different distributions (as long as both distributions share the same support). On the contrary, Precision and Recall curves do capture the fundamental differences between two distributions in the transition zone between the two extreme values. Another concern with published PR metrics relates to their statistical consistency which is either provably lacking or at best only partially studied. A notable exception is (Pillutla et al., 2023), in which a formal and theoretical framework is presented in the context of divergence frontiers introduced by (Djulonga et al., 2020). The authors propose several scalar summary metrics for such performance curves, and conduct a theoretical analysis to provide estimation error bounds for  $f$ -divergences. It is also worth mentioning (Kim et al., 2023) in which the authors propose a metric called topological Precision and Recall (topP&R) along with its consistency analysis. However this latter is limited to extreme PR metrics and relies on strong assumptions (no tail in the distributions – see their assumption A3) to obtain meaningful results.

### 1.3 Contributions and Link to Prior Works

In this work, we propose a novel approach for evaluating generative models with Precision and Recall, through the prism of binary classification. We also exploit the link between Precision and Recall and the Total Variation (TV) distance in order to make a fine grained statistical analysis on several of our proposed metrics. This analysis includes a non asymptotic convergence analysis for our KDE based methods which formally unveils the elephant in the room that is the curse of dimensionality.

In section 4 we express how our novel framework actually coincides with the methods from the relevant literature on the extreme values. In particular, we detail how our approach restricted to these extreme values can be viewed as a dual approach of common metrics. In section 5 we first lead experiments on a controlled setting in which we can compare our proposed methods to the known ground truth. We then conduct experiments on real data and compare them to the results of articles from the same literature. Finally, in order to come to a conclusion concerning these “real world” experiments, we construct a hybrid experiment in which ground truth can effectively be computed and a qualitative conclusion can be made on our extended metrics.

## 2 A New Framework for Evaluating Generative Models

In this section, we formally present our novel framework for evaluating generative models. This framework stems from the dual approach proposed in Simon et al. (2019) related to binary classification. Before presenting this formulation, let us first introduce the notations.

### 2.1 Notations

- $\langle \cdot, \cdot \rangle$  is the standard scalar product on  $\mathbb{R}^d$  and  $\| \cdot \|$  the associated norm;

- $\Omega \subset \mathbb{R}^d$ : a measurable space;
- $\mathcal{P}(\Omega)$ : set of probability measures on  $\Omega$ ;
- $P, Q \in \mathcal{P}(\Omega)$  (e.g. real and generated distributions);
- $\mathcal{X}$  (resp.  $\mathcal{Y}$ ): finite training set of  $N$  samples from  $P$  (resp.  $Q$ ) ( $\#\mathcal{X} = \#\mathcal{Y} = N$ );
- $\mathcal{X}'$  (resp.  $\mathcal{Y}'$ ): finite testing set of  $N$  samples from  $P$  (resp.  $Q$ ) ( $\#\mathcal{X}' = \#\mathcal{Y}' = N$ );
- $\bar{P}, \bar{Q}$ : empirical distributions formed by the  $N$  training samples;
- $\hat{P}, \hat{Q}$ : estimates of the probability  $P$  and  $Q$ ;

We also recall the following definitions.

**Definition 1 (TV norm)** *The total variation (TV) norm of a signed measure  $\mu$  over  $\Omega$  is*

$$\|\mu\|_{\text{TV}} := \frac{1}{2} \sup_A (\mu(A) - \mu(A^c))$$

where  $A$  is a measurable set and  $A^c$  is the complement of  $A$ .  
The TV distance between  $P, Q \in \text{probaSet}$  is

$$\|P - Q\|_{\text{TV}} = \int_{dP \geq dQ} d(P - Q) = \sup_A (P(A) - Q(A)).$$

When the probability measures have densities w.r.t to the Lebesgue measure, i.e.,  $dP(x) = p(x)dx$  and  $dQ(x) = q(x)dx$ , the TV distance can be expressed as the half  $L^1(\Omega)$  norm,

$$\|p - q\|_{\text{TV}} = \frac{1}{2} \|p - q\|_1 = \frac{1}{2} \int_{\Omega} |p(x) - q(x)| dx.$$

**Definition 2 (Subgaussianity)** *A probability distribution  $P$  on  $\mathbb{R}^d$  is said to be subgaussian with constant  $K$  iff  $\forall \theta \in \mathbb{R}^d$*

$$\mathbb{E}_{X \sim P} [\exp(\langle \theta, X - \mathbb{E}[X] \rangle)] \leq \exp\left(\frac{1}{2} K^2 \|\theta\|^2\right). \quad (1)$$

We will say that  $X \sim P$  is a  $K$ -subgaussian random vector. Besides the smallest value of  $K^2$  for which  $X$  is  $K$ -subgaussian is denoted  $\text{Var}_G(X)$  and is called the subgaussian variance of  $X$  (with the convention  $\text{Var}_G(X) = +\infty$  iff  $X$  is not subgaussian).

**Definition 3** *The Sobolev space  $W^{1,1}$  on  $\Omega \subset \mathbb{R}^d$  is defined by:*

$$W^{1,1}(\Omega) := \{p \in L^1(\Omega) : \nabla p \in L^1(\Omega)\}$$

where  $\nabla p$  is a priori defined in the distributional sense. This space is equipped with the semi-norm

$$\|p\|_{W^{1,1}(\Omega)} := \int_{\Omega} \|\nabla p(x)\| dx.$$

## 2.2 PR Curve Starting Point

**The initial framework** As initially described in Sajjadi et al. (2018) the PR set between  $P$  and  $Q$  is defined as the set of non-negative pairs  $(\alpha, \beta)$  such that  $\exists \mu \in \mathcal{P}(\Omega)$  verifying both  $P \geq \beta\mu$  and  $Q \geq \alpha\mu$ . This set is denoted as  $\text{PRD}(P, Q)$ . In a nutshell, the two conditions translate the fact that the “probe” distribution  $\mu$  can simultaneously “extract” some mass  $\beta$  from  $P$  and  $\alpha$  from  $Q$ . Note that the PR set is included within  $[0, 1]^2$  and it is quite regular: it is topologically closed and also a lower set for the component-wise order. The latter means that  $\forall \alpha' \leq \alpha, \forall \beta' \leq \beta, (\alpha, \beta) \in \text{PRD}(P, Q) \implies (\alpha', \beta') \in \text{PRD}(P, Q)$ . As a result, this set is fully characterized by its (upper-right) Pareto frontier denoted by  $\partial\text{PRD}(P, Q)$  which can be parameterized as  $\partial\text{PRD}(P, Q) = \{(\alpha_\lambda, \beta_\lambda), \lambda \in \bar{\mathbb{R}}^+\}$  with

$$\begin{aligned} \alpha_\lambda &= (\lambda P \wedge Q)(\Omega) = \frac{\lambda + 1}{2} - \|\lambda P - Q\|_{\text{TV}} \\ \beta_\lambda &= \frac{\alpha_\lambda}{\lambda} \end{aligned} \tag{2}$$

where  $\wedge$  is the minimum operator between two measures. The link with TV can directly be adapted from (Verine, 2024, Lemma 4.4.1). Note that since Recall can be deduced from Precision as  $\beta_\lambda = \frac{\alpha_\lambda}{\lambda}$ , we will often solely focus on the Precision  $\alpha_\lambda$  value.

The information captured by this whole curve encompasses both extreme Precision and Recall values corresponding to  $\alpha_\infty$  and  $\beta_0$  which play a central role in the literature starting from Kynkäänniemi et al. (2019). In addition, it also describes how similarly the mass is distributed within the common support of  $P$  and  $Q$  (see (Siry et al., 2023) for details). This curve can also be characterized in a **dual form, based on a specific two-sample classification problem** as developed in Simon et al. (2019).

**Dual approach description of PR curves** This dual approach allows to express the Precision and Recall pairs  $(\alpha_\lambda, \beta_\lambda)$  for a given  $\lambda$  as a minimization problem over the hypothesis class  $\mathcal{F}$  which is composed of all binary classifiers on  $\Omega$ . The expression under minimization is a linear combination of the false negative rate (FNR) and the false positive rate (FPR) of the binary classifier  $f$ , that is to say the probability that a sample  $\mathcal{Y} \sim Q$  was classified as a sample from  $P$  for FNR (resp.  $\mathcal{X} \sim P$  classified as coming from  $Q$  for FPR). With these notations set once and for all, the dual approach gives:

$$\alpha_\lambda(P, Q) = \min_{f \in \mathcal{F}} \{\lambda \text{fpr}(f) + \text{fnr}(f)\} \tag{3}$$

where the false positive and negative rates write:

$$\text{fpr}(f) = \int (1 - f)dP \quad \text{and} \quad \text{fnr}(f) = \int f dQ. \tag{4}$$

## 2.3 A New Framework for Estimating $(\alpha_\lambda, \beta_\lambda)$

While (3) provides a rigorous alternative definition of Precision, it is intractable to estimate in practice as minimization occurs on the the entire set of binary classifiers  $\mathcal{F}$ . To tackle this issue, Simon et al. (2019) propose to approximate the optimal Bayes classifier  $f^*$  through the optimization over a parameterized family  $\mathcal{F}_\theta$  of deep neural networks. While this method does yield relevant Precision and Recall curves, it is computationally intensive

since it requires training one classifier per point on the curve. In the proposed approach, we propose to replace the neural network parameterization by a family of non-parametric classifiers, *e.g.* based on nearest-neighbors. For a specific classification method  $M$ , the corresponding family takes the form:

$$\begin{aligned}\mathcal{F}^M &= \{f_\gamma^M : \gamma \in [0, +\infty]\} \\ f_\gamma^M(z) &= \mathbf{1}_{\{\gamma c(\mathcal{X}, z) \geq c(\mathcal{Y}, z)\}},\end{aligned}\tag{5}$$

where  $c(\mathcal{X}, z)$  (resp.  $c(\mathcal{Y}, z)$ ) denotes a matching score between some sample  $z$  and samples from  $\mathcal{X}$  (resp from  $\mathcal{Y}$ ). The key idea is that a classifier will attribute a sample  $z$  to distribution  $P$  if it “matches” better the samples from  $P$  than it does those from  $Q$ . Concrete instances of scores corresponding to standard classification rules will be given in an upcoming section.

Adapting the dual form (3) to this new family of classifiers  $\mathcal{F}^M$ , we obtain the following estimators of the Precision and Recall values based on Empirical Risk Minimization:

$$\hat{\alpha}_\lambda^M := \min_{\gamma \in \mathbb{R}^+} \{\lambda \overline{\text{fpr}}(f_\gamma^M) + \overline{\text{fnr}}(f_\gamma^M)\},\tag{6}$$

where  $\overline{\text{fpr}}$  (resp.  $\overline{\text{fnr}}$ ) are the empirical estimates of the false positive (resp. negative) rate (4).

In addition to these empirical estimators, we introduce for the theoretical analysis conducted in the next part another estimator for the Precision  $\alpha_\lambda$ . This estimator, denoted as  $\hat{\alpha}_\lambda^{\text{TV}}$ , is related to the TV distance between estimated distributions:

**Definition 4 (TV based plug-in estimator of Precision)**

$$\hat{\alpha}_\lambda^{\text{TV}} := \alpha_\lambda(\hat{P}, \hat{Q}) = \frac{\lambda + 1}{2} - \|\lambda \hat{P} - \hat{Q}\|_{\text{TV}},\tag{7}$$

where  $\hat{P}, \hat{Q}$  are the estimated distributions based on samples  $\mathcal{X}, \mathcal{Y}$  and the selected method  $M$ . This estimator can be thought of as the “ground truth value” of Precision at  $\lambda$  for the estimated distributions  $\hat{P}, \hat{Q}$  and is in fact the plugin-in TV estimator of Precision at  $\lambda$ . While this estimator is not tractable in practice, it is useful in the statistical analysis that will be detailed in Section 3. This notation emphasizes the fact that here the Precision is estimated directly from the TV distance between estimated measures, which is different from the proposed classifier-based approach  $\hat{\alpha}_\lambda^M$  defined in (6).

## 2.4 Instantiating our Framework on kNN and KDE Families

In this section, we present effective families —based on the  $k$ -nearest neighbors method (kNN) and KDE— of classifiers  $\mathcal{F}^M$  and their advantages. This section will be complemented in Section 4 with some extensions to encompass other metrics from the literature.

**kNN family** An intuitive way of defining a binary classifier which aims at discriminating between  $P$  and  $Q$  from sample points is to define a majority vote: for a given test sample, it simply consists in comparing the number of closest samples from  $\mathcal{X}$  and from  $\mathcal{Y}$ . We introduce the following family of kNN binary classifiers

$$f_\gamma^{\text{kNN}}(z) = \mathbf{1}_{\gamma \#\{x \in \mathcal{X} / x \in B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}(z)\} \geq \#\{y \in \mathcal{Y} / y \in B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}(z)\}},$$

where the parameter  $\gamma$  is related to the ratio of samples originating from  $Q$ . For  $\gamma = 1$ , this binary classifier is an indicator function that outputs 1 when there are more samples from  $\mathcal{X}$  falling in the kNN radius ball of  $z$  than there are from  $\mathcal{Y}$ . This parameterized classifier will be discussed in depth in Sections 3.1 and 4.

**KDE family** Another family which naturally emerges when thinking of classification based on inter-sample distances is Kernel Density Estimation (KDE). We propose the simplest instance of KDE which uses fixed bandwidth. This method uses a uniform kernel and may be defined as follows:

$$f_\gamma^{\text{KDE}}(z) = \mathbb{1}_{\frac{\hat{p}(z)}{\hat{q}(z)} \geq \frac{1}{\gamma}}, \quad (8)$$

with estimated density  $\hat{p}(z) \propto \sum_{x \in \mathcal{X}} \mathbb{1}_{\|x-z\| \leq \sigma}$  (similarly for  $\hat{q}$ ) using a positive bandwidth parameter  $\sigma$ .

**Corresponding estimators** With the parametrized families of classifiers for kNN and KDE we can now define the corresponding Precision metrics as:

$$\begin{aligned} \hat{\alpha}_\lambda^{\text{kNN}}(P, Q) &= \min_{\gamma \in \mathbb{R}^+} \left\{ \lambda \text{fpr}(f_\gamma^{\text{kNN}}) + \text{fnr}(f_\gamma^{\text{kNN}}) \right\}, \\ \hat{\alpha}_\lambda^{\text{KDE}}(P, Q) &= \min_{\gamma \in \mathbb{R}^+} \left\{ \lambda \text{fpr}(f_\gamma^{\text{KDE}}) + \text{fnr}(f_\gamma^{\text{KDE}}) \right\}. \end{aligned}$$

## 2.5 Extending the Classification Framework

**Introduction of splitting** Based on the previous classification standpoint, it appears natural to split the dataset into two subsets: a “training set” ( $\mathcal{X} \cup \mathcal{Y}$ ) which is used to define the family of classifiers  $f_\gamma^M$  and a “test set” ( $\mathcal{X}' \cup \mathcal{Y}'$ ) used to evaluate the empirical error rates as defined above. In this setting, the conventional empirical estimation of FPR and FNR for a classifier  $f$  simply write:

$$\overline{\text{fpr}}(f) := \frac{1}{|\mathcal{X}'|} \sum_{x \in \mathcal{X}'} (1 - f(x)) \quad \text{and} \quad \overline{\text{fnr}}(f) := \frac{1}{|\mathcal{Y}'|} \sum_{y \in \mathcal{Y}'} f(y).$$

By splitting the dataset, the law of large numbers applies which is crucial for the consistency of  $\hat{\alpha}_\lambda^M$  which will be shown in Section 3.

Note however that because of splitting, it is possible that none of the classifier  $f_\gamma^M$  ensures a null FPR. As a result, it is possible that  $\hat{\alpha}_\lambda^M > 1$ . As a remedy, in our experiments, we always complement  $\mathcal{F}^M$  with the trivial classifiers  $f \equiv 1$  and  $f \equiv 0$  that predict either  $P$  or  $Q$  uniformly.

From now on in the splitting setting  $\mathcal{X} \cup \mathcal{Y}$  will denote the training set used to fit the classifiers and  $\mathcal{X}' \cup \mathcal{Y}'$  the set on which the error rates are evaluated.

**Hyper-parameter  $k$**  In addition to the introduction of splitting, we consider modifying the hyper-parameter  $k$  for the kNN balls approach. In the kNN literature (see e.g. Devroye et al. (2013)), it is known that as the number of samples  $N$  gets bigger,  $k$  must also increase but at a slower rate (this will be a key element to ensure the consistency of the kNN estimator in Theorem 5). We therefore consider for each approach, setting  $k \propto \sqrt{N}$ , which is a standard choice of parameters in kNN classification and allows to benefit from the consistency Theorem 5.

### 3 Theoretical Consistency Guarantees

In this section, we will first prove asymptotic consistency of our kNN-based estimator. Then we will turn to establishing non-asymptotic consistency results for our KDE-based estimator. The second result highlights the curse of dimensionality as the upper-bound depends exponentially on the dimension. These convergence analysis go way further than the ones in the relevant literature which includes examples where metrics are not even consistent.

#### 3.1 Asymptotic Consistency with kNN

We first prove asymptotic statistical consistency of Precision and Recall estimators when using our kNN method defined in Section 2.4. Once again, we will be focusing on  $\alpha_\lambda$  as the results can be deduced for  $\beta_\lambda$  using the relationship  $\beta_\lambda = \alpha_\lambda/\lambda$ .

**Theorem 5** *Let  $\lambda \in \bar{\mathbb{R}}^+$ ,  $k \geq 3$  and  $N = \#\mathcal{X} = \#\mathcal{Y}$ . If  $k \rightarrow \infty$  and  $\frac{k}{N} \rightarrow 0$  as  $N \rightarrow \infty$ , and denoting*

$$\Gamma_\lambda^* = \arg \min_{\gamma} \lim_{k \rightarrow \infty, \frac{k}{N} \rightarrow 0} \mathbb{E}[\lambda \text{fpr}(f_\gamma^{\text{kNN}}) + \text{fnr}(f_\gamma^{\text{kNN}})],$$

*then, the following hold:*

1.  $\lambda \in \Gamma_\lambda^*$ .
2.  $\mathbb{E}[\hat{\alpha}_\lambda^{\text{kNN}}] \rightarrow \alpha_\lambda$  assuming that data split was used.

The proof is provided in Appendix A and is similar to the standard Bayes consistency results of the kNN classifier (see, *e.g.* (Devroye et al., 2013, Chapters 5 and 6)). It is essentially adapted to the fact that the risk is class weighted *i.e.*  $R_\lambda(f) = \lambda \text{fpr}(f) + \text{fnr}(f)$  instead of the classical one  $R(f) = \frac{1}{2}(\text{fpr}(f) + \text{fnr}(f))$ .

In short, this theorem highlights that with increasing number of samples  $N$ , if the number of nearest neighbors  $k$  is chosen to grow large but more slowly than  $N$  so that it remains negligible in comparison, then our kNN-based estimator of Precision tends towards the ground-truth Precision in expectation.

#### 3.2 Non-Asymptotic Consistency with KDE: TV counterpart

In this section, we derive an intermediate result about the optimal bound associated to the *risk* of the estimated Precision based on the KDE. More precisely, we consider the equivalent formulation (2) of the Precision  $\alpha_\lambda(P, Q)$  and the corresponding plugin-in estimator (7) based on the KDE of bandwidth  $\sigma$ .

For simplicity, we use the same Parzen-window kernel  $k_\sigma(z) \propto \mathbb{1}\{\|z\| \leq \sigma\}$  to estimate both densities, and we consider the estimator  $\hat{P} := \bar{P} * k_\sigma$  (respectively  $\hat{Q} := \bar{Q} * k_\sigma$ ), where  $\bar{P}$  is the empirical distribution of  $N$  *i.i.d.* samples  $\mathcal{X}$  drawn from  $P$  (resp.  $\bar{Q}$  is the empirical distribution of  $N$  samples  $\mathcal{Y}$  from  $Q$ ).

From the expression of  $\alpha_\lambda$  in (2) and using standard statistical arguments, it is well-known that if one wants to establish precise non-asymptotic convergence rates, one has to

trade on the generality w.r.t to the class of distributions. Indeed, any estimator of  $\alpha_\lambda(P, Q)$  can be trivially turned into an estimator of  $\|\lambda P - Q\|_{\text{TV}}$  and vice versa. In that sense, the estimation difficulty of both functionals is the same. In fact, the difficulty of estimating  $\|P - Q\|_{\text{TV}}$  (i.e  $\lambda = 1$ ) is not much studied in the literature but one can conjecture that it is related to the difficulty of estimating  $P$  and  $Q$  under the TV norm (more precise statements will be made later on in the paper). It is important to know that in Devroye and Györfi (1990), the authors exhibit a scenario in which for any number of samples  $N$ , and any estimation scheme, the estimation of a worst-case distribution  $P$  is  $1/2$  away from the actual target distribution in terms of TV norm. Such scenario relies on the potential existence of a singular component in the distribution (i.e. a component that is neither discrete nor absolutely continuous). Therefore in order to obtain any non-asymptotic consistency result, one usually requires some regularity and tail assumptions on the target distribution. In our case, we will assume Sobolev-type regularity and subgaussianity (see Definition 2).

**Theorem 6** *Assume that both  $P$  and  $Q$  are  $K$ -subgaussian and have densities w.r.t. Lebesgue measure that belong to the Sobolev ball of radius  $R$ . Set  $\sigma \propto \frac{1}{N^{\frac{1}{2+d}}}$ . Then*

$$\mathbb{E}|\hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda(P, Q)| = O\left(\frac{\lambda + 1}{N^{\frac{1}{2+d}}}\right), \quad (9)$$

where the constant in the big  $O$  notation depends only on  $K$ ,  $R$  and the dimension  $d$  (not on  $P$  and  $Q$  nor  $\lambda$  and  $N$ ).

Denoting  $\mathcal{M}_{R,K}$  the set of probability distributions pairs  $(P, Q)$  that are both  $K$ -subgaussian and whose densities are in the Sobolev ball of radius  $R$ , we obtain the minimax upper-bound

$$\inf_{\hat{\alpha}_\lambda} \sup_{(P,Q) \in \mathcal{M}_{R,K}} \mathbb{E}_{P^{\otimes N} \times Q^{\otimes N}} |\hat{\alpha}_\lambda - \alpha_\lambda(P, Q)| = O\left(\frac{\lambda + 1}{N^{\frac{1}{2+d}}}\right). \quad (10)$$

In the remainder of this subsection, we will prove the theorem through a series of intermediate lemmas.

The following lemma highlights how the estimation error made on the Precision and Recall metrics can be transposed to a TV error estimation made on the distributions under study.

**Lemma 7**

$$\mathbb{E} [|\hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda(P, Q)|] \leq \lambda \mathbb{E} [\|\hat{P} - P\|_{\text{TV}}] + \mathbb{E} [\|\hat{Q} - Q\|_{\text{TV}}]. \quad (11)$$

**Proof** Indeed

$$\begin{aligned} \mathbb{E} [|\hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda(P, Q)|] &= \mathbb{E} \left[ \left| \|\lambda \hat{P} - \hat{Q}\|_{\text{TV}} - \|\lambda P - Q\|_{\text{TV}} \right| \right] \\ &\leq \mathbb{E} \left[ \|\lambda \hat{P} - \hat{Q} - (\lambda P - Q)\|_{\text{TV}} \right] \quad (\text{reverse triangle inequality}) \\ &\leq \lambda \mathbb{E} [\|\hat{P} - P\|_{\text{TV}}] + \mathbb{E} [\|\hat{Q} - Q\|_{\text{TV}}] \quad (\text{direct triangle inequality}). \end{aligned}$$

■

Because of the previous bound, one can now consider the TV risk of the KDE estimator of  $P$  (or  $Q$ ). To do so, it is standard to use the bias-deviation decomposition which follows from the triangular inequality.

**Lemma 8 (Bias-Deviation decomposition)**

$$\mathbb{E} \left[ \|\hat{P} - P\|_{\text{TV}} \right] \leq \underbrace{\mathbb{E} \left[ \|\hat{P} - \mathbb{E}[\hat{P}]\|_{\text{TV}} \right]}_{\text{deviation}} + \underbrace{\mathbb{E} \left[ \|\mathbb{E}[\hat{P}] - P\|_{\text{TV}} \right]}_{\text{bias}}. \quad (12)$$

Besides, for the considered estimator  $\hat{P} := \bar{P} * k_\sigma$ , we have

$$\mathbb{E}[\hat{P}] = P * k_\sigma. \quad (13)$$

**Lemma 9 (Deviation bound)** For  $P$  a  $K$ -subgaussian probability distribution defined in  $\mathbb{R}^d$ ,  $\bar{P}$  an empirical distribution based on  $N$  samples of  $P$  and a Parzen-window kernel  $k_\sigma$  of width  $\sigma > 0$  we have

$$\mathbb{E}[\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}] \leq \sqrt{\Gamma \left( \frac{d}{2} + 1 \right)} 8^{\frac{d}{4}} \left( \frac{\eta}{d} + \frac{K^2}{\sigma^2} \right)^{\frac{d}{4}} \frac{1}{2\sqrt{N}}, \quad (14)$$

where  $\eta$  is a universal constant independent from  $d$  (see Proposition 21 for details) and  $\Gamma$  is Euler's gamma function.

See B.3 for the proof.

**Remark 10** Lemma 9 is stated for the Parzen-window kernel. Nonetheless the proof can be easily extended to any kernel  $k$  obeying the following growth condition:

$$\exists c > 0, \forall z \in \mathbb{R}^d, k^2(z) \leq c\tilde{k}(z),$$

where  $\tilde{k}$  is non-negative subgaussian kernel (i.e. corresponds to a density that is subgaussian). In particular, this condition does not require the original kernel  $k$  to be non-negative, which is an important flexibility in order to control the bias under high-order smoothness conditions.

In what follows, we consider probability distributions having densities which lie in a Sobolev space  $W^{1,1}(\mathbb{R}^d)$ .

**Lemma 11 (Bias bound)** Let  $R > 0$  and assume that  $P$  has a density  $p \in W^{1,1}(\mathbb{R}^d)$  w.r.t. Lebesgue measure such that  $\|p\|_{W^{1,1}(\mathbb{R}^d)} \leq R$ . Then

$$\|P - P * k_\sigma\|_{\text{TV}} \leq \frac{R}{2}\sigma.$$

See B.2 for the proof.

**Remark 12** *Again this lemma can be extended to more general kernels. For first-order smoothness, any kernel  $k$  (not necessarily non negative) is fine, as long as  $\int \|z\| |k(z)| dz < +\infty$  (this quantity would of course appear in the upper bound). For higher order smoothness  $s$ , the previous condition is replaced by  $\int \|z\|^s |k(z)| dz < +\infty$  along with assumptions of vanishing signed moments up to order  $s-1$ . The proof is again similar, but rely on a higher order Taylor expansion. In this context, the bias term is controlled as  $O(\sigma^s)$ . This kind of result is standard in the approximation folklore and is mentioned merely for completeness.*

**Proof** [of Theorem 6] First, combining Lemma 9 and Lemma 11, and optimizing for  $\sigma$  in the resulting bound, the best bias-deviation trade of is achieved for the devised choice  $\sigma \propto \frac{1}{N^{\frac{1}{2s+d}}}$ . Plugging this into Lemma 8 and then using Lemma 7 yields the first claim of Theorem 6. The second claim follows from the first one as the left hand side of (10) is obviously upper bounded by the one of (9).  $\blacksquare$

**Remark 13** *The result of Theorem 6 is consistent with known results on density estimation under  $L^1$  or TV norm. Roughly speaking, for smoothness order  $s$  (e.g in terms of Sobolev spaces or other regularity spaces), the minimax lower bound rate is known to be  $\frac{C}{N^{\frac{s}{2s+d}}}$  for  $C > 0$ ; see Tsybakov (2009). Following the discussion in Remark 12, to obtain an estimator for larger smoothness orders, it is well known that the Parzen-Window (constant) kernel is not sufficient. Instead one needs to use more regular kernels with vanishing signed moments up to order  $s-1$ , which implies that the kernel cannot be non-negative anymore.*

*That being said, one should realize from the proof of Lemma 7 that what really matters for our purpose is consistent estimation of  $\alpha_\lambda(P, Q)$ , or equivalently, of  $\|\lambda P - Q\|_{\text{TV}}$ . Even for the simple case of  $\lambda = 1$ , less is known about the consistency rate of  $\|P - Q\|_{\text{TV}}$  except that it is necessarily faster than the rate for density estimation under TV norm. This distinction between the convergence rate for an estimation under a functional vs the estimation of the functional itself is classical (see e.g. (Lepski et al., 1999)). Usually the rates differ by a slow correcting factor (e.g. logarithmic or even sub-logarithmic in  $N$ ).*

*On the other hand, the dependence we have obtained for the prefactor w.r.t the value of  $\lambda$  is  $\lambda + 1$ . This stems from the use of the triangular inequality. This approach is quite crude, since it would imply an infinite pre-factor when  $\lambda \rightarrow \infty$ , while the consistency bound is necessarily lower than  $\frac{1}{2}$  for all  $\lambda \in [0, +\infty]$ . Indeed, we know that the estimation target  $\alpha_\lambda \in [0, 1]$ , so that trivial estimator  $\hat{\alpha}_\lambda = \frac{1}{2}$  display the prescribed error of  $\frac{1}{2}$  (see the discussion on the onset of Section 3.2).*

*Intuitively, if one considers a common support hypothesis, then the consistency error is zero for  $\lambda \rightarrow \infty$  since we know for sure that  $\alpha_\infty(P, Q) = 1$  in that case. So one should lean towards thinking of  $\lambda = 1$  as the most difficult configuration. That being said, despite being widespread in density estimation, the common support assumption is not well adapted to a setting where one desires to estimate  $\alpha_\infty$ . Indeed one would not seek to estimate a value that carries no information (as is the case for  $\alpha_\infty$  under the common support assumption). Studying the consistency of estimating  $\alpha_\lambda(P, Q)$  under more general settings and its dependence on  $\lambda$  is therefore a very interesting avenue that we leave to a future work.*

Another interesting direction of future research is proving that the rate of Theorem 6 is (nearly) optimal. We are currently working towards this direction by establishing a minimax lower bound that would reveal that the rate of Theorem 6 is optimal up to a logarithmic factor.

### 3.3 Bias of the KDE method

In the previous section we obtained an optimal bound on the risk associated to the estimated Precision using the plug-in estimator  $\hat{\alpha}_\lambda^{\text{TV}}$  based on the TV distance using KDE. However, this analysis does not account for the practical way the Precision is estimated, based the evaluation of a classifier as proposed in (6), rather than a distance between densities. Furthermore, the TV norm  $\|\lambda\hat{P} - \hat{Q}\|_{\text{TV}}$  is intractable in practice. In this section, we show that we can still use this result to directly derive a bound on the bias associated to the KDE classifier (8) used to estimate  $\hat{\alpha}_\lambda^{\text{KDE}}$ .

Recall first that the proposed approach is based on a “data splitting” strategy introduced in 2 where the classifier is evaluated on a different set than the training one. Consequently, let us denote by  $(\mathcal{X}, \mathcal{Y})$  (resp.  $(\mathcal{X}', \mathcal{Y}')$ ) random and independent samples from the distributions  $(P, Q)$  which designate the training set (resp. the evaluation set). Namely,  $\mathcal{X}, \mathcal{X}' \sim P^{\otimes N}$  and  $\mathcal{Y}, \mathcal{Y}' \sim Q^{\otimes N}$ . The distributions  $\hat{P}, \hat{Q}$  are the estimated distributions based on the empirical samples  $\mathcal{X}, \mathcal{Y}$ .

#### Proposition 14 (Bound on the bias of the KDE-classifier estimator)

$$|\mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} - \alpha_\lambda]| \leq 2\lambda\mathbb{E}\|\hat{P} - P\|_{\text{TV}} + 2\mathbb{E}\|\hat{Q} - Q\|_{\text{TV}}, \quad (15)$$

where once again  $\hat{P} := k_\sigma * \bar{P}$  and  $\hat{Q} := k_\sigma * \bar{Q}$  are the KDE estimators of  $P$  and  $Q$ .

See B.4 for the proof.

**Remark 15** We have seen in Theorem 6, that the plug-in estimator displays a fairly good convergence rate for its estimation risk (if the bandwidth  $\sigma$  is set appropriately), namely  $O\left(\frac{\lambda+1}{N^{\frac{1}{2+d}}}\right)$ . The same goes for the classifier KDE estimate (which can be implemented in practice), but in that case, we only control the bias. To get a guarantee in terms of risk, one would need to also control the mean average deviation for this estimator, and show that it is  $O\left(\frac{\lambda+1}{N^{\frac{1}{2+d}}}\right)$ . This result is more involved due to the minimization over  $\gamma$ .

## 4 A Bridge Towards the Existing Literature

In this section we establish a link between the proposed framework described in Section 2 and scalar Precision and Recall metrics from the state of the art. We show that these metrics can be expressed as the extreme values of Precision and Recall curves when carefully selecting a parameterized family  $\mathcal{F}_M$ .

#### 4.1 Previous Work Under our Setting

In this section, we formally describe a few published metrics related to extreme Precision and Recall ( $\alpha_\infty$  and  $\beta_0$ ). Once again, we will focus on the estimate of  $\alpha_\infty(P, Q)$  because swapping the role of  $P$  and  $Q$  entails the extreme Recall  $\beta_0(P, Q)$ .

Recall that we consider a finite set  $\mathcal{X}$  of examples from  $P$  and others in  $\mathcal{Y}$  sampled from  $Q$  with  $\#\mathcal{X} = \#\mathcal{Y} = N$ . In the literature, no splitting occurs that is to say  $\mathcal{X}' = \mathcal{X}$  and  $\mathcal{Y}' = \mathcal{Y}$ .

**IPR** Proposed in Kynkäänniemi et al. (2019), the Improved Precision and Recall metric is given by

$$\hat{\alpha}_\infty^{iPR} := \frac{1}{N} \sum_{y \in \mathcal{Y}} \mathbb{1}_{\exists x \in \mathcal{X}, y \in B_{\text{kNN}}^{\mathcal{X}}(x)}, \quad (16)$$

where  $\mathcal{X}$  and  $\mathcal{Y}$  are the observed samples from  $P$  and  $Q$  respectively, and  $B_{\text{kNN}}^{\mathcal{X}}(x)$  represents the kNN ball around  $x$  computed within the set  $\mathcal{X}$ . This value can be interpreted as the empirical estimate of  $Q(\text{supp}(P))$  where samples from  $\mathcal{Y}$  are used to estimate the  $Q$  probability, and those from  $\mathcal{X}$  are used to estimate the support of  $P$  as the union of kNN balls.

**Coverage** First proposed as an estimate of  $\beta_0$  in Naeem et al. (2020), it was also adapted to  $\alpha_\infty$  in Khayatkhoei and AbdAlmageed (2023).

$$\hat{\alpha}_\infty^{cov} := \frac{1}{N} \sum_{y \in \mathcal{Y}} \mathbb{1}_{\exists x \in \mathcal{X}, x \in B_{\text{kNN}}^{\mathcal{Y}}(y)}. \quad (17)$$

Note that compared to Equation (16), the condition  $y \in B_{\text{kNN}}^{\mathcal{X}}(x)$  is merely replaced by  $x \in B_{\text{kNN}}^{\mathcal{Y}}(y)$ . This metric was proposed to leverage the sensitivity to outliers of **IPR**.

**EAS** Khayatkhoei and AbdAlmageed (2023) propose to combine both previous estimates by taking their minimum

$$\hat{\alpha}_\infty^{eas} := \min(\hat{\alpha}_\infty^{iPR}, \hat{\alpha}_\infty^{cov}). \quad (18)$$

**PRC** Proposed in Cheema and Uner (2023) Precision and Recall Cover is an extension of coverage:

$$\hat{\alpha}_\infty^{PRC} = \frac{1}{N} \sum_{y \in \mathcal{Y}} \mathbb{1}_{\#\{x \in \mathcal{X} / x \in B_{\text{kNN}}^{\mathcal{Y}}(y)\} \geq k'}, \quad (19)$$

where  $k' \in \mathbb{N}^*$  is an additional hyper-parameter: setting  $k' = 1$  makes this estimator identical to  $\alpha_\infty^{cov}$ .

**PPR** Last, Probabilistic Precision and Recall was proposed in Park and Kim (2023):

$$\hat{\alpha}_\infty^{PPR} := \frac{1}{N} \sum_{y \in \mathcal{Y}} \left( 1 - \prod_{x \in \mathcal{X}} (1 - \tau(\|y - x\|)) \right), \quad (20)$$

where  $\tau(d) = \max(0, 1 - \frac{d}{R})$  is a fixed bandwidth tent kernel and  $R$  a hyper-parameter defining the width of the kernel. This method allows to take into account the stochasticity of the samples by accounting for an uncertainty on the borders of the estimated support.

## 4.2 Classification Interpretation

One may notice that every estimators mentioned in the previous section reads as

$$\hat{\alpha}_\infty^M = \overline{\text{fnr}}(f_\infty^M),$$

where  $\overline{\text{fnr}}$  is the empirical FNR,  $M$  denotes the selected method (e.g. **iPR**, **Cov**) and  $f_\infty^M$  is a classifier specific to the approach which takes as an input any  $z \in \Omega$ . In particular, one has  $f_\infty^{iPR}(z) = \mathbb{1}_{\#\{x \in \mathcal{X}/z \in B_{\text{kNN}}^\mathcal{X}(x)\} \geq 1}$  and  $f_\infty^{\text{cov}}(z) = \mathbb{1}_{\#\{x \in \mathcal{X}/x \in B_{\text{kNN}}^\mathcal{Y}(z)\} \geq 1}$ .

In most cases (all except PRC when  $k' > 1$ ), by design  $\overline{\text{fpr}}(f_\infty^M) = 0$  because the classifier is equal to 1 on training samples from  $\mathcal{X}$  —which are the same used to evaluate the error rates. This is reminiscent of the form of  $\alpha_\lambda$  in Equation 3 when  $\lambda \rightarrow \infty$ :

$$\alpha_\infty = \min_{f \in \mathcal{F} \text{ s.t. } \text{fpr}(f)=0} \text{fnr}(f).$$

In fact, one can analyze each approach  $M$ , as the instantiation over a carefully selected family of classifiers  $\mathcal{F}^M$  as defined in Definition 5.

Note that there is no unique way of defining the class of functions  $\mathcal{F}^M$  from the extreme classifiers  $f_\infty^M$  and  $f_0^M$ . Yet, we shall now see that natural families emerge for both **iPR** and **Cov**.

**IPR** In that case, we set

$$f_\gamma^{iPR}(z) = \mathbb{1}_{\gamma \#\{x \in \mathcal{X}/z \in B_{\text{kNN}}^\mathcal{X}(x)\} \geq \#\{y \in \mathcal{Y}/z \in B_{\text{kNN}}^\mathcal{Y}(y)\}}.$$

Note that this classifier is a Kernel Density Estimator (KDE) of the form

$$f_\gamma^{iPR}(z) = \mathbb{1}_{\frac{\hat{p}(z)}{\hat{q}(z)} \geq \frac{1}{\gamma}}$$

where  $\hat{p}(z) \propto \sum_{x \in \mathcal{X}} \mathbb{1}_{B_{\text{kNN}}^\mathcal{X}(x)}(z)$  and similarly for  $\hat{q}(z)$ . It is therefore a KDE classifier with adaptive bandwidth.

**Coverage** Here we set

$$f_\gamma^{\text{cov}}(z) = \mathbb{1}_{\gamma \#\{x \in \mathcal{X}/x \in B_{\text{kNN}}^\mathcal{Y}(z)\} \geq \#\{y \in \mathcal{Y}/y \in B_{\text{kNN}}^\mathcal{X}(z)\}}.$$

This resembles to a classical kNN classifier up to a minor difference: it is more standard to use the same kNN structure for both classes, that is  $B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}$  rather than using a separate one per class. Interestingly, one can verify that the condition  $\exists x \in \mathcal{X}$  s.t.  $x \in B_{\text{kNN}}^\mathcal{Y}(y)$  is in fact equivalent<sup>1</sup> to  $\exists x \in \mathcal{X}$  s.t.  $x \in B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}(y)$ . In other words, without split our **kNN** method matches the extreme values of Coverage.

**Remark 16 (On symmetry)** *For symmetry reasons between Precision and Recall, we use the previous definitions of  $f_\gamma^M$  for  $\gamma \geq 1$  and favor a strict inequality over a loose one for  $\gamma < 1$ .*

1. One direction is trivial since  $B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}(y) \subset B_{\text{kNN}}^\mathcal{Y}(y)$ , the other requires a bit more reasoning: assuming  $\exists x \in B_{\text{kNN}}^\mathcal{Y}(y)$  one may consider in particular the  $x$  closest to  $y$  and conclude that it belongs to  $B_{\text{kNN}}^{\mathcal{X} \cup \mathcal{Y}}(y)$ .

**Remark 17 (Fixed vs adaptive bandwidth)** *Adaptive bandwidth KDE estimators based on  $k$ NN distances—such as those used in **iPR**—may seem like a natural choice. However, in such settings, outliers can significantly inflate the local bandwidth, distorting the estimator’s output by attributing excessive support to regions that should lie outside the true distribution. Using a fixed bandwidth as in our **KDE** estimator (see Section 2.4) helps in leveraging this issue.*

**Remark 18 (Splitting)** *While the metrics from the literature were defined out of the splitting context, extending them is straightforward. This is done by computing the empirical error rates on “test set”  $\mathcal{X}' \cup \mathcal{Y}'$  different from the “training set”  $(\mathcal{X} \cup \mathcal{Y})$  used to define the classifiers.*

### 4.3 The Correct Version of Extreme Precision

In the literature, the accepted expression of the extreme Precision is

$$\alpha_\infty(P, Q) := \lim_{\lambda \rightarrow \infty} \alpha_\lambda(P, Q) = Q(\text{supp}(P)). \quad (21)$$

While this formula is used as a theoretical basis in the literature, it is flimsy and requires to be amended mainly because the support of a distribution is defined up to null sets for that distribution. In the incriminated identity, the issue stems from the fact that adding a  $P$ -null set can change the  $Q$ -mass of the set, and therefore the right-hand side is not well characterized. Correcting identity (21) requires to clarify a few notions beforehand.

**Definition 19 (support and co-support)** *Let  $A$  be a measurable subset of  $\Omega$  and  $A^c$  its complement. We say that  $A$  is a*

- *support of  $P$ , denoted<sup>2</sup>  $A = \text{supp}(P)$ , iff  $P(A^c) = 0$ .*
- *co-support of  $P$  and  $Q$  denoted  $A = \text{cosupp}(P, Q)$  iff  $((P \wedge Q)(A^c) = 0$  and  $\forall B \subset A, P(B) = 0 \Leftrightarrow Q(B) = 0$ ).*

As the reader may notice, the second notion is characterized up to sets that are simultaneously  $P$  and  $Q$  null. More precisely, we have the following result.

**Proposition 20** *All co-supports of  $P$  and  $Q$  have the same  $Q$ -mass and*

$$\alpha_\infty(P, Q) = Q(\text{cosupp}(P, Q)).$$

**Proof** See proof in Appendix C. ■

**Example 1** *The following example illustrates the difference between the non-corrected and the corrected variants. Let  $P = N(0, \text{Id})$  and let  $Q$  be a singular distribution, e.g.,  $\delta_x$  for some  $x \in \mathbb{R}^d$ . Then the entire PR curve is trivial and  $\alpha_\infty(P, Q) = 0$ . However,  $\text{supp}(P) = \mathbb{R}^d$ , so  $Q(\text{supp}(P)) = 1$ . In contrast,  $\text{cosupp}(P, Q) = \emptyset$ , hence  $Q(\text{cosupp}(P, Q)) = 0$ .*

---

2. This is a slight abuse of notations.

A natural follow-up question is whether either variant (based on support or co-support) provides a meaningful Precision metric. We believe the answer is negative for both. The issue is that these quantities tend to saturate at 1 even in situations where Precision should be considered low. Beyond the classical example of shifted Gaussians (which yields perfect Precision regardless of the shift), the example above—and slight variations of it—illustrate this limitation.

Consider the following interpretation. Let  $P$  be a white noise distribution over images, and let  $Q$  be a Dirac distribution concentrated on a single image (e.g., a dog). A meaningful Precision metric should reflect whether samples from  $Q$  resemble typical samples from  $P$ . Clearly, a dog image is not representative of white noise, so assigning a Precision score of 1 would be incongruous.

Moreover, this issue persists (albeit to a lesser extent) with the corrected formulation. If  $Q$  is instead a narrow Gaussian centered at the dog image (i.e., the dog image slightly perturbed by a zero-mean Gaussian noise), both  $Q(\text{supp}(P))$  and  $Q(\text{cosupp}(P, Q))$  still evaluate to 1, despite the lack of alignment with the notion of Precision. For this reason, we argue that extreme metrics such as  $\alpha_\infty$  are not suitable estimation targets. Consequently, studying their convergence rates is of limited interest. Instead, we advocate evaluating full PR curves and, if needed, summarizing them using appropriate statistics.

More specifically, the previously mentioned caveats reflect the influence of the tails of  $P$  and  $Q$  on the extreme metrics, even when those tails decay very fast. The previous example of white-noise ( $P = \mathcal{N}(0, \text{Id})$ ) vs noisy-dog images ( $Q = \mathcal{N}(x, \sigma \text{Id})$ ) is illuminating. Whatever the noise scale  $\sigma$  (be it extremely small), it remains that  $P$  and  $Q$  have full co-support and therefore  $\alpha_\infty(P, Q) = 1$ . The first negative impact of this observation is that both  $\alpha_\infty$  and  $\beta_0$  provide a very weak characterization of the relation between  $P$  and  $Q$ . The second negative impact concerns the estimation of  $\alpha_\infty$ : namely, the tails of  $P$  and  $Q$  are elusive based on empirical samples, making this extreme Precision the most challenging to evaluate. Both of these observations have gone unnoticed by the previous approaches starting from Kynkäänniemi et al. (2019) as they purposely focused on the extreme values. In addition to those two arguments, we would like to highlight that estimating the mass of a support is not a standard topic in machine learning. As a result, taking the binary classification dual standpoint brings much more useful hindsight to design estimators correctly.

#### 4.4 Distilling Curves with Scalar Metrics

While we have just highlighted the importance using Precision and Recall curves over scalar Precision and Recall metrics, practitioners may find it useful to summarize the PR curve using two metrics that respectively reflect Precision and Recall, thereby trading exhaustiveness for conciseness.

This can be particularly helpful for simplifying model comparisons. In this context, we previously discussed how focusing solely on extreme Precision or Recall values is often far from ideal. We therefore discuss four alternative approaches and examine their behavior in a scenario combining mode dropping, mode invention, and mode re-weighting (e.g., Figure 3a).

An empirical evaluation of the advantages and limitations of all four alternatives would be a valuable future direction. For example, one could study how these metrics behave

with respect to hyperparameters in state-of-the-art generative models, such as the truncation trick in GANs, the temperature in auto-regressive model and normalizing flows, or the guidance scale factor in diffusion models. Such experiments are presented in Sections 5.2 and 5.4.

**AuC** Area under the curve (AuC) is a standard metric which synthesizes a curve by simply outputting its area. This metric is particularly used in binary classification. Such a metric is proposed in Pillutla et al. (2021) where the MAUVE metric is the area under the curve which is defined by type I and type II errors on generated and real distributions. In our case the AuC simply denotes as  $\int \alpha_\lambda(\beta_\lambda)d\beta_\lambda$ . The values of AuC range from 0 to 1. An AuC of 0 means that the PR curve boils down to the point (0, 0) and is synonymous of the two distributions under study having disjoint supports. On the contrary, AuC=1 means that the two distributions under study are the same. This score behaves as an averaging metric which tends to soften extreme behaviors for example giving the same score to curves with very large Recall and very small Precision and conversely meaning it fails to distinguish high Recall and high Precision because a symmetrized version of a curve would give exactly the same AuC value. However, AuC complements the extreme values as it synthesizes what happens between those two extreme values.

**F-scores** Proposed by Sajjadi et al. (2018), the  $F_b$  score is defined as:

$$F_b = \max_{\lambda \in [0, +\infty]} \frac{b^2 + 1}{\frac{b^2}{\alpha_\lambda} + \frac{1}{\beta_\lambda}}.$$

As  $b \rightarrow \infty$ ,  $F_b \nearrow \alpha_\infty$ , and conversely, as  $b \rightarrow 0$ ,  $F_b \searrow \beta_0$ . Based on this, Sajjadi et al. (2018) proposed using both  $F_b$  and  $F_{1/b}$  with a large value of  $b$  (specifically,  $b = 8$ ). Although not explicitly discussed in their original work, it can be understood that these metrics will not be sensitive to rapidly decaying infinite tails. However, note that the  $F_b$  score is a weighted harmonic mean computed from a single (optimal) point on the PR curve, meaning it can remain unchanged even when the curve is modified elsewhere. As a result, when  $b$  is large,  $F_b$  and  $F_{1/b}$  mainly capture mode dropping or invention but are only superficially informative about mode re-weighting.

**PR median** Another alternative is to consider  $(\alpha_{\bar{\lambda}}, \beta_{\bar{\lambda}})$ , where  $\bar{\lambda}$  is chosen such that the line  $\alpha = \bar{\lambda}\beta$  divides the area under the PR curve into two equal parts. In cases of pure mode dropping or invention, we have  $\alpha_{\bar{\lambda}} = \alpha_\infty$  and  $\beta_{\bar{\lambda}} = \beta_0$ . By contrast, when infinite tails induce sharp transitions to  $\alpha_\infty = 1$  (as seen in Figure 3a),  $\alpha_{\bar{\lambda}} < 1$  will be only mildly affected by the fast-decaying tails, similar to  $F_b$ . Unlike  $F_b$ , however,  $\alpha_{\bar{\lambda}}$  is significantly influenced by the presence of transitions due to mode re-weighting (purple transition in Figure 3a). As a result, these metrics may be preferable when one wishes to account for both mode re-weighting and mode dropping/invention.

**PR@ $\epsilon$**  When  $P$  and  $Q$  share the same support, the extreme values will always equal 1, providing little useful information about the distributions under study. However, in the case of Gaussian distributions, as the means of  $P$  and  $Q$  diverge, the PR curve tends toward the point (0, 0). Capturing the behavior “just before” reaching these extremes can provide more detailed insights. This can be done by computing Precision at a fixed value

of Recall, denoted  $\alpha@{\epsilon}$ , or vice versa. In our case, we set  $\epsilon = 0.05$ . This idea is inspired by recommender systems or information retrieval evaluation, where one might measure how many of the top- $k$  selected items are relevant. In short,  $\alpha@{\epsilon}$  (or equivalently  $\beta@{\epsilon}$ ) refers to the maximum  $\alpha$  such that  $(\alpha, \epsilon) \in \text{PRD}(P, Q)$  if such a point exists and 0 otherwise.

## 5 Experiments

In this section, we first revisit several experiments on toy datasets previously proposed in the literature where the ground truth is available and results are easy to interpret (Section 5.1). We then apply our metrics on data generated by StyleGAN-V2 model reproducing a key experiment from Kynkäänniemi et al. (2019) and computing PR curves based on embeddings (Section 5.2). However, analyzing experiments conducted solely on real-world data is inherently challenging due to the absence of clear ground truth, making it difficult to rigorously evaluate the performance of the metrics. To address this, we introduce a hybrid experiment: starting from real samples, we fit Gaussian distributions and generate synthetic samples following the corresponding Gaussian laws. This hybrid setup provides a controlled reference against which we can meaningfully compare our metrics (Section 5.3). We believe that some encouraging results in the literature are the effect of a combination of negative bias—due to correlation as train and test samples are the same—which could be compensated by the positive bias due to the sub-optimal selected classifier. These experiments showcase the pros and cons of the various extensions, yet without declaring a clear winner.

In short, in the toy experiments, **Cov** and **kNN** perform far better than **iPR** which fails to recover the middle part of the ground truth curve. In the StyleGAN experiment, **iPR** exhibits a very different behavior than **Cov** and **kNN** which in any event is hard to analyze. Finally, the proposed hybrid setting illustrates the advantages of each method in different situations. Additional experiments are presented in Appendix D including the specific case  $P = Q$  advocating the use of data splitting.

In all the experiments, we denote  $n$  as the total number of samples per distribution. Therefore, considering the previous notations,  $n = 2N$  when splitting occurs otherwise  $n = N$ .

### 5.1 Toy Examples

In all toy experiments, we use  $n = 10\text{K}$  samples in  $\mathbb{R}^{64}$  for the various estimators under study, matching the experimental settings of previous works. In each case, the distributions  $P$  and  $Q$  are known analytically, and the ground truth PR curve can be easily estimated since the Bayes classifier is the likelihood ratio classifier  $f_{\lambda}^*(z) = \mathbb{1}_{\frac{d_Q}{d_P}(z) \leq \lambda}$  (Simon et al., 2019). To achieve high-accuracy ground truth curves, we apply a large-sample Monte Carlo estimation ( $n^{GT} = 100\text{K}$ ) and estimate:

$$\hat{\alpha}_{\lambda}^{GT} = \lambda \overline{\text{fpr}}(f_{\lambda}^*) + \overline{\text{fnr}}(f_{\lambda}^*).$$

Based on this PR curve, we can either visually assess the quality of an estimator or use a scalar indicator to summarize performance. In particular, we propose using the Intersection over Union (IoU) score between the ground truth curve and the estimator under review.

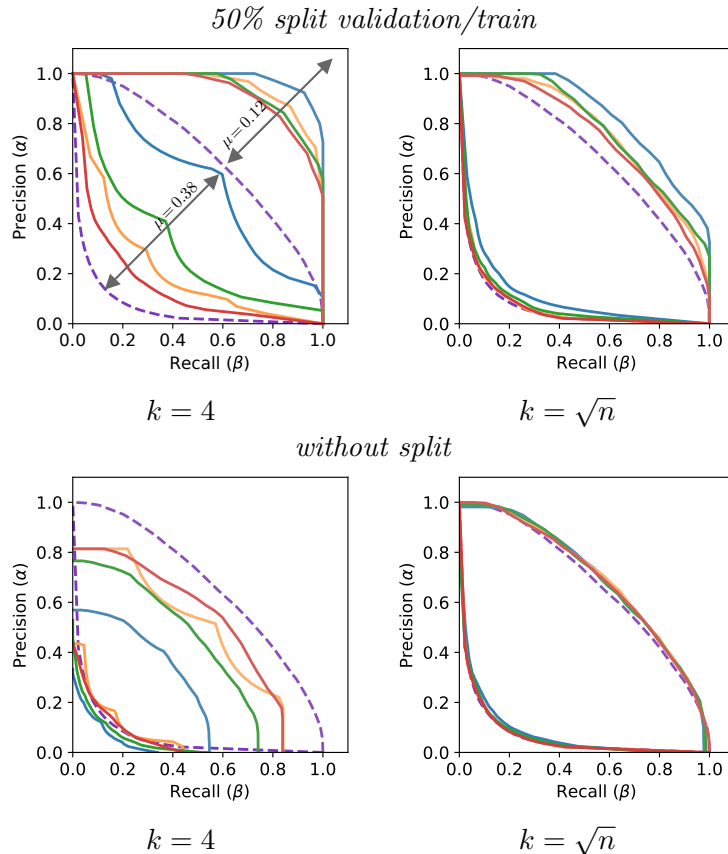


Figure 2: **Comparing two shifted Gaussians.** The ground truth PR curve ( - -GT) is compared to empirical estimates from various NN-classifiers: -iPR, -kNN, -KDE, and -Cov. Here  $P \sim \mathcal{N}(0, \mathbb{I}_d)$  and  $Q \sim \mathcal{N}(\mu \mathbf{1}_d, \mathbb{I}_d)$  with  $d = 64$  dimensions and  $\mu = \frac{1}{\sqrt{d}} \approx .12$  or  $\mu = \frac{3}{\sqrt{d}} \approx .38$ .  $n = 10\text{K}$  points are sampled using  $k = 4$  or  $k = \sqrt{n}$  for NN comparison, with or without dataset validation/train split. (Each curve is obtained by averaging 10 PR curves from different sets of random samples.)

This index ranges from 0 to 1, with values closer to 1 indicating stronger similarity to the ground truth.

### 5.1.1 SHIFTED GAUSSIAN

Inspired by experiments in related works, we consider the case where  $P$  and  $Q$  are two Gaussian distributions with increasing shifts. Although we sample a large number of points, we run each experiment 10 times to obtain robust interpretations of the metrics. Unlike Kynkäänniemi et al. (2019); Naeem et al. (2020); Park and Kim (2023), we once again emphasize that the extreme values of the curves should always equal 1, as the two Gaussian distributions have a full co-support.

We evaluate four different methods across four different shifts of the fake Gaussian. The resulting curves are presented for two shifts in Figure. 2 and summarize the full experiment results and the corresponding IoU scores in Appendix D Table 1. In these settings, the standard deviation across experiments is below  $10^{-2}$  on IoU scores. We vary the nearest neighbor parameter  $k \in \{4, \sqrt{10K}\}$  and we test our metrics with and without split through the experiment. As described in Remark 18, we can apply the splitting framework to the methods from the literature which initially didn’t encompass it. In this split setting, for each distribution  $P, Q$  with  $n$  sampled points, half of them are selected for defining the parameterized family of functions and the other half is used to compute the error rates which yield the Precision and Recall curve.


As it was already highlighted in Naeem et al. (2020), we can observe how iPR noticeably underestimates the ground truth in the no-split and  $k=4$  setting (it’s initial definition by the original authors). In some settings, the PR curves differ significantly from the ground truth. In prior works, the number of nearest neighbors  $k$  used for manifold estimation was set to 3 (Kynkäänniemi et al., 2019) or 5 (Naeem et al., 2020). Here, we first use  $k = 4$  and then, motivated by Th. 5, set  $k = \sqrt{n}$ . Interestingly, we observe in the former case that, as expected, estimated PR curves are more pessimistic in the split setting, meaning they underestimate the similarity between the two distributions compared to the validation/train split. This effect is so pronounced that the extreme values of the curves do not reach one, as they theoretically should. We observe that results are more robust with  $k = \sqrt{n}$ , and that differences between split and no-split scenarios become marginal.

Given the theoretical guarantees in the split case, we retain the combination of split and  $k = \sqrt{n}$  for the subsequent toy experiment (additional results are provided in the Appendix D.1).


### 5.1.2 GAUSSIAN MIXTURE MODELS

As argued by Luzi et al. (2023), neural representations—such as Inception features—are better approximated by Gaussian mixture models (GMMs) than by pure Gaussians as stated in the original paper introducing FID. This more general modeling also allows us to illustrate key phenomena: mode *dropping* (a mode present in  $P$  but not in  $Q$ ), mode *invention* (a mode present in  $Q$  but not in  $P$ ), and mode *re-weighting* (shared modes between  $P$  and  $Q$  with different weights), as shown in Figure 3a. The corresponding theoretical PR curves are depicted in Figure 3b. The distributions  $P$  and  $Q$  used for experiments are here defined as a mixture of 4 modes:  $P = \sum_{\ell} p_{\ell} \mathcal{N}(\mu_{\ell} \mathbf{1}_d, \mathbb{I}_d)$  and  $Q = \sum_{\ell} q_{\ell} \mathcal{N}(\mu_{\ell} \mathbf{1}_d, \mathbb{I}_d)$ , with  $d = 64$  and  $\mu_{\ell} \in \{0, -5, 3, 5\}$ . The weights differ between  $P$  and  $Q$ :  $p_{\ell} \in \{0.2, 0.2, 0.6, 0\}$  and  $q_{\ell} \in \{0, 0.5, 0.1, 0.4\}$ . In this synthetic experiment, we sample points from two GMMs and apply splitting with  $k = \sqrt{n}$  (see Figure 3c).


The **kNN**, **KDE**, and **Cov** methods perform well relative to the ground truth, while **iPR** tends to overestimate the PR curve, particularly failing to capture the re-weighting transitions. Transitions to the extreme values are noticeably not well estimated by the **KDE** approach. Indeed, the **KDE** classifier, with non-adapting kernel width, is less discriminative on Gaussian tail densities, while being more robust than **iPR**. The tails effects are amplified for **iPR** because low density samples produce large kNN balls. Additional experiments described in D.2 confirm these observations.



(a) Gaussian Mixture Models  $P$  and  $Q$  showing mode dropping (only in  $P$ ), mode inventing (only  $Q$ ), and mode re-weighting (in both but distributed differently).



(b) Expected coarse shape of the PR-curve (solid black). Due to the infinite tails of the Gaussian modes, the vertical and horizontal transitions are theoretically smooth and reach 1 (dashed curve).



(c) **Comparing two Gaussian mixtures.** The ground truth PR curve ( - -GT) is compared to empirical estimates from various NN-classifiers:  $-iPR$ ,  $-kNN$ ,  $-KDE$ , and  $-Cov$ . Notice that the method **Cov** and **kNN** overlap (and are very close to the ground truth).

Figure 3: Comparison of Gaussian Mixture Models.

### 5.1.3 VARIABILITY STUDY

As the access to a large number of samples might be a bottleneck, the question of the variability of the metrics’ output might be risen. In Figure 4, the impact of the size sample  $n$  on the variability of the evaluation curves is shown. These figures illustrate the consistency of the proposed method based on robust classifiers when increasing the number of samples (Theorem 5). Empirically, using  $10K$  points reduces sufficiently the variability to make comparison between curves reliable, which is in line with the standard usage for generative model evaluation (Heusel et al., 2017; Sajjadi et al., 2018).

Figure 4 complements Section 5.1.3 about variability. Average curves are obtained by computing the empirical mean of  $N = 100$  PR curves obtained different random  $n$ -samples (with  $n = 10^4$ ), *i.e.*

$$(\bar{\alpha}(\lambda), \bar{\beta}(\lambda)) = \frac{1}{n} \sum_{i=1}^n (\hat{\alpha}_i(\lambda), \hat{\beta}_i(\lambda)),$$

where  $(\hat{\alpha}_i(\lambda), \hat{\beta}_i(\lambda))$  are the independent estimates (associated with the separate runs). Deviation from average curves are materialized with two curves

$$(\alpha_{\pm\sigma}(\lambda), \beta_{\pm\sigma}(\lambda)) = (\bar{\alpha}(\lambda) \pm \sigma_{\alpha}(\lambda), \bar{\beta}(\lambda) \pm \sigma_{\beta}(\lambda)),$$

with empirical estimator

$$\sigma_{\alpha}(\lambda)^2 = \frac{1}{n} \sum_{i=1}^n (\hat{\alpha}_i(\lambda) - \bar{\alpha}(\lambda))^2.$$

and similarly for the Recall deviation  $\sigma_{\beta}$ . **NB:** By construction, for a fixed  $\lambda$ ,  $\hat{\alpha}_i(\lambda) = \lambda \hat{\beta}_i(\lambda)$ . This ratio is automatically inherited for the average curves  $(\bar{\alpha}(\lambda), \bar{\beta}(\lambda))$  and the standard deviation curves  $(\alpha_{\pm\sigma}(\lambda), \beta_{\pm\sigma}(\lambda))$ . Whereas the methods’ Precision and Recall have different biases, their deviation are very similar in their behavior as shown in Figure 4.

## 5.2 Experiment with StyleGAN

In this section, we test our metrics in the scenario of a generative model trained on some “real-world” data. To do so, we use the same approach as in Kynkäänniemi et al. (2019) where the model StyleGAN<sup>3</sup> is trained on the Flickr Face High-Quality dataset (FFHQ).

**Image Embedding** In the previous experiments on synthetic data, we noticed how high dimension made the estimation of Precision and Recall metrics challenging. However, images from the FFHQ dataset live in very high dimension (*i.e.*  $d \approx 3.10^6$ ). As it is standard in the literature, image samples are compared via neural representations from pre-trained deep networks. In addition to mitigating the curse of dimensionality by mapping images to a new space of smaller dimension, it allows to define relevant metrics that can be aligned with human perception (such as LPIPS). In this experimental section, we make use of several popular feature embedding (supervised or non-supervised): InceptionV3 (Szegedy et al., 2016) ( $d = 2048$ ), VGG16 (Simonyan and Zisserman, 2015) ( $d = 4096$ ), VGG16 with random weights as proposed in Naeem et al. (2020), and DINOv2 (Oquab et al., 2024) ( $d = 384$ ).

3. In our case, we use StyleGAN v2 implementation (Karras et al., 2020).




Figure 4: **Influence of sample size  $n$ .** The setting is the same as Figure 2 for a translation of  $\mu = .21$  between two Gaussian in dimension  $d = 64$  (with splitting and  $k = \sqrt{n}$ ). Solid (respectively transparent) curves correspond to the empirical average (resp. deviations) of 100 PR curves computed from random samples. In this experiment, we use splitting with a factor 0.5.

**Filtering of generated samples *via* truncation** Recall that StyleGAN is a latent generative model built upon a specific feed-forward neural network. Starting from a constant input, convolutional layers are modulated by a latent variable that is obtained in two stages: a Gaussian random variable  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}_d, \mathbb{I}_d)$  is first sampled (Z space) that is then encoded using a MLP mapping (W space):  $\mathbf{w} = M(\mathbf{z})$ . Note that in StyleGAN-V1 (Karras et al., 2019) additional random noise is also employed for modulation. In Kynkäänniemi et al. (2019), several truncation methods parametrized by a scalar  $\Psi$  are investigated to improve automatically the quality of generated samples. The truncation parameter  $\Psi$  indicates how close to the average the latent variables of the model are linearly remapped after being sampled from a normal distribution. The most efficient one in terms of a trade-off between

simplicity and efficiency consists in using an interpolation to the mean:  $\mathbf{w} \mapsto \Psi(\mathbf{w} - \bar{\mathbf{w}}) + \bar{\mathbf{w}}$ , where  $\bar{\mathbf{w}} = \mathbb{E}M(\mathbf{z})$ . In the Improved Precision and Recall paper, the authors’ intuition was that as  $\Psi \rightarrow 0$ , the generated distribution collapses to a single point (the point with highest probability density) thus always generating—slight variations of—the same image but which is the most probable. Conversely, when  $\Psi \rightarrow 1$ , the model is less constrained and its latent variables are varied, yielding images with high diversity yet with more artifacts.

We reproduced the results of the truncation experiment on StyleGAN from Kynkäänniemi et al. (2019) using our metrics. As described in Section 4.4, practitioners might want to have Precision and Recall curves summarized to scalar metrics (Area Under Curve or AuC,  $F_b$ ,  $\alpha@0.05$ ). We computed those synthesized metrics including the extreme value  $\alpha_\infty$  as done in the StyleGAN experiment of Kynkäänniemi et al. (2019)). These metrics are shown in Figures 5, 6 and 7 as a function of the truncation parameter  $\Psi$ , for respectively **iPR**, **Cov** and **kNN**.

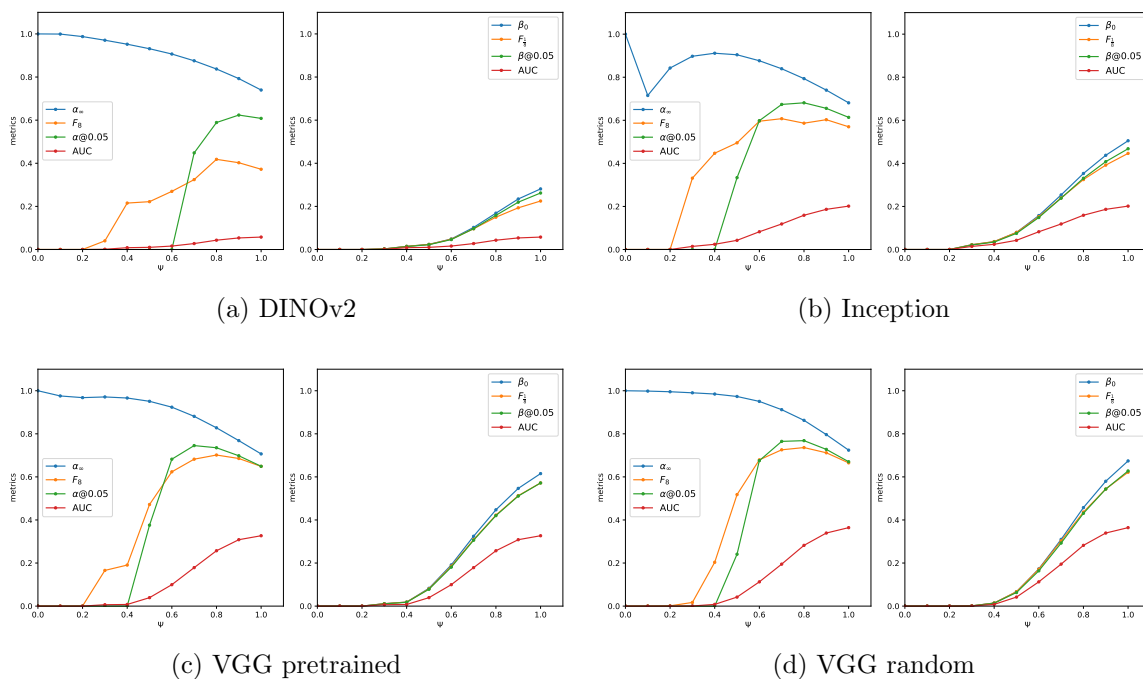


Figure 5: Truncation experiment with the **iPR** metric, using the same parameters as in the original article.

Regarding the extreme Precision metric  $\alpha_\infty$  for the **iPR** method, as expected we obtain the same global behavior (mostly monotonic decreasing) as in Kynkäänniemi et al. (2019). Notice that the only slight difference is for the Inception embedding, which is probably due to the fact that we use the deterministic StyleGAN-V2 which makes generated samples collapse to a single image when  $\Psi \rightarrow 0$ , making estimation of Precision more challenging.

The comparison of **Cov** with **kNN** and **iPR** highlights a different behavior for  $\alpha_\infty$  when  $\Psi \rightarrow 0$  which we argue to be the desired one, as done previously by Djolonga et al. (2020). In this setting, we find ourselves comparing the real distribution  $P$ , to the distribution




Figure 6: Truncation experiment with the **Coverage** metric, using the same parameters as in the original article.




Figure 7: Truncation experiment with the **kNN** method (split=0.5,  $k=\sqrt{n}$ ).

$Q_\Psi$ , which collapses to a Dirac. Using the cosupport concept introduced in Definition. 19, the extreme Precision writes  $\alpha_\infty(P, Q) = Q(\text{cosupp}(P, Q))$ . In this case  $\text{cosupp}(P, Q) = \emptyset$  yielding  $\alpha_\infty(P, Q) = 0$ , which is not the value yielded by **iPR**. On the other hand, the other metrics all yield  $\alpha_\infty = 0$  when  $\Psi = 0$  which indeed is the expected behavior.

Note that other metrics (AuC,  $F_b$  and  $\alpha@{\varepsilon}$ ) share the same behavior across all embeddings (DINOv2, VGG, Inception), and all methods, including **iPR**.

**Impact of the embedding model** Interestingly, a randomly initialized VGG embedding yields very similar result to what is obtained with DINOv2 or VGG networks trained on ImageNet, as previously reported by Naeem et al. (2020). In the mean time, we might question the relevance of using models pretrained on simple and cropped datasets such as ImageNet to evaluate models on complex datasets. Using a random model seems relevant as it does not induce a bias from any given dataset. To go further, we might wonder how many output dimensions are necessary to capture all the relevant information from the images. Consistently removing a large number of dimensions from the output of a randomly initialized neural network might help yielding relevant results while not suffering from the curse of dimensionality.

### 5.3 Hybrid Experiment: Sampling Gaussians from StyleGAN Statistics

In this section we present a hybrid experiment standing between the “real” StyleGAN truncation experiment and the “toy experiment” where we made two gaussian distributions evolve with respect to one another.

Due to the fact that the DINOv2 model produces embeddings in lower dimension than the other models—thus making computations less intensive—we will use DINOv2 embeddings for the remainder of this experiment. In this experimental setting, we select the 50 000 samples embeddings that come from the FFHQ dataset and are ran through our embedding model. We will project all the features in the same subspace which is yielded by the first  $d$  principal components of FFHQ embeddings. Our experiment focuses on varying the number of selected dimensions  $d$ . For DINOv2 dimension  $d$  lies in  $[1, 384]$ .

For each given truncation value of  $\Psi \in \{0.0, 0.1, \dots, 1.0\}$ , we select the embeddings of the 50 000 samples generated on StyleGAN (discussed in 5.2), trained on FFHQ and truncated with parameter  $\Psi$ . We project the given embeddings in the same  $d$  dimension subspace determined by FFHQ. In this setting, we therefore first reduce the complexity of the problem and transpose it in a controlled setting where we associate each distribution to a Gaussian distribution determined by the mean and covariance of the projected embeddings. This experimental setting makes it possible to compare the PR curves we compute to the ground truth which is estimated by a Monte Carlo simulation applied on the Bayes likelihood ratio classifier—which is explicitly known.

For each generated curve which is defined by a subspace dimension  $d$ , a truncation value  $\Psi$  and a method  $M$ , we compare the method’s curve to the ground truth estimation using the intersection over union (IoU) score. Given a value of  $\Psi$ , we propose to denote  $G_\Psi^d$  the multivariate and non isotropic Gaussian distribution yielded by the projected embeddings in dimension  $d$ . Similarly,  $G_{\text{FFHQ}}^d$  describes the pending FFHQ distribution in dimension  $d$ .

We propose three different experimental settings, detailed in the following table, each representing a particular scenario designed to reveal how certain approaches are favored or disadvantaged under specific conditions.

Description	Section	$P^d$	$Q^d$	Favorable Bias
FFHQ vs truncated distribution	5.3.1	$G_{FFHQ}^d$	$G_{\Psi}^d$	underestimation
Distribution vs itself	5.3.2	$G_{\Psi}^d$	$G_{\Psi}^d$	overestimation
Differently weighted GMMs	5.3.3	$GMM_{\Psi, \Psi'}$	$GMM_{\Psi, \Psi''}$	unknown

### 5.3.1 COMPARING PROJECTED FFHQ TO PROJECTED TRUNCATED DISTRIBUTION

In this setting, we transpose exactly the previous StyleGAN experiment to a controlled setting where, for a given method we compare  $G_{\Psi}^d$  to  $G_{FFHQ}^d$ . As we now know, in high dimension, the curves are more complex to estimate. Besides, because Gaussians in high dimensions tend to be singular with respect to one another, the ground truth curve tends towards 0. Therefore this experiment **favors methods that tend to underestimate the PR curve** as it is the case for Kynkäänniemi et al. (2019).

Similarly to the previous workflows we presented, we compute the IoU score between the generated curves and the estimated ground truth and average the various computed IoU over different  $\Psi$  values.

The results of this experiment are shown in Figure 8. We can observe how indeed the dimension of the ambient space makes it in practice hard to estimate Precision and Recall, especially for the extreme values. As we select more and more eigenvalues, the IoU score effectively decreases until reaching approximately 0. For example if we focus on the **kNN with split** method we introduced earlier on, we see that when selecting the 2 largest eigenvalues, the IoU score is around 0.9 and this value decreases to as low as 0.3 when selecting the largest 50 eigenvalues. **iPR** turns out to be a clear winner in this setting as it is known to underestimate the ground truth PR.




Figure 8: Evolution of IoU score between ground truth and various Precision and Recall metrics based on Gaussian samples of DINOv2.

### 5.3.2 COMPARING A TRUNCATED AND PROJECTED DISTRIBUTION TO ITSELF

In this setting, for each truncation  $\Psi$  value, the generated distribution  $G_{\Psi}^d$  is compared to itself. We know that in this setting, **iPR** tends to considerably underestimate the PR curve which should be the entire going through (1, 1). This experiment should leave an advantage to methods that tend to overestimate the true PR curve.

The results of this experiment are shown in Figure 9. We indeed observe that **iPR** decreases very fast and that the other methods give more appropriate results. In this setting, we illustrate the consistency of the kNN based estimators, **kNN** and **Cov**. Interestingly, in this setting, **kNN** without split performs globally better than coverage (note that the number  $k$  of neighbors differs between these two approaches so that their estimates are not strictly identical even for extreme values  $\lambda = 0$  or  $\lambda = \infty$ ).

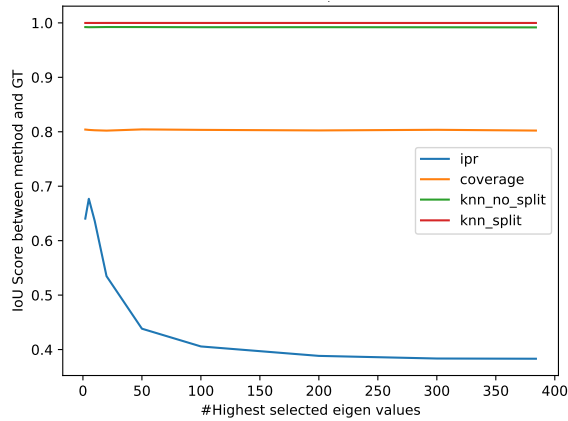


Figure 9: Evolution of IoU score when comparing a distribution to itself with various Precision and Recall metrics.

### 5.3.3 TRUNCATED AND PROJECTED GAUSSIAN MIXTURE

In this experiment, we turn back to the GMM experiment described in Section 5.1.2. Now, we consider the case where  $P^d \propto G_{\Psi=0.3}^d + G_{\Psi=0.5}^d$  and  $Q^d \propto G_{\Psi=0.3}^d + G_{\Psi=0.9}^d$  are two Gaussian mixtures sharing a mode. As the modes have all equal weight, the theoretical PR curve is approximately be a square going through (0.5, 0.5). This setting neither favors overestimation nor underestimation.

The results of this experiment are shown in Figure 10 where kNN based methods have very similar results and have an IoU score that tends very fast towards 0. This time the comparison between **kNN** without split and **Cov** is reversed as the latter now has the upper hand by a small margin. On the other hand, the **iPR** method even though having a kink at the start gives a constant and reasonable IoU score with the ground truth curve.

As the two distributions are Gaussian Mixtures, the extreme values should be equal to 1. Yet as the dimension increases, the two distributions become singular with respect to one another and the ground truth curve tends towards a square passing through (0.5, 0.5).




Figure 10: Evolution of IoU score between ground truth and various Precision and Recall metrics in the Gaussian mixture of StyleGAN truncated embeddings setting.

#### 5.4 Experiments with a conditional diffusion model

This section is dedicated to experiments with a state-of-the-art generative flow-matching model. We use Stable Diffusion (Esser et al., 2024), a conditional text-to-image generative model based on rectified flows. More specifically, we use the public model `SD-3.5-medium` (with approximately three billion trainable parameters plus CLIP encoders for text conditioning) from the HuggingFace repository (Stability AI, 2025). During inference, for a given text prompt, an ordinary differential equation solver is used to synthesize iteratively a new image based on the velocity field predicted by the trained model conditioned on the text prompt. To improve the image quality and prompt adherence, this diffusion process may be enhanced by extrapolating the estimated velocity using the classifier-free guidance technique of Ho and Salimans (2021). This effect is controlled by the so-called “guidance scale” factor.

In this context, we experiment two different configurations which allow to showcase the proposed metrics for a conditional generative model, when either changing the text conditioning (§ 5.4.1) or the guidance scale (§ 5.4.2), by comparing samples from a generic vs specific guidance. To do so, in both case, two datasets of 40k images are created at  $512 \times 512$  pixels resolution with 25 inference steps. Similarly to Section 5.2, we then use DINOv2 (Oquab et al., 2024) to extract embeddings from the images.

##### 5.4.1 GENERIC VS. SPECIFIC PROMPTING

In this setting, we first generate a reference dataset with the generic prompt “butterfly” corresponding to distribution  $P$ , and a second one (distribution  $Q$ ) from the more specific prompt “blue butterfly”. This setting allows us to evaluate how restricting the prompt affects diversity (Recall). In order to affect the Precision, we also fixed the guidance scale to 7.5 (instead of the default value of 5) in order to amplify the adherence to the prompt, in this case the “blue” aspect of the generated image.

The experimental performance curves associated with this first experiment are presented in Figure 11. As the prompt of distribution  $Q$  describes a subset of distribution  $P$ , we expect

in this setting a drop in Recall. Now, because of the high guidance during synthesis, we may also expect an impact on Precision. All the curves behave accordingly, but in noticeable different ways. In accordance with the previous experiments on StyleGAN, the **kNN** curve tends to overestimate extreme Precision and Recall (estimated with split datasets), but with a sharp transition along the axes, also suggesting a drop of Precision and Recall. The effect of drop in Precision and Recall is highlighted by the curves summary  $F_8 = 0.44$  and  $F_{1/8} = 0.52$ . In particular the **iPR** curve shows a significant drop in extreme Recall to  $\beta_0^{\text{iPR}} = 0.21$  and a moderate for Precision to  $\alpha_\infty^{\text{iPR}} = 0.65$  while, conversely, the **Cov** curve exhibits a more important drop of extreme Precision than Recall ( $\alpha_\infty^{\text{Cov}} = 0.14$  vs  $\beta_0^{\text{Cov}} = 0.42$ ).


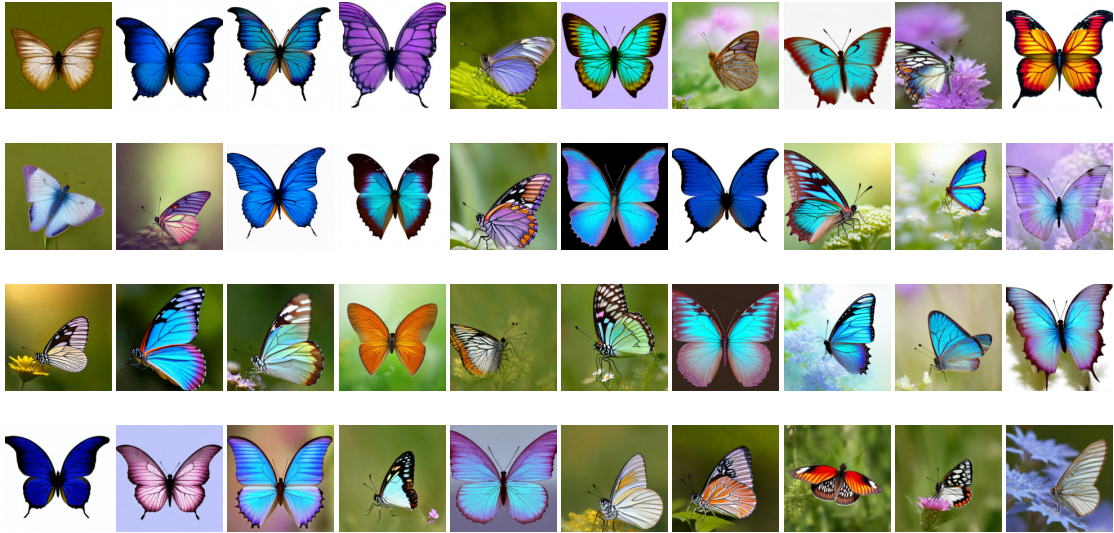



Figure 11: A comparison of the Precision and Recall metrics estimated with various classifiers: **iPR**, **kNN** and **Cov**.

It is important to underline that, even in this controlled setting, we do not have access to the ground-truth to assess which method is the most relevant. As a tentative to do so, it is interesting to visually inspect the images, some of which are showed in Figure 12, with the following observations. As expected, the variability of the “**blue butterfly**” images is much less in terms of color (even the blue colors of butterflies fall within a very narrow portion of the hue space compared to the “**butterfly**” dataset) and pose compared to those of the general class “**butterfly**” thus affecting the Recall. However, blue butterflies in the general class are mostly displayed on a white background, while samples from the restricted dataset seems to appear with a more diverse background (such as blue color gradient) which may also explain the simultaneous drop in Precision. To illustrate this observation, Figure 13 displayed the first 40 samples of blue butterflies from the unrestricted “**butterfly**” dataset,

selected manually. Of course, this analysis strongly depends on the feature extractor ability to retain information about color, pose and background.



(a) Samples from the “butterfly” dataset.



(b) Samples from the “blue butterfly” dataset.

Figure 12: Random images from each generated datasets, with guidance scale set to 7.5.

#### 5.4.2 VARYING GUIDANCE SCALE

In this setting, we generate several datasets using different guidance scales and analyze their impact on the resulting Precision and Recall curves. The guidance scale parameter can be




Figure 13: First 40 images of blue butterflies from the unrestricted “butterfly” dataset, selected by a human.

interpreted as an inverse temperature: a higher guidance scale reduces image variability while increasing adherence to the text prompt.

In the following experiment, 40k images are generated for each dataset, using guidance scales ranging from 1 to 15, with the same prompt : “a beautiful landscape with mountains and rivers”. Note that in order to make sure the samples are independent, a different random seed is used for the generation of each dataset. Random samples are shown in Figure 14.

All datasets are then compared with the reference dataset corresponding to a standard guidance scale of 5. The summary metrics ( $AuC$ ,  $\alpha_\infty$ ,  $\alpha@_\epsilon$ ,  $F_b$ ,  $\beta_0$ ,  $\beta@_\epsilon$ ,  $F_{1/b}$ ) of the corresponding PR curves are displayed in Figure 15.

As expected, when spanning guidance factors, metrics of Precision and Recall increase until a maximum (1 in in theory) when they reach the reference guidance scale of 5, and then decrease. When the guidance scale is low, the drop in Precision and Recall is due to a lower adherence to the prompt, resulting here in a larger diversity of scenery and a lower image quality. For larger guidance scale, the decrease in Precision and Recall is likely due to an excessive match of the prompt (*e.g.* resulting in images with a more vibrant color palette).

Interestingly, this experiment also highlights a limitation of the **iPR** method first reported in (Naeem et al., 2020, Section 4.1.1): when comparing two identical Gaussian distributions, **iPR** strongly underestimates the extreme Precision and Recall (*i.e.*  $\alpha_\infty^{\mathbf{iPR}} = 0.66$  and  $\beta_0^{\mathbf{iPR}} = 0.67$  instead of the perfect score 1). In Section 5.1.1, while we focused the discussion on the core idea of evaluating how Precision and Recall curves are affected by two Gaussian distributions shifting away from one another, we also had observed how extreme Precision and Recall were severely underestimated with the **iPR** method when the shift was absent. This behavior was then observed in the StyleGAN Truncation experiment in




Figure 14: Generated images of “a beautiful landscape with mountains and rivers” for varying guidance scales (indicated on the left).

section 5.3.2, where Gaussian densities were fitted on the dataset of generated StyleGAN images, with an IoU of 0.66 instead of 1. Here in this setting, for a guidance scale of 5, the two random datasets are directly drawn from the same non-Gaussian distribution, and yet we observe again the same phenomenon with  $\alpha_{\infty}^{\text{iPR}} = 0.73$  and  $\beta_0^{\text{iPR}} = 0.74$  instead of 1. Accordingly to previous experiments, **Cov** tends to give a better estimate in this setting ( $\alpha_{\infty}^{\text{Cov}} = 0.97$  and  $\beta_0^{\text{Cov}} = 0.97$ ), while **kNN** with data-splitting (which tends to overestimate extreme values) provides perfect estimates ( $\alpha_{\infty}^{\text{kNN}} = 1 = \beta_0^{\text{kNN}}$ ).

### 5.5 Summary of the experiments

These experiments show in turn that different methods have an advantage over each other in some specific settings. Because of the lack of information in the general open setting, we prefer comparing metrics in controlled environments such as 5.1 or 5.3 in which we can estimate the ground truth and compare our estimators to it. It turns out that the extension of **iPR**, even though not theoretically established, display some good empirical features in high dimension. While this observation diverges with the current literature, it should be understood as referring to the overall shape of the extended curve. In contrast, prior studies focused exclusively on the extreme value, due to inherent limitations in the PR metrics under investigation.

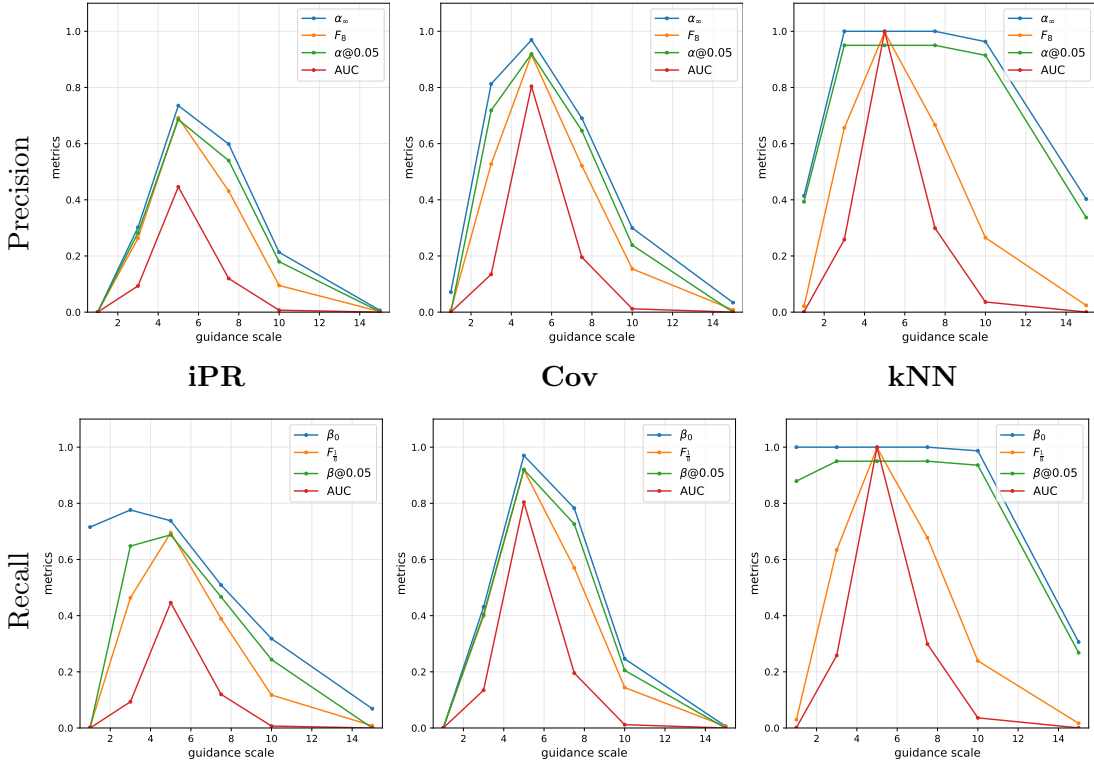


Figure 15: Experiment with generated images using the same conditioning prompt (“a beautiful landscape with mountains and rivers”) with different guidance scales (1, 2.5, 5, 7.5, 10, 15). Generated datasets are compared with a reference dataset using a fixed guidance scale of 5. Summarizing metrics for Precision and Recall curves w.r.t to the guidance scale. The embeddings are based on DINOv2.

## 6 Conclusion

In this work, we have described a fresh perspective on the evaluation of generative models by alleviating the links between PR curves, the binary classification standpoint and total variation distance. We have proposed a new framework for evaluating generative models that describes an entire curve, rather than being limited to scalar metrics, as is often done in the recent literature. This classification standpoint opened new practical perspectives (splitting, choice of hyperparameter  $k$ ). Benefiting from the splitting setting, we statistically analyzed the estimators of PR we introduced. This analysis allowed to highlight the crucial role played by the dimension in the upper bounds of the estimation errors. We then showed how to bridge common scalar PR metrics from the literature to our framework, thus stemming an entire curve from initially 2 scalar metrics. This extension can be viewed as a dual approach to those common metrics as the extreme values of our extended curves match rigorously the scalar metrics. Finally, we have studied the empirical behavior of the obtained variants in the light of several toy datasets and real world experiments with complex datasets. We consolidated the experiments by introducing a hybrid setting between

the two previous kinds of dataset which allowed to display a complex case in which the ground truth metrics are known.

Our main messages are the following. First, computing non extreme PR values is crucial because of essential issues in the extreme values which are related to their sensitivity to the distribution tails. Then, the curves themselves allow to describe more finely how the masses of the two distributions under comparison differ on their modes. This is useful in practice in order to tackle the case where a model generates data from the target support but with re-weighted masses. We highlighted the negative impact of dimension in the quality of our estimators. Yet, as the extreme values of our extensions match some of the methods we extended, this approximation upper bound also holds for the initial metrics. On the experimental side, there is no method clearly winning over the others as different behaviors emerge in varying scenarios (high dimension, several modes in the distributions ...) If employing a data split is theoretically appealing and allowed our statistical analyses, its empirical impact is less marked since the negative bias resulting from the lack of split can sometimes advantageously compensate the positive bias caused by the restricted hypothesis class. However, this benefit is not consistent over all experiments.

This work offers several interesting perspectives and could benefit from some extensions. A first step would be to focus on an asymptotic analysis for kNN methods in order to elucidate the convergence rate of the estimator. In addition, the proposed KDE asymptotic analysis derived an upper bound for the estimation error which yielded a minimax upper bound. An extension of this analysis could be developed to establish a minimax lower bound.

## Appendix A. Proof of kNN consistency

**Proof** To establish the proof, we need only show that  $R_\lambda(f_\lambda^{kNN}) \rightarrow \alpha_\lambda$  as  $k \rightarrow \infty$  and  $\frac{k}{n} \rightarrow 0$ . This will effectively imply both items in the theorem since  $\alpha_\lambda$  is the associated Bayes risk (Simon et al., 2019). To establish this limit, the first step is to show that for fixed  $k$  then  $\lim_{n \rightarrow \infty} R_\lambda(f_\gamma^{kNN})$  is equal to

$$\begin{aligned} & 2\lambda\mathbb{E}[\eta(Z)\mathbb{P}\{\text{Binom}(k, \eta(Z)) < \frac{k}{\gamma+1}|Z\}] + \\ & 2\mathbb{E}[(1 - \eta(Z))\mathbb{P}\{\text{Binom}(k, \eta(Z)) > \frac{k}{\gamma+1}|Z\}], \end{aligned} \quad (22)$$

where  $Z = UX + (1 - U)Y$  with  $X \sim P$ ,  $Y \sim Q$  and  $U$  is a fair coin random variable so that  $Z \sim \frac{P+Q}{2}$  and  $\eta(Z) := \mathbb{P}(U = 1|Z) = \frac{dP}{d(P+Q)}(Z)$  (respectively  $1 - \eta(Z) := \mathbb{P}(U = 0|Z) = \frac{dQ}{d(P+Q)}(Z)$ ). The demonstration of Equation (22) follows the same argument as in Devroye et al. (2013)[Theorem 5.2] (up to the occurrence of  $\lambda$  and  $\gamma$  weights) and is not repeated here for the sake of conciseness.

Now, taking  $\gamma = \lambda$  we want to show that the previous expression tends to  $\alpha_\lambda$  which for the Recall equals  $(\lambda P \wedge Q)(\Omega)$  or expressed otherwise as  $2\mathbb{E}[\lambda\eta(Z) \wedge (1 - \eta(Z))]$ .

Equation (22) can be reformulated as  $\lim_{n \rightarrow \infty} R_\lambda(f_\lambda^{kNN}) = 2\mathbb{E}[\mu_\lambda(\eta(Z))]$  with

$$\mu_\lambda(p) = \lambda p \mathbb{P}\{\text{Binom}(k, p) < \frac{k}{\lambda+1}\} + (1 - p) \mathbb{P}\{\text{Binom}(k, p) > \frac{k}{\lambda+1}\}. \quad (23)$$

So that it suffices to show that  $\forall p \in [0, 1]$ ,  $\mu_\lambda(p) \rightarrow \lambda p \wedge (1 - p)$ .

Let's proceed by cases, starting by considering  $\lambda p < (1 - p)$  which is also equivalent to  $p < \frac{1}{\lambda+1}$ . In that case we need to show that  $2\mu_\lambda(p) \rightarrow \lambda p$ . Denoting  $q_\lambda(p) = \mathbb{P}\{\text{Binom}(k, p) > \frac{k}{\lambda+1}\}$ , we have

$$\begin{aligned} \mu_\lambda(p) &= \lambda p (1 - q_\lambda(p)) + (1 - p) q_\lambda(p) \\ &= \lambda p + q_\lambda(p) (1 - (\lambda + 1)p). \end{aligned} \quad (24)$$

Using Hoeffding's inequality (i.e.  $\forall t > 0$ ,  $\mathbb{P}\{\text{Binom}(k, p) - kp > t\} \leq \exp(-2kt^2)$ ) and obtain

$$\begin{aligned} q_\lambda(p) &= \mathbb{P}\{\text{Binom}(k, p) - kp > k(\frac{1}{\lambda+1} - p)\} \\ &\leq \exp\left(-2k\left(\frac{1}{\lambda+1} - p\right)^2\right). \end{aligned} \quad (25)$$

Note that the assumption  $p < \frac{1}{\lambda+1}$  is crucial to apply Hoeffding's inequality (because  $t$  needs to be positive). The right hand side converges to 0 as  $k \rightarrow \infty$  because by assumption  $p \neq \frac{1}{\lambda+1}$ .

The case where  $\lambda p > (1 - p)$  (or  $p > \frac{1}{\lambda+1}$ ) is similar and is left to the reader. In that case, we obtain  $\mu_\lambda(p) \rightarrow (1 - p)$ . There remains the case of equality, that is  $\lambda p = 1 - p = \frac{1}{\lambda+1}$ . In that case, even without taking the limit, one can check that  $\mu_\lambda(p) = \lambda p$ , which concludes the proof.  $\blacksquare$

## Appendix B. Proof of KDE results

### B.1 Subgaussian Random Variables

#### Proposition 21 (Properties on subgaussian random variables)

1. If  $X$  and  $Y$  are independent random vectors in  $\mathbb{R}^d$ , then  $\text{Var}_G(X + Y) \leq \text{Var}_G(X) + \text{Var}_G(Y)$ .
2. If  $X$  is uniformly distributed on the unit ball of  $\mathbb{R}^d$ , then  $\text{Var}_G(X) \leq \frac{\eta}{d}$  where  $\eta > 0$  is a universal constant independent from  $d$

**Proof** The first item is easily obtained from the standard factorization of the exponential combined with the independence assumption. For the second item, see (Vershynin, 2018, Theorem 3.4.5) which establishes the result for the uniform distribution on the sphere and henceforth on the ball as well.  $\blacksquare$

**Lemma 22** Let  $X \in \mathbb{R}^d$  be a  $K$ -subgaussian and centered random vector, and  $W \sim \mathcal{N}(0, I_d)$ , then

$$\mathbb{E}[\exp(a\|X\|^2)] \leq \mathbb{E}[\exp(K^2a\|W\|^2)] = \frac{1}{(1 - 2K^2a)^{\frac{d}{2}}}$$

(the rightmost identity being valid as long as  $1 - 2K^2a > 0$  i.e.  $a < \frac{1}{2K^2}$ ).

In particular, for  $a = \frac{1}{4K^2}$  one obtains

$$\mathbb{E}[\exp(a\|X\|^2)] \leq 2^{\frac{d}{2}}.$$

**Proof** NB: this lemma is a sharper version<sup>4</sup> of Goldfeld et al. (2021, Equation 7) with a direct and basic proof. The first part is referred as the ‘‘Gaussian replacement technique’’ in Vershynin (2018)). First notice that  $\forall x \in \mathbb{R}^d$ ,  $\mathbb{E}[\exp(W^T x)] = \exp\left(\frac{\|x\|^2}{2}\right)$ , and apply it to  $x := \sqrt{2a}X$  (and condition on  $X$  so that it can be considered constant):

$$\begin{aligned} \mathbb{E}[\exp(a\|X\|^2)] &= \mathbb{E}[\mathbb{E}[\exp(W^T \sqrt{2a}X) | X]] = \mathbb{E}[\exp(W^T \sqrt{2a}X)] \\ &= \mathbb{E}[\mathbb{E}[\exp(W^T \sqrt{2a}X) | W]] \leq \mathbb{E}[\exp(\frac{1}{2}K^2 2a\|W\|^2)]. \quad (\text{subgaussianity of } X) \end{aligned}$$

This ends the proof of the first part (i.e. the inequality).

Let us turn to the second part (i.e. the identity),

$$\begin{aligned} \mathbb{E}[\exp(aK^2\|W\|^2)] &= \prod_{i=1}^d \underbrace{\mathbb{E}[\exp(aK^2W_i^2)]}_{= \frac{1}{\sqrt{2\pi}} \int \exp(-\frac{1-2aK^2}{2}w^2)dw} \\ &= \frac{1}{(1 - 2aK^2)^{\frac{d}{2}}} \quad \text{if } a < \frac{1}{2K^2}. \end{aligned}$$

---

4. In their work, for  $a = \frac{1}{4K^2}$  Goldfeld et al obtain an upper bound of the form  $(\exp(\frac{3}{4}))^{\frac{d}{2}}$  which is slightly loose in comparison to our bound since  $\exp(\frac{3}{4}) \approx 2.117$ . Actually our bound is tight, because it is an identity for isotropic Gaussians.

■

## B.2 TV upper bound on bias

**Proof** We have

$$\begin{aligned}
p(x) - (p * k_\sigma)(x) &= \int_y (p(x)k_\sigma(x-y) - p(y)k_\sigma(x-y)) dy \\
&= \int_y (p(x) - p(y)) \frac{k(\frac{x-y}{\sigma})}{\sigma^d} dy \\
&= \int_z (p(x) - p(x - \sigma z)) k(z) dz \\
&= \int_z k(z) \int_{t=0}^1 \frac{d}{dt} p(x - \sigma z t) dz dt \\
&= -\sigma \int_z k(z) \int_{t=0}^1 \nabla p(x - \sigma z t)^T z dt dz.
\end{aligned}$$

Hence,

$$\begin{aligned}
2\|P - P * k_\sigma\|_{\text{TV}} &= \int |p(x) - (p * k)(x)| dx \leq \sigma \int_z |k(z)| \int_{t=0}^1 \underbrace{\int |\nabla p(x - \sigma z t)^T z| dx}_{\leq \int_{x \in \mathbb{R}^d} \|\nabla p(x - \sigma z t)\| \|z\| dx = \|z\| \int \|\nabla p(v)\| dv} dt dz \\
&\leq \sigma \int \|z\| |k(z)| dz \int \|\nabla p(v)\| dv \\
&\leq \sigma R.
\end{aligned}$$

Last, we can notice that for the constant kernel that is both non-negative and supported on the unit ball  $\int \|z\| |k(z)| dz \leq \int k(z) dz = 1$ . ■

## B.3 Proof of Lemma 9

**Proof** We have

$$\begin{aligned}
\mathbb{E}[2\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}] &= \int \mathbb{E}[|\bar{P} * k_\sigma(z) - P * k_\sigma(z)|] dz \\
&= \mathbb{E} \left[ \frac{|\bar{P} * k_\sigma(Z) - P * k_\sigma(Z)|}{f_a(Z)} \right] \\
&\leq \mathbb{E} \left[ \left( \frac{\bar{P} * k_\sigma(Z) - P * k_\sigma(Z)}{f_a(Z)} \right)^2 \right]^{\frac{1}{2}},
\end{aligned}$$

where  $Z \sim f_a(z) dz$  is an artificially introduced random vector in  $\mathbb{R}^d$  independent of the samples and  $f_a$  is a density that can be chosen arbitrarily as long as it is never null. In our case, we use

$$f_a(z) := \mathcal{N}(z; 0, \frac{1}{2a} I) = \left(\frac{a}{\pi}\right)^{\frac{d}{2}} \exp(-a\|z\|^2).$$

Note that in the last inequality is merely the Cauchy-Schwarz one.

Now using the law of total expectation

$$\begin{aligned} \mathbb{E} \left[ \left( \frac{\bar{P} * k_\sigma(Z) - P * k_\sigma(Z)}{f_a(Z)} \right)^2 \right] &= \mathbb{E} \left[ \mathbb{E} \left[ \left( \frac{\bar{P} * k_\sigma(Z) - P * k_\sigma(Z)}{f_a(Z)} \right)^2 \middle| Z \right] \right] \\ &= \mathbb{E} \left[ \underbrace{\text{Var} \left[ \frac{\bar{P} * k_\sigma(Z)}{f_a(Z)} \middle| Z \right]}_{:= (*)} \right]. \end{aligned}$$

Since the samples  $X_1, \dots, X_N$  are independent and  $Z$  is independent from them, the  $X_k$ 's remain independent given  $Z$  is observed. Therefore, we have

$$(*) = \frac{1}{N} \text{Var} \left[ \frac{\delta_{X_1} * k_\sigma(Z)}{f_a(Z)} \middle| Z \right] = \frac{1}{N} \frac{1}{f_a(Z)^2} \text{Var} [k_\sigma(Z - X_1) | Z],$$

which, plugged into the total second moment above, shows that

$$\begin{aligned} \mathbb{E}[2\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}]^2 &\leq \frac{1}{N} \int \frac{1}{f_a(z)} \text{Var} [k_\sigma(z - X_1)] dz \\ &\leq \frac{1}{N} \int \frac{1}{f_a(z)} \mathbb{E} [k_\sigma^2(z - X_1)] dz. \end{aligned}$$

In our case,  $k_\sigma(y) = \frac{1}{\text{vol}_{B_1}(d)\sigma^d} \mathbb{1}_{\|y\| \leq \sigma}$  and therefore  $k_\sigma^2(y) = \frac{1}{\text{vol}_{B_1}(d)\sigma^d} k_\sigma(y)$ . Thus,

$$\begin{aligned} \mathbb{E}[2\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}]^2 &\leq \frac{1}{N \text{vol}_{B_1}(d)\sigma^d} \mathbb{E} \left[ \int \frac{1}{f_a(z)} k_\sigma(z - X_1) dz \right] \\ &= \frac{1}{N \text{vol}_{B_1}(d)\sigma^d} \mathbb{E} \left[ \frac{1}{f_a(Y + X_1)} \right]. \end{aligned}$$

where  $Y \sim k_\sigma(y)dy$  is independent from  $X_1$ .

Given that the constant kernel corresponds to a uniform distribution on a ball of radius  $\sigma$ , Prop 21.2 implies it has a subgaussian variance smaller than  $\frac{\eta\sigma^2}{d}$ . In turn, Prop 21.1 implies that  $S = Y + X_1$  is also sub-gaussian with constant  $K_S$  following  $K_S^2 = K^2 + \frac{\eta}{d}\sigma^2$  where the constant  $\eta$  does not depend on the ambient dimension  $d$ .

Then we will use Lemma 22 with  $a := \frac{1}{4K_S^2}$  (which turns out to be the optimal choice here since it minimizes  $\frac{1}{a(1-2K_S^2a)}$ ):

$$\begin{aligned} \mathbb{E}[2\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}]^2 &\leq \frac{\left(\frac{\pi}{a}\right)^{\frac{d}{2}}}{N \text{vol}_{B_1}(d)\sigma^d} \mathbb{E} [\exp(a\|S\|^2)] \\ &\leq \left( \frac{\pi}{a(1-2K_S^2a)} \right)^{\frac{d}{2}} \frac{1}{N \text{vol}_{B_1}(d)\sigma^d} \quad (\text{Lemma 22}) \\ &\leq (8K_S^2\pi)^{\frac{d}{2}} \frac{1}{N \text{vol}_{B_1}(d)\sigma^d}. \end{aligned}$$

Last, we take the square root of this inequality and use both  $\text{vol}_{B_1}(d) = \frac{\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}$  and  $K_S^2 = K^2 + \frac{\eta}{d}\sigma^2$  to obtain

$$\mathbb{E}[2\|\bar{P} * k_\sigma - P * k_\sigma\|_{\text{TV}}] \leq \sqrt{\Gamma\left(\frac{d}{2}+1\right)} 8^{\frac{d}{4}} \left(\frac{\eta}{d} + \frac{K^2}{\sigma^2}\right)^{\frac{d}{4}} \frac{1}{\sqrt{N}}.$$

■

#### B.4 Proof KDE bias term

**Proof** Before starting the proper proof, let us mention that we will somewhat involve a relative bias between the classifier KDE estimator  $\hat{\alpha}_\lambda^{\text{KDE}}$  and the plug-in KDE estimator  $\hat{\alpha}_\lambda^{\text{TV}}$ . To bound this relative bias we will need to set the latter in a form similar to the former. Recall that from the definition (6), using empirical false positive (resp. negative) rate  $\overline{\text{fpr}}$  (resp.  $\overline{\text{fnr}}$ ), we have

$$\hat{\alpha}_\lambda^{\text{KDE}} = \min_\gamma \lambda \overline{\text{fpr}}(\hat{f}_\gamma^{\text{KDE}}) + \overline{\text{fnr}}(\hat{f}_\gamma^{\text{KDE}}),$$

where the empirical likelihood-ratio classifier parameterized by  $\gamma$  is estimated from training data  $\{\mathcal{X}, \mathcal{Y}\}$  using

$$\hat{f}_\gamma^{\text{KDE}}(z) = \mathbb{1}_{\gamma \hat{P}(z) \geq \hat{Q}(z)}.$$

Now, we recall that the empirical FNR and FPR are estimated on the evaluation data  $\{\mathcal{X}', \mathcal{Y}'\}$

$$\overline{\text{fpr}}(f) = \frac{1}{N} \sum_{x \in \mathcal{X}'} (1 - f(x)) \quad \text{and} \quad \overline{\text{fnr}}(f) = \frac{1}{N} \sum_{y \in \mathcal{Y}'} f(y).$$

On the other hand, the plug-in estimator is

$$\hat{\alpha}_\lambda^{\text{TV}} = \frac{1}{2}(\lambda + 1) - \|\lambda \hat{P} - \hat{Q}\|_{\text{TV}} = \lambda \widehat{\text{fpr}}(\hat{f}_\lambda^{\text{KDE}}) + \widehat{\text{fnr}}(\hat{f}_\lambda^{\text{KDE}}).$$

where the estimated FNR and FPR depend on the training data  $\{\mathcal{X}, \mathcal{Y}\}$  are defined as

$$\widehat{\text{fpr}}(f) = \int_x (1 - f(x)) d\hat{P}(x) \quad \text{and} \quad \widehat{\text{fnr}}(f) = \int_y f(y) d\hat{Q}(y).$$

Let us now proceed with the proof. First, due to the splitting between train/test samples, the bias of the considered estimator is positive. Indeed,

$$\begin{aligned} \mathbb{E} \hat{\alpha}_\lambda^{\text{KDE}} - \alpha_\lambda &= \mathbb{E} \left[ \mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] - \alpha_\lambda \right] \\ &= \mathbb{E} \left[ \mathbb{E} [\lambda \overline{\text{fpr}}(f_{\gamma^*}^{\text{KDE}}) + \overline{\text{fnr}}(f_{\gamma^*}^{\text{KDE}}) | \mathcal{X}, \mathcal{Y}] - \alpha_\lambda \right] \\ &= \mathbb{E} \left[ \underbrace{\lambda \overline{\text{fpr}}(f_{\gamma^*}^{\text{KDE}}) + \overline{\text{fnr}}(f_{\gamma^*}^{\text{KDE}})}_{\geq 0} - \alpha_\lambda \right] \geq 0, \end{aligned}$$

where we have used that for any classifier  $f$ , conditionally to the training set, the expectations of  $\overline{\text{fpr}}(f)$  and  $\overline{\text{fnr}}(f)$  are just the population  $\text{fpr}(f)$  and  $\text{fnr}(f)$ , as well as the fact that the population excess-risk is non-negative.

Therefore

$$\begin{aligned} |\mathbb{E} \hat{\alpha}_\lambda^{\text{KDE}} - \alpha_\lambda| &= \mathbb{E} \hat{\alpha}_\lambda^{\text{KDE}} - \alpha_\lambda \\ &= \mathbb{E} [\mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] - \alpha_\lambda] \\ &= \mathbb{E} [\mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] - \hat{\alpha}_\lambda^{\text{TV}} + \hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda] \quad \text{using } \hat{\alpha}_\lambda^{\text{TV}} \text{ defined in (4)} \\ &\leq \underbrace{\mathbb{E} [\mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] - \hat{\alpha}_\lambda^{\text{TV}}]}_{\text{relative bias w.r.t to plug-in}} + \underbrace{\mathbb{E} |\hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda|}_{\text{TV plug-in risk}}. \end{aligned}$$

Regarding the second term, we have already seen in Proposition 7 that

$$\mathbb{E} |\hat{\alpha}_\lambda^{\text{TV}} - \alpha_\lambda| \leq \lambda \mathbb{E} \|\hat{P} - P\|_{\text{TV}} + \mathbb{E} \|\hat{Q} - Q\|_{\text{TV}}.$$

It remains to show that the relative bias is also subject to the same upper bound. Using once again the fact that the expectation of the empirical FPR (resp. FNR) on an independent evaluation set is the FPR (resp. FNR), we can upper bound the expected estimated Precision

$$\begin{aligned} \mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] &= \mathbb{E} \left[ \min_\gamma \lambda \overline{\text{fpr}}(\hat{f}_\gamma^{\text{KDE}}) + \overline{\text{fnr}}(\hat{f}_\gamma^{\text{KDE}}) | \mathcal{X}, \mathcal{Y} \right] \\ &\leq \lambda \mathbb{E} [\overline{\text{fpr}}(\hat{f}_\lambda^{\text{KDE}}) | \mathcal{X}, \mathcal{Y}] + \mathbb{E} [\overline{\text{fnr}}(\hat{f}_\lambda^{\text{KDE}}) | \mathcal{X}, \mathcal{Y}] \quad (\text{picking } \gamma = \lambda \text{ instead of } \gamma^*) \\ &= \lambda \text{fpr}(\hat{f}_\lambda^{\text{KDE}}) + \text{fnr}(\hat{f}_\lambda^{\text{KDE}}). \end{aligned}$$

So we can bound the classifier relative bias using

$$\begin{aligned} \mathbb{E} [\hat{\alpha}_\lambda^{\text{KDE}} | \mathcal{X}, \mathcal{Y}] - \hat{\alpha}_\lambda^{\text{TV}} &\leq \lambda \int (1 - \hat{f}_\lambda^{\text{KDE}}) d(P - \hat{P}) + \int \hat{f}_\lambda^{\text{KDE}} d(Q - \hat{Q}) \\ &= \lambda \int \hat{f}_\lambda^{\text{KDE}} d(P - \hat{P}) + \int \hat{f}_\lambda^{\text{KDE}} d(Q - \hat{Q}) \\ &= \lambda \int \mathbf{1}_{\lambda \hat{P} \geq \hat{Q}} d(P - \hat{P}) + \int \mathbf{1}_{\lambda \hat{P} \geq \hat{Q}} d(Q - \hat{Q}) \\ &\leq \lambda \|P - \hat{P}\|_{\text{TV}} + \|Q - \hat{Q}\|_{\text{TV}}. \end{aligned}$$

where we have used the definition of  $\|P - \hat{P}\|_{\text{TV}} := \sup_A P(A) - \hat{P}(A) = \sup_A \int \mathbf{1}_A d(P - \hat{P})$ . ■

## Appendix C. Proof of Theorem 20: co-support

**Proof** First, if  $A, A'$  are two co-supports. Then  $Q(A) = Q(A \cap A') = Q(A')$ . Indeed, if  $Q(A) > Q(A \cap A')$  then letting  $B = A \setminus A' \subset A$  one obtains  $Q(B) > 0$ . Yet  $B \subset A'^c$  so that  $P(B) < P(A'^c) = 0$  yielding a contradiction.

Second, let us exhibit a co-support  $C$  that verifies  $Q(C) = \alpha_\infty$ . In Simon et al. (2019), it is shown that for  $\lambda = +\infty$ , Equation 3 can be restated as  $\alpha_\infty = \min_{A \text{ s.t. } P(A^c)=0} Q(A)$ . Let  $A^*$  one of the minimizers. Without further care on the minimizer,  $A^*$  should be merely a particular support of  $P$  but could still not be a co-support. We therefore need to filter out any part of the space that charges  $P$  but not  $Q$  (which will make it a co-support without affecting its  $Q$ -mass). To do so, we consider  $C = A^* \setminus \cap_{\lambda>0} \{\lambda P > Q\}$ . First, the monotone convergence theorem implies that  $Q(C) = Q(A^*) = \alpha_\infty$ .

It remains to show that  $C$  is indeed a co-support. Notice that  $P \wedge Q(C^c) = (P \wedge Q)(A^{*c} \cup \cap_{\lambda>0} \{\lambda P > Q\}) \leq P(A^{*c}) + Q(\cap_{\lambda>0} \{\lambda P > Q\}) = 0$  (because of the constraint on  $A^*$  for the first summand, and by the monotone convergence theorem again for the other summand).

Besides let  $B \subset C \subset (\cap_{\lambda>0} \{\lambda P > Q\})^c = \cup_{\lambda>0} \{\lambda P \leq Q\}$  so that  $Q(B) = 0 \implies P(B) = 0$ . Conversely, if  $P(B) = 0$  let us show that  $Q(B) = 0$ . To do so let us reason by contradiction, by assuming that  $Q(B) > 0$ . Then  $A = A^* \setminus B$  verifies the constraint  $P(A) = 0$  and  $Q(A) = Q(A^*) - Q(B)$  (because  $B \subset C \subset A^*$ ). Then  $Q(A) < Q(A^*)$  would contradict the definition of  $A^*$ . ■

## Appendix D. Additional experimental results

### D.1 Gaussians shifts

Table. 1 complements Section 5.1.1 and Figure 2 by comparing various estimated PR curves with respect to the ground truth, using average IoU scores. While this observation contrasts with findings reported in the literature, it should be understood as referring to the overall shape of the curve. In contrast, prior studies focused exclusively on the extremes, due to inherent limitations in the PR metrics under investigation.

### D.2 Gaussian mixture comparison

Figure 16 complements Section 5.1.2 and Figure 3c with additional curves for different setting (w/ and w/o splitting,  $k = 4$  or  $k = 100$ ).

		shift $\mu$	iPR	kNN	KDE	Cov
with 50 % split	$k = 4$	0.12	0.69	0.71	0.72	0.73
		0.21	0.42	0.49	0.49	0.55
		0.29	0.24	0.38	0.34	0.48
		0.38	0.13	0.33	0.24	0.48
	$k = \sqrt{n}$	0.12	0.81	0.87	0.84	0.92
		0.21	0.69	0.84	0.78	0.90
		0.29	0.65	0.84	0.75	0.90
		0.38	0.63	0.84	0.75	0.93
without split	$k = 4$	0.12	0.43	0.7	0.62	0.76
		0.21	0.55	0.81	0.68	0.84
		0.29	0.62	0.79	0.68	0.77
		0.38	0.55	0.61	0.62	0.63
	$k = \sqrt{n}$	0.12	0.91	0.93	0.94	0.96
		0.21	0.88	0.93	0.92	0.97
		0.29	0.84	0.92	0.90	0.95
		0.38	0.83	0.91	0.90	0.96

Table 1: **Mean IoU scores** for shifted Gaussians. Standard deviations are  $< 10^{-2}$  with  $n = 10K$ .

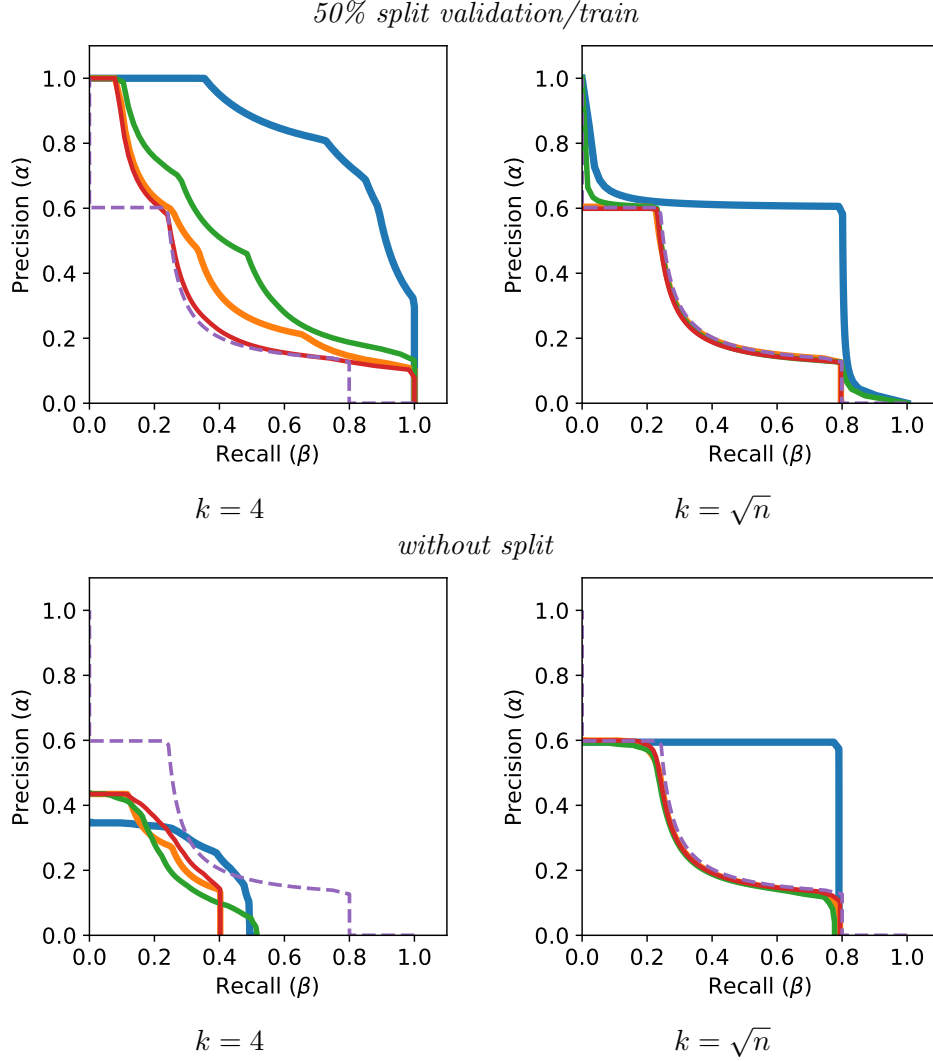


Figure 16: **Comparing two Gaussian mixtures.** This figure complements Figure 3c. The ground truth PR curve ( - -GT) is compared to empirical estimates from various NN-classifiers: -iPR, -kNN, -KDE, and -Cov. Here  $P$  and  $Q$  are two GMMs sharing the same modes (centered at  $\mu_k$ ):  $P = \sum_{\ell} p_{\ell} \mathcal{N}(\mu_{\ell} \mathbf{1}_d, \mathbb{I}_d)$  and  $Q = \sum_{\ell} q_{\ell} \mathcal{N}(\mu_{\ell} \mathbf{1}_d, \mathbb{I}_d)$  with  $d = 64$  dimensions and  $\mu_{\ell} \in \{0, -5, 3, 5\}$ . However,  $P$  and  $Q$  have different weights ( $p_{\ell}$  and  $q_{\ell}$ ):  $p_{\ell} \in \{0.3, 0.2, 0.5, 0\}$  vs  $q_{\ell} \in \{0, 0.5, 0.2, 0.3\}$ .  $n = 1\text{k}$  points are sampled and split in half between validation and train, and  $k = \sqrt{n}$ .

## References

- Fasil Cheema and Ruth Urner. Precision recall cover: A method for assessing generative models. In *International Conference on Artificial Intelligence and Statistics*, pages 6571–6594. PMLR, 2023.
- Luc Devroye and Laszlo Györfi. No Empirical Probability Measure can Converge in the Total Variation Sense for all Distributions. *The Annals of Statistics*, 18:1496–1499, 1990.
- Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- Josip Djolonga, Mario Lucic, Marco Cuturi, Olivier Bachem, Olivier Bousquet, and Sylvain Gelly. Precision-recall curves using information divergence frontiers. In *International Conference on Artificial Intelligence and Statistics*, pages 2550–2559. PMLR, 2020.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling Rectified Flow Transformers for High-Resolution Image Synthesis. In *Proceedings of the 41st International Conference on Machine Learning*, pages 12606–12633. PMLR, July 2024.
- Ziv Goldfeld, Kristjan Greenewald, Yury Polyanskiy, and Jonathan Weed. Convergence of Smoothed Empirical Measures with Applications to Entropy Estimation. *IEEE Transactions on Information Theory*, 66:4368–4391, 2021. doi: 10.1109/TIT.2020.2975480. URL <http://arxiv.org/abs/1905.13576>.
- 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.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021. URL <https://openreview.net/forum?id=qw8AKxfYbI>.
- Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- Mahyar Khayatkhoei and Wael AbdAlmageed. Emergent asymmetry of precision and recall for measuring fidelity and diversity of generative models in high dimensions. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023.

- Pum Jun Kim, Yoojin Jang, Jisu Kim, and Jaejun Yoo. TopP&R: Robust Support Estimation Approach for Evaluating Fidelity and Diversity in Generative Models. In *37th Conference on Neural Information Processing Systems (NeurIPS 2023)*, 2023.
- 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.
- Oleg Lepski, Arkady Nemirovski, and Vladimir Spokoiny. On estimation of the  $l_r$  norm of a regression function. *Probability theory and related fields*, 113:221–253, 1999.
- Lorenzo Luzi, Carlos Ortiz Marrero, Nile Wynar, Richard G Baraniuk, and Michael J Henry. Evaluating generative networks using gaussian mixtures of image features. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 279–288, 2023.
- 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.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khali-dov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dino2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal*, pages 1–31, 2024.
- Dogyun Park and Suhyun Kim. Probabilistic precision and recall towards reliable evaluation of generative models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20099–20109, 2023.
- Krishna Pillutla, Swabha Swayamdipta, Rowan Zellers, John Thickstun, Sean Welleck, Yejin Choi, and Zaid Harchaoui. MAUVE: Measuring the Gap Between Neural Text and Human Text using Divergence Frontiers. In *NeurIPS*, 2021. URL <https://openreview.net/pdf?id=Tqx7nJp7PR>.
- Krishna Pillutla, Lang Liu, John Thickstun, Sean Welleck, Swabha Swayamdipta, Rowan Zellers, Sewoong Oh, Yejin Choi, and Zaid Harchaoui. MAUVE Scores for Generative Models: Theory and Practice. *Journal of Machine Learning Research*, 2023.
- 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.
- 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.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1409.1556>.

Rodrigue Siry, Ryan Webster, Loic Simon, and Julien Rabin. On the theoretical equivalence of several trade-off curves assessing statistical proximity. *Journal of Machine Learning Research*, 24(185):1–34, 2023. URL <http://jmlr.org/papers/v24/21-0607.html>.

Stability AI. Stable diffusion 3.5 medium, 2025. URL <https://huggingface.co/stabilityai/stable-diffusion-3.5-medium>.

Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826. IEEE, 2016. ISBN 978-1-4673-8851-1. doi: 10.1109/CVPR.2016.308. URL <http://ieeexplore.ieee.org/document/7780677/>.

Alexandre B Tsybakov. Nonparametric estimators. *Introduction to Nonparametric Estimation*, pages 1–76, 2009.

Alexandre Verine. *Quality and Diversity in Generative Models through the Lens of F-Divergences*. PhD thesis, Université Paris sciences et lettres, 2024.

Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.