

Appendix

A More Background on Topological Data Analysis

Topological data analysis (TDA) [19] is a recent and emerging field of data science that relies on topological tools to infer relevant features for possibly complex data. A key object in TDA is persistent homology, which quantifies salient topological features of data by observing them in multi-resolutions.

A.1 Persistent Homology

Persistent homology. *Persistent homology* is a multiscale approach to represent the topological features. For a filtration \mathcal{F} and for each $k \in \mathbb{N}_0 = \mathbb{N} \cup \{0\}$, the associated k -th persistent homology $PH_k\mathcal{F}$ is a collection of k -th dimensional homologies $\{H_k\mathcal{F}_\delta\}_{\delta \in \mathbb{R}}$ equipped with homomorphisms $\{\phi_k^{a,b} : H_k\mathcal{F}_a \rightarrow H_k\mathcal{F}_b\}_{a \leq b}$ induced by the inclusion $\mathcal{F}_a \subset \mathcal{F}_b$.

Persistence diagram. For the k -th persistent homology $PH_k\mathcal{F}$, the set of filtration levels at which a specific homology appears is always an interval $[b, d) \subset [-\infty, \infty]$. The corresponding k -th persistence diagram is a multiset of points $(\mathbb{R} \cup \{\infty\})^2$, consisting of all pairs (b, d) where $[b, d)$ is the interval of filtration values for which a specific homology appears in $PH_k\mathcal{F}$.

A.2 Statistical Inference of Persistent Homology

As discussed above, a homological feature with a long life-length is important information in topology while the homology with a short life-length can be treated as non-significant information or noise. The confidence band estimator provides the confidence set from the features that only include topologically and statistically significant (statistically considered as elements in the population set) under a certain level of confidence. And to build a confidence set, we first need to endow a metric on the space of persistence diagrams.

Bottleneck distance. The most fundamental metric to measure the distance between two persistence diagrams is the *bottleneck distance*.

Definition A.1. The *bottleneck distance* between two persistence diagrams D_1 and D_2 is defined by

$$d_B(D_1, D_2) = \inf_{\gamma \in \Gamma} \sup_{p \in D_1} \|p - \gamma(p)\|_\infty,$$

where the set Γ consists of all the bijections $\gamma : D_1 \cup \text{Diag} \rightarrow D_2 \cup \text{Diag}$, and Diag is the diagonal $\{(x, x) : x \in \mathbb{R}\} \subset \mathbb{R}^2$ with infinite multiplicity.

One way of constructing the confidence set uses the superlevel filtration of the kernel density estimator and the bootstrap confidence band. Let $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ as given points cloud, then the probability for the distribution of points can be estimated via KDE defined as following: $\hat{p}_h(x) := \frac{1}{nh^d} \sum_{i=1}^n K\left(\frac{x-X_i}{h}\right)$ where h is the bandwidth and d as a dimension of the space. We compute \hat{p}_h and \hat{p}_h^* , which are the KDE of \mathcal{X} and the KDE of bootstrapped samples \mathcal{X}^* , respectively. Now, given the significance level α and $h > 0$, let confidence band $q_{\mathcal{X}}$ be bootstrap bandwidth of a Gaussian Empirical Process [32, 33], $\sqrt{n} \|\hat{p}_h - \hat{p}_h^*\|_\infty$. Then it satisfies $P(\sqrt{n} \|\hat{p}_h - p_h\|_\infty < q_{\mathcal{X}}) \geq 1 - \alpha$, as in Proposition D.2 in Section D. Then $\mathcal{B}_{d_B}(\hat{\mathcal{P}}_h, c_{\mathcal{X}})$, the ball of persistent homology centered at $\hat{\mathcal{P}}_h$ and radius $c_{\mathcal{X}} = q_{\mathcal{X}}/\sqrt{n}$ in the bottleneck distance d_B , is a valid confidence set as $\liminf_{n \rightarrow \infty} \mathbb{P}(\mathcal{P} \in \mathcal{B}_{d_B}(\hat{\mathcal{P}}_h, c_{\mathcal{X}})) \geq 1 - \alpha$. This confidence set has further interpretation that in the persistence diagram, homological features that are above twice the radius $2c_{\mathcal{X}}$ from the diagonal are simultaneously statistically significant.

B Johnson-Lindenstrauss Lemma

Johnson-Lindenstrauss Lemma is stated as follows:

Lemma B.1 (Johnson-Lindenstrauss Lemma). [34, Lemma 1]

Let $0 < \epsilon < 1$ and $\mathcal{X} \subset \mathbb{R}^D$ be a set of points with size n . Then for $d = O(\min\{n, D, \epsilon^{-2} \log n\})$, there exists a linear map $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ such that for all $x, y \in \mathcal{X}$,

$$(1 - \epsilon) \|x - y\|^2 \leq \|f(x) - f(y)\|^2 \leq (1 + \epsilon) \|x - y\|^2. \quad (5)$$

Remark B.2. This Johnson-Lindenstrass Lemma is optimal for linear maps, and nearly-optimal even if we allow non-linear maps: there exists a set of points $\mathcal{X} \subset \mathbb{R}^D$ with size n , such that d should satisfy: (a) $d = \Omega(\epsilon^{-2} \log n)$ for a linear map $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ satisfying (5) to exist [35, Theorem 3], and (b) $d = \Omega(\epsilon^{-2} \log(\epsilon^2 n))$ for a map $f : \mathbb{R}^D \rightarrow \mathbb{R}^d$ satisfying (5) to exist [36, Theorem 1].

C Denoising topological features from outliers

Using the bootstrap bandwidth $c_{\mathcal{X}}$ as the threshold is the key part of our estimators TopP&R for robustly estimating $\text{supp}(P)$. When the level set $\hat{p}_h^{-1}[c_{\mathcal{X}}, \infty)$ is used, the homology of $\hat{p}_h^{-1}[c_{\mathcal{X}}, \infty)$ consists of homological features whose (birth) $\geq c_{\mathcal{X}}$ and (death) $\leq c_{\mathcal{X}}$, which are the homological features in skyblue area in Figure A1. In this example, we consider three types of homological noise, though there can be many more corresponding to different homological dimensions.

- There can be a 0-dimensional homological noise of (birth) $< c_{\mathcal{X}}$ and (death) $< c_{\mathcal{X}}$, which is the red point in the persistence diagram of Figure A1. This noise corresponds to the orange connected component on the left. As in the figure, this type of homological noise usually corresponds to outliers.
- There can be a 0-dimensional homological noise of (birth) $> c_{\mathcal{X}}$ and (death) $> c_{\mathcal{X}}$, which is the green point in the persistence diagram of Figure A1. This noise corresponds to the connected component surrounded by the green line on the left. As in the figure, this type of homological noise lies within the estimated support, not like the other two.
- There can be a 1-dimensional homological noise of (birth) $< c_{\mathcal{X}}$ and (death) $< c_{\mathcal{X}}$, which is the purple point in the persistence diagram of Figure A1. This noise corresponds to the purple loop on the left.

These homological noises satisfy either their (birth) $< c_{\mathcal{X}}$ and (death) $< c_{\mathcal{X}}$ or their (birth) $> c_{\mathcal{X}}$ and (death) $> c_{\mathcal{X}}$ simultaneously with high probability, so those homological noises are removed in the estimated support $\hat{p}_h^{-1}[c_{\mathcal{X}}, \infty)$, which is the blue area in the left and the skyblue area in the right in Figure A1.

We would like to further emphasize that homological noises are not restricted to 0-dimension lying outside the estimated support (red point in the persistence diagram of Figure A1). 0-dimensional homological noise inside the estimated support (green point in the persistence diagram of Figure A1), 1-dimensional homological noise can also arise, and the bootstrap bandwidth $c_{\mathcal{X}}$ allows to simultaneously filter them.

D Assumptions on distributions and kernels

For distributions, we assume that the order of probability volume decay $P(\mathcal{B}(x, r))$ is at least r^d .

Assumption A1. For all $x \in \text{supp}(P)$ and $y \in \text{supp}(Q)$,

$$\liminf_{r \rightarrow 0} \frac{P(\mathcal{B}(x, r))}{r^d} > 0, \quad \liminf_{r \rightarrow 0} \frac{Q(\mathcal{B}(y, r))}{r^d} > 0.$$

Remark D.1. Assumption A1 is analogous to Assumption 2 of Kim et al. [37], but is weaker since the condition is pointwise on each $x \in \mathbb{R}^d$. And this condition is much weaker than assuming a density on \mathbb{R}^d : for example, a distribution supported on a low-dimensional manifold satisfies Assumption A1. This provides a framework suitable for high dimensional data, since many times high dimensional data lies on a low dimensional structure hence its density on \mathbb{R}^d cannot exist. See Kim et al. [37] for a more detailed discussion.

For kernel functions, we assume the following regularity conditions:

Assumption A2. Let $K : \mathbb{R}^d \rightarrow \mathbb{R}$ be a nonnegative function with $\|K\|_1 = 1$, $\|K\|_{\infty} < \infty$, $\|K\|_2 < \infty$, and satisfy the following:

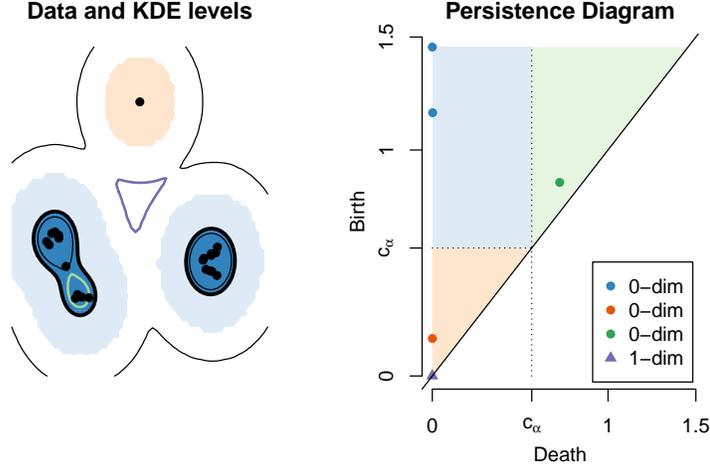


Figure A1: To robustly estimate the support, we use the bootstrap bandwidth c_α to filter out topological noise (orange) and keep topological signal (skyblue). Then TopP&R is computed on this support.

- (i) $K(0) > 0$.
- (ii) K has a compact support.
- (iii) K is Lipschitz continuous and of second order.

Assumption A2 allows to build a valid bootstrap confidence band for kernel density estimator (KDE). See Theorem 12 of [26] or Theorem 3.4 of [38]

Proposition D.2 (Theorem 3.4 of [38]). *Let $\mathcal{X} = \{X_1, \dots, X_n\}$ be IID from a distribution P . For $h > 0$, let \hat{p}_h, \hat{p}_h^* be kernel density estimator for \mathcal{X} and its bootstrap \mathcal{X}^* , respectively, and for $\alpha \in (0, 1)$, let c_α be the α bootstrap quantile from $\sqrt{nh^d} \|\hat{p}_h - \hat{p}_h^*\|_\infty$. For $h_n \rightarrow 0$,*

$$\mathbb{P}\left(\sqrt{nh_n^d} \|\hat{p}_{h_n} - p_{h_n}\|_\infty > c_\alpha\right) = \alpha + \left(\frac{\log n}{nh_n^d}\right)^{\frac{4+d}{4+2d}}.$$

Assumption A1, A2 ensures that, when the bandwidth $h_n \rightarrow 0$, average KDEs are bounded away from 0.

Lemma D.3. *Let P be a distribution satisfying Assumption A1. Suppose K is a nonnegative function satisfying $K(0) > 0$ and continuous at 0. Suppose $\{h_n\}_{n \in \mathbb{N}}$ is given with $h_n \geq 0$ and $h_n \rightarrow 0$. Then for all $x \in \text{supp}(P)$,*

$$\liminf_{n \rightarrow \infty} p_{h_n}(x) > 0.$$

Proof. Since $K(0) > 0$ and K is continuous at 0, there is $r_0 > 0$ such that for all $y \in \mathcal{B}(0, r_0)$, $K(y) \geq \frac{1}{2}K(0) > 0$. And hence

$$\begin{aligned} p_h(x) &= \int \frac{1}{h^d} K\left(\frac{x-y}{h}\right) dP(y) \geq \int \frac{K(0)}{2h^d} \mathbf{1}\left(\frac{x-y}{h} \in \mathcal{B}(0, r_0)\right) dP(y) \\ &\geq \frac{K(0)}{2h^d} P(\mathcal{B}(x, r_0h)). \end{aligned}$$

Hence as $h_n \rightarrow 0$,

$$\liminf_{n \rightarrow \infty} p_{h_n}(x) > 0.$$

□

Before specifying the rate of convergence, we introduce the concept of reach. First introduced by [39], the reach is a quantity expressing the degree of geometric regularity of a set. Given a closed subset $A \subset \mathbb{R}^d$, the medial axis of A , denoted by $\text{Med}(A)$, is the subset of \mathbb{R}^d consisting of all the points that have at least two nearest neighbors on A .

$$\text{Med}(A) = \{x \in \mathbb{R}^d \setminus A : \exists q_1 \neq q_2 \in A, \|q_1 - x\| = \|q_2 - x\| = d(x, A)\},$$

where $d(x, A) = \inf_{q \in A} \|q - x\|$ denotes the distance from a generic point $x \in \mathbb{R}^d$ to A . The reach of A is then defined as the minimal distance from A to $\text{Med}(A)$.

Definition D.4. The reach of a closed subset $A \subset \mathbb{R}^d$ is defined as

$$\text{reach}(A) = \inf_{q \in A} d(q, \text{Med}(A)) = \inf_{q \in A, x \in \text{Med}(A)} \|q - x\|.$$

Now, for specifying the rate of convergence, we first assume that distributions have densities away from 0 and ∞ .

Assumption A3. P and Q have Lebesgue densities p and q that, there exists $0 < p_{\min} \leq p_{\max} < \infty$, $0 < q_{\min} \leq q_{\max} < \infty$ with for all $x \in \text{supp}(P)$ and $y \in \text{supp}(Q)$,

$$p_{\min} \leq p(x) \leq p_{\max}, \quad q_{\min} \leq q(x) \leq q_{\max}.$$

We also assume weak geometric assumptions on the support of the distributions $\text{supp}(P)$ and $\text{supp}(Q)$, being bounded and having positive reach.

Assumption A4. We assume $\text{supp}(P)$ and $\text{supp}(Q)$ are bounded. And the support of P and Q have positive reach, i.e. $\text{reach}(\text{supp}(P)) > 0$ and $\text{reach}(\text{supp}(Q)) > 0$.

Sets with positive reach ensure that the volume of its tubular neighborhood grows in polynomial order, which is Theorem 26 from [40] and originally from [39].

Proposition D.5. Let A be a set with $\text{reach}(A) > 0$. Let $A_r := \{x \in \mathbb{R}^d : d(x, A) \leq r\}$. Then for $r < \text{reach}(A)$, there exists $a_0, \dots, a_d \in \mathbb{R} \cup \{\infty\}$ that satisfies

$$\lambda_d(A_r) = \sum_{k=0}^{d-1} a_k r^k,$$

where λ_d is the usual Lebesgue measure of \mathbb{R}^d .

E Details and Proofs for Section 4

For a random variable X and $\alpha \in (0, 1)$, let $q_{X, \alpha}$ be upper α -quantile of X , i.e., $\mathbb{P}(X \geq q_\alpha) = \alpha$, or equivalently, $\mathbb{P}(X < q_\alpha) = 1 - \alpha$. Let \tilde{p}_h be the KDE on \mathcal{X}^0 . For a finite set \mathcal{X} , we use the notation $c_{\mathcal{X}, \alpha}$ for α -bootstrap quantile satisfying $\mathbb{P}(\|\hat{p}_{\mathcal{X}, h_n} - \hat{p}_{\mathcal{X}^b, h_n}\|_\infty < c_{\mathcal{X}, \alpha} | \mathcal{X}) = 1 - \alpha$, where \mathcal{X}^b is the bootstrap sample from \mathcal{X} . For a distribution P , we use the notation $c_{P, \alpha}$ for α -quantile satisfying $\mathbb{P}(\|\hat{p}_{h_n} - p_{h_n}\|_\infty < c_{P, \alpha}) = 1 - \alpha$, where \hat{p}_{h_n} is kernel density estimator of IID samples from P . Hence when \mathcal{X} is not IID samples from P , the relation of Proposition D.2 may not hold.

Claim E.1. Let $Z \sim \mathcal{N}(0, 1)$ be a sample from standard normal distribution, then for $\alpha \in (0, 1)$,

$$q_{Z, \alpha} = \Theta\left(\sqrt{\log(1/\alpha)}\right).$$

And for $0 \leq \delta < \alpha < 1$,

$$q_{Z, \alpha} - q_{Z, \alpha + \delta} = \left(\delta \alpha \sqrt{\log(1/\alpha)}\right)$$

Proof. Let $Z \sim \mathcal{N}(0, 1)$, then for all $x > 2$,

$$\frac{1}{2x} \exp\left(-\frac{x^2}{2}\right) \leq \mathbb{P}(Z > x) \leq \frac{1}{x} \exp\left(-\frac{x^2}{2}\right).$$

Then $x = C\sqrt{\log(1/\alpha)}$ gives

$$\frac{\alpha^{C/2}}{2C\sqrt{\log(1/\alpha)}} \leq \mathbb{P}(Z > x) \leq \frac{\alpha^{C/2}}{C\sqrt{\log(1/\alpha)}}.$$

And hence

$$q_{Z,\alpha} = \Theta\left(\sqrt{\log(1/\alpha)}\right),$$

and in particular,

$$\exp\left(-\frac{q_{Z,\alpha}^2}{2}\right) = \Theta\left(\alpha\sqrt{\log(1/\alpha)}\right).$$

Now for $0 \leq \delta < \alpha < 1$,

$$\begin{aligned} (q_{Z,\alpha} - q_{Z,\alpha+\delta}) \exp\left(-\frac{q_{Z,\alpha}^2}{2}\right) &\leq \int_{q_{Z,\alpha+\delta}}^{q_{Z,\alpha}} \exp\left(-\frac{t^2}{2}\right) dt = \delta \\ &\leq (q_{Z,\alpha} - q_{Z,\alpha+\delta}) \exp\left(-\frac{q_{Z,\alpha+\delta}^2}{2}\right). \end{aligned}$$

And hence

$$q_{Z,\alpha} - q_{Z,\alpha+\delta} \leq \delta \exp\left(-\frac{q_{Z,\alpha+\delta}^2}{2}\right) = \Theta\left(\delta(\alpha + \delta)\sqrt{\log(1/\alpha)}\right),$$

and

$$q_{Z,\alpha} - q_{Z,\alpha+\delta} \geq \delta \exp\left(-\frac{q_{Z,\alpha}^2}{2}\right) = \Theta\left(\delta\alpha\sqrt{\log(1/(\alpha + \delta))}\right).$$

Then from $\delta < \alpha$,

$$q_{Z,\alpha} - q_{Z,\alpha+\delta} = \Theta\left(\delta\alpha\sqrt{\log(1/\alpha)}\right).$$

□

Claim E.2. Let X, Y be random variables, and $0 \leq \delta < \alpha < 1$. Suppose there exists $c > 0$ satisfying

$$\mathbb{P}(|X - Y| > c) \leq \delta. \quad (6)$$

Then,

$$q_{X,\alpha+\delta} - c \leq q_{Y,\alpha} \leq q_{X,\alpha-\delta} + c.$$

Proof. Note that $q_{X,\alpha-\delta}$ satisfies $\mathbb{P}(X > q_{X,\alpha-\delta}) = \alpha - \delta$. Then from (6),

$$\begin{aligned} \mathbb{P}(Y > q_{X,\alpha-\delta} + c) &\leq \mathbb{P}(Y > q_{X,\alpha-\delta} + c, |X - Y| \leq c) + \mathbb{P}(|X - Y| > c) \\ &\leq \mathbb{P}(X > q_{X,\alpha-\delta}) + \delta = \alpha. \end{aligned}$$

And hence

$$q_{Y,\alpha} \leq q_{X,\alpha-\delta} + c.$$

Then changing the role of X and Y gives

$$q_{X,\alpha+\delta} - c \leq q_{Y,\alpha}.$$

□

Lemma E.3. (i) Under Assumption 1, 2 and A2,

$$\|\hat{p}_h - \tilde{p}_h\|_\infty \leq \frac{\rho_n \|K\|_\infty}{h^d}.$$

(ii) Under Assumption 1, 2 and A2, with probability $1 - \delta$,

$$c_{\mathcal{X}^0,\alpha+\delta} - O\left(\frac{\rho_n}{h^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh^{2d}}}\right) \leq c_{\mathcal{X},\alpha} \leq c_{\mathcal{X}^0,\alpha-\delta} + O\left(\frac{\rho_n}{h^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh^{2d}}}\right).$$

(iii) Suppose Assumption 1, 2, A2 hold, and let $\alpha, \delta_n \in (0, 1)$. Suppose $nh_n^d \rightarrow \infty$, $\delta_n^{-1} =$

$$O\left(\left(\frac{\log n}{nh_n^d}\right)^{\frac{4+d}{4+2d}}\right) \text{ and } nh_n^{-d} \rho_n^2 \delta_n^{-2} \rightarrow 0. \text{ Then with probability } 1 - \alpha - 4\delta_n,$$

$$\|\hat{p}_h - p_h\|_\infty < c_{\mathcal{X},\alpha} \leq c_{P,\alpha-3\delta_n}.$$

Proof. (i)

First, note that

$$\hat{p}_h(x) - \tilde{p}_h(x) = \frac{1}{nh^d} \sum_{i=1}^n \left(K\left(\frac{x - X_i}{h}\right) - K\left(\frac{x - X_i^0}{h}\right) \right).$$

Then under Assumption A2,

$$\begin{aligned} \|\hat{p}_h - \tilde{p}_h\|_\infty &\leq \frac{1}{nh^d} \sum_{i=1}^n \left\| K\left(\frac{\cdot - X_i}{h}\right) - K\left(\frac{\cdot - X_i^0}{h}\right) \right\|_\infty \\ &\leq \frac{1}{nh^d} \sum_{i=1}^n \|K\|_\infty I(X_i \neq X_i^0). \end{aligned}$$

Then from Assumption 2, $\sum_{i=1}^n I(X_i \neq X_i^0) \leq n\rho_n$, and hence

$$\|\hat{p}_h - \tilde{p}_h\|_\infty \leq \frac{\|K\|_\infty \rho_n}{h^d}.$$

(ii)

Let $\mathcal{X}_b, \mathcal{X}_b^0$ be bootstrapped samples of $\mathcal{X}, \mathcal{X}^0$ with the same sampling with replacement process. Let $\hat{p}_h^b, \tilde{p}_h^b$ be KDE of \mathcal{X}_b and \mathcal{X}_b^0 , respectively. And, note that

$$\left| \|\hat{p}_h - \hat{p}_h^b\|_\infty - \|\tilde{p}_h - \tilde{p}_h^b\|_\infty \right| \leq \|\hat{p}_h - \tilde{p}_h\|_\infty + \|\hat{p}_h^b - \tilde{p}_h^b\|_\infty.$$

Let L_b be the number of elements where \mathcal{X}_b and \mathcal{X}_b^0 differ, i.e., $L_b = |\mathcal{X}_b \setminus \mathcal{X}_b^0| = |\mathcal{X}_b^0 \setminus \mathcal{X}_b|$, then $L_b \sim \text{Binomial}(n, \rho_n)$, and

$$\|\hat{p}_h^b - \tilde{p}_h^b\|_\infty \leq \frac{\|K\|_\infty L_b}{nh^d}.$$

And hence,

$$\left| \|\hat{p}_h - \hat{p}_h^b\|_\infty - \|\tilde{p}_h - \tilde{p}_h^b\|_\infty \right| \leq \frac{\|K\|_\infty (n\rho_n + L_b)}{nh^d}. \quad (7)$$

Then by Hoeffding's inequality, with probability $1 - \delta$,

$$L_b \leq n\rho_n + \sqrt{\frac{n \log(1/\delta)}{2}}.$$

by applying this to (7), with probability $1 - \delta$,

$$\left| \|\hat{p}_h - \hat{p}_h^b\|_\infty - \|\tilde{p}_h - \tilde{p}_h^b\|_\infty \right| \leq O\left(\frac{\rho_n}{h^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh^{2d}}}\right).$$

Hence applying Claim E.2 implies that, with probability $1 - \delta$,

$$c_{\mathcal{X}^0, \alpha + \delta} - O\left(\frac{\rho_n}{h^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh^{2d}}}\right) \leq c_{\mathcal{X}, \alpha} \leq c_{\mathcal{X}^0, \alpha - \delta} + O\left(\frac{\rho_n}{h^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh^{2d}}}\right).$$

(iii)

Let $\tilde{\delta}_n := O\left(\left(\frac{\log n}{nh^d}\right)^{\frac{4+d}{4+2d}}\right)$ be from RHS of Proposition D.2, then for large enough n , $\delta_n \geq \tilde{\delta}_n$. Note that Proposition D.2 implies that for all $\alpha \in (0, 1)$,

$$c_{P, \alpha + \delta_n} \leq c_{\mathcal{X}_0, \alpha} \leq c_{P, \alpha - \delta_n}. \quad (8)$$

Now, $nh_n^d \rightarrow \infty$ implies that $\sqrt{nh_n^d}(\tilde{p}_{h_n} - p_{h_n})$ converges to a Gaussian process, and then Claim E.2 implies that $c_{P, \alpha} = \Theta\left(\sqrt{\frac{\log(1/\alpha)}{nh_n^d}}\right)$ and

$$c_{P,\alpha} - c_{P,\alpha+\delta_n} = \Theta \left(\delta_n \alpha \sqrt{\frac{\log(1/\alpha)}{nh_n^d}} \right). \quad (9)$$

Then under Assumption 2, since $nh_n^{-d} \rho_n^2 \delta_n^{-2} = o(1)$ and $n\rho_n \geq 1$ implies $h_n^{-d} \rho_n \delta_n^{-2} = o(1)$, and then

$$\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} = O \left(\delta_n \alpha \sqrt{\frac{\log(1/\alpha)}{nh_n^d}} \right). \quad (10)$$

Then (8), (9), (10) implies that

$$c_{P,\alpha+3\delta_n} \leq c_{P,\alpha+2\delta_n} - O \left(\delta_n \alpha \sqrt{\frac{\log(1/\alpha)}{nh_n^d}} \right) \leq c_{\mathcal{X}^0,\alpha+\delta_n} - O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right), \quad (11)$$

and

$$c_{\mathcal{X}^0,\alpha-\delta_n} + O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right) \leq c_{P,\alpha-2\delta_n} + O \left(\delta_n \alpha \sqrt{\frac{\log(1/\alpha)}{nh_n^d}} \right) \leq c_{P,\alpha-3\delta_n}. \quad (12)$$

Now, from the definition of $c_{P,\alpha+3\delta_n}$, with probability $1 - \alpha - 3\delta_n$,

$$\|\hat{p}_h - p_h\|_\infty \leq c_{P,\alpha+3\delta_n}. \quad (13)$$

And (ii) implies that, with probability $1 - \delta_n$,

$$c_{\mathcal{X}^0,\alpha+\delta_n} - O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right) \leq c_{\mathcal{X},\alpha} \leq c_{\mathcal{X}^0,\alpha-\delta_n} + O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right). \quad (14)$$

Hence by combining (11), (12), (13), (14), with probability $1 - \alpha - 4\delta_n$,

$$\begin{aligned} \|\hat{p}_h - p_h\|_\infty &\leq c_{P,\alpha+3\delta_n} \\ &\leq c_{\mathcal{X}^0,\alpha+\delta_n} - O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right) \\ &\leq c_{\mathcal{X},\alpha} \\ &\leq c_{\mathcal{X}^0,\alpha-\delta_n} + O \left(\frac{\rho_n}{h_n^d} + \sqrt{\frac{\rho_n \log(1/\delta)}{nh_n^{2d}}} \right) \\ &\leq c_{P,\alpha-3\delta_n}. \end{aligned}$$

□

Corollary E.4. *Suppose Assumption 1, 2, A2 hold, and let $\alpha \in (0, 1)$. Suppose $\delta_n^{-1} = O \left(\left(\frac{\log n}{nh_n^d} \right)^{\frac{4+d}{4+2d}} \right)$ and $nh_n^{-d} \rho_n^2 \delta_n^{-2} \rightarrow 0$. Then with probability $1 - \alpha - 4\delta_n$,*

(i) *Let $\delta_n \in (0, 1)$ and suppose $nh_n^d \rightarrow \infty$, $\delta_n^{-1} = O \left(\left(\frac{\log n}{nh_n^d} \right)^{\frac{4+d}{4+2d}} \right)$ and $nh_n^{-d} \rho_n^2 \delta_n^{-2} \rightarrow 0$. Then with probability $1 - \alpha - 4\delta_n$,*

$$p_{h_n}^{-1}[2c_{P,\alpha-3\delta_n}, \infty) \subset \hat{p}_{h_n}^{-1}[c_{\mathcal{X},\alpha}, \infty) \subset \text{supp}(P_{h_n}).$$

(ii) *Let $\delta_m \in (0, 1)$ and suppose $mh_m^d \rightarrow \infty$, $\delta_m^{-1} = O \left(\left(\frac{\log m}{mh_m^d} \right)^{\frac{4+d}{4+2d}} \right)$ and $mh_m^{-d} \rho_m^2 \delta_m^{-2} \rightarrow 0$. Then with probability $1 - \alpha - 4\delta_m$,*

$$q_{h_m}^{-1}[2c_{Q,\alpha-3\delta_m}, \infty) \subset \hat{q}_{h_m}^{-1}[c_{\mathcal{Y},\alpha}, \infty) \subset \text{supp}(Q_{h_m}).$$

Proof. (i)

Lemma E.3 (iii) implies that with probability $1 - \alpha - 4\delta_n$, $\|\hat{p}_h - p_h\| < c_{\mathcal{X},\alpha} \leq c_{P,\alpha-3\delta_n}$. This implies

$$p_{h_n}^{-1}[2c_{P,\alpha-3\delta_n}, \infty) \subset \hat{p}_{h_n}^{-1}[c_{\mathcal{X},\alpha}, \infty) \subset \text{supp}(P_{h_n}).$$

(ii)

This can be proven similarly to (i). □

Claim E.5. For a nonnegative measure μ and sets A, B, C, D ,

$$\mu(A \cap B) - \mu(C \cap D) \leq \mu(A \setminus C) + \mu(B \setminus D).$$

Proof.

$$\begin{aligned} \mu(A \cap B) - \mu(C \cap D) &\leq \mu((A \cap B) \setminus (C \cap D)) = \mu((A \cap B) \cap (C^c \cup D^c)) \\ &= \mu((A \cap B) \cap C^c) \cup (A \cap B) \cap D^c) \\ &\leq \mu((A \cap B) \setminus C) + \mu(A \cap B \setminus D) \\ &\leq \mu(A \setminus C) + \mu(B \setminus D). \end{aligned}$$

□

From here, let P_n and Q_m be the empirical measures on \mathcal{X} and \mathcal{Y} , respectively, i.e., $P_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$ and $Q_m = \frac{1}{m} \sum_{j=1}^m \delta_{Y_j}$.

Lemma E.6. *Suppose Assumption 1, 2 hold. Let $A \subset \mathbb{R}^d$. Then with probability $1 - \delta$,*

$$|P_n - P|(A) \leq \rho_n + \sqrt{\frac{\log(2/\delta)}{2n}},$$

and in particular,

$$P_n(A) \leq P(A) + \rho_n + \sqrt{\frac{\log(2/\delta)}{n}}.$$

Proof. Let P_n^0 be the empirical measure on \mathcal{X}^0 , i.e., $P_n^0 = \frac{1}{n} \sum_{i=1}^n \delta_{X_i^0}$. By using Hoeffding's inequality,

$$\mathbb{P}(|(P_n^0 - P)(A)| \geq t) \leq 2 \exp(-2nt^2),$$

and hence with probability $1 - \delta$,

$$|P_n^0 - P|(A) \leq \sqrt{\frac{\log(2/\delta)}{2n}}.$$

And $|P_n - P_n^0|(A)$ is expanded as

$$|P_n - P_n^0|(A) = \frac{1}{n} \sum_{i=1}^n |I(X_i \in A) - I(X_i^0 \in A)|.$$

Under Assumption 2, $\sum_{i=1}^n I(X_i \neq X_i^0) \leq n\rho_n$, and hence

$$\begin{aligned} |P_n - P_n^0|(A) &= \frac{1}{n} \sum_{i=1}^n |I(X_i \in A) - I(X_i^0 \in A)| \\ &\leq \frac{1}{n} \sum_{i=1}^n I(X_i \neq X_i^0) = \rho_n. \end{aligned}$$

Therefore, with probability $1 - \delta$,

$$|P_n - P|(A) \leq \rho_n + \sqrt{\frac{\log(2/\delta)}{2n}}.$$

□

Claim E.7. Suppose Assumption 1, 2, A1, A2 hold, and let $\delta_n, \delta_m \in (0, 1)$. Suppose $nh_n^d \rightarrow \infty$, $mh_m^d \rightarrow \infty$, $\delta_n^{-1} = O\left(\left(\frac{\log n}{nh_n^d}\right)^{\frac{4+d}{4+2d}}\right)$, $\delta_m^{-1} = O\left(\left(\frac{\log m}{mh_m^d}\right)^{\frac{4+d}{4+2d}}\right)$, $nh_n^{-d}\rho_n^2\delta_n^{-2} \rightarrow 0$, and $nh_m^{-d}\rho_m^2\delta_m^{-2} \rightarrow 0$. Let

$$B_{n,m} := (\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \cup q_{h_m}^{-1}(0, 2c_Q) \cup \text{supp}(P_{h_n}) \setminus \text{supp}(P). \quad (15)$$

(i) With probability $1 - 2\alpha - 4\delta_n - 8\delta_m$,

$$\begin{aligned} & |Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q))| \\ & \leq C \left(Q(B_{n,m}) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right). \end{aligned}$$

(ii) With probability $1 - 2\alpha - 8\delta_m$,

$$|Q_m(\hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(Q))| \leq C \left(Q(q_{h_m}^{-1}(0, 2c_Q)) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right).$$

(iii) With probability $1 - 2\alpha - 4\delta_n - 9\delta_m$,

$$\begin{aligned} & |Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q(\text{supp}(P))| \\ & \leq C \left(Q(B_{n,m}) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right). \end{aligned}$$

(iv) With probability $1 - 2\alpha - 9\delta_m$,

$$|Q_m(\hat{q}_{h_m}^{-1}[c_Y, \infty)) - 1| \leq C \left(Q(q_{h_m}^{-1}(0, 2c_Q)) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right).$$

(v) As $n, m \rightarrow \infty$, $B_{n,m} \rightarrow \emptyset$. And in particular,

$$Q(B_{n,m}) \rightarrow 0.$$

Proof. (i)

From Corollary E.4 (i) and (ii), with probability $1 - 2\alpha - 4(\delta_n + \delta_m)$,

$$\begin{aligned} Q_m(p_{h_n}^{-1}[2c_P, \infty) \cap q_{h_m}^{-1}[2c_Q, \infty)) & \leq Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) \\ & \leq Q_m(\text{supp}(P_{h_n}) \cap \text{supp}(Q_{h_m})). \end{aligned} \quad (16)$$

Then from the first inequality of (16), combining with Claim E.5 gives

$$\begin{aligned} & Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q)) \\ & \geq Q_m(p_{h_n}^{-1}[2c_P, \infty) \cap q_{h_m}^{-1}[2c_Q, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q)) \\ & \geq - (Q_m(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) + Q_m(\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty))). \end{aligned} \quad (17)$$

And from the second inequality of (16), combining with Claim E.5 gives

$$\begin{aligned} & Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q)) \\ & \leq Q_m(\text{supp}(P_{h_n}) \cap \text{supp}(Q_{h_m})) - Q_m(\text{supp}(P) \cap \text{supp}(Q)) \\ & \leq Q_m(\text{supp}(P_{h_n}) \setminus \text{supp}(P)) + Q_m(\text{supp}(Q_{h_m}) \setminus \text{supp}(Q)). \end{aligned} \quad (18)$$

And hence combining (17) and (18) gives that, with probability $1 - 2\alpha - 4(\delta_n + \delta_m)$,

$$\begin{aligned} & |Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q))| \\ & \leq \max \{ Q_m(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) + Q_m(\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty)), \\ & \quad Q_m(\text{supp}(P_{h_n}) \setminus \text{supp}(P)) + Q_m(\text{supp}(Q_{h_m}) \setminus \text{supp}(Q)) \}. \end{aligned} \quad (19)$$

Now we further bound the upper bound of (19). From Lemma E.6, with probability $1 - \delta_m$,

$$Q_m(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \leq Q(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) + \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}. \quad (20)$$

And similarly, with probability $1 - \delta_m$,

$$\begin{aligned} Q_m(\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty)) &\leq Q(\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty)) + \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}} \\ &= Q(q_{h_m}^{-1}(0, \infty)) + \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}. \end{aligned} \quad (21)$$

And similarly, with probability $1 - \delta_m$,

$$Q_m(\text{supp}(P_{h_n}) \setminus \text{supp}(P)) \leq Q(\text{supp}(P_{h_n}) \setminus \text{supp}(P)) + \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}. \quad (22)$$

And similarly, with probability $1 - \delta_m$,

$$\begin{aligned} Q_m(\text{supp}(Q_{h_m}) \setminus \text{supp}(Q)) &\leq Q(\text{supp}(Q_{h_m}) \setminus \text{supp}(Q)) + \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}} \\ &= \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}. \end{aligned} \quad (23)$$

Hence by applying (20), (21), (22), (23) to (19), with probability $1 - 2\alpha - 4\delta_n - 8\delta_m$,

$$\begin{aligned} &|Q_m(\hat{p}_{h_n}^{-1}[c_X, \infty) \cap \hat{q}_{h_m}^{-1}[c_Y, \infty)) - Q_m(\text{supp}(P) \cap \text{supp}(Q))| \\ &\leq C \left(Q((\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \cup q_{h_m}^{-1}(0, 2c_Q) \cup \text{supp}(P_{h_n}) \setminus \text{supp}(P)) \right. \\ &\quad \left. + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right) \\ &= C \left(Q(B_{n,m}) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right), \end{aligned}$$

where $B_{n,m}$ is from (15).

(ii)

This can be done similarly to (i).

(iii)

Lemma E.6 gives that with probability $1 - \delta_m$,

$$|Q_m(\text{supp}(P) \cap \text{supp}(Q)) - Q(\text{supp}(P))| \leq \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}.$$

Hence combining this with (i) gives the desired result.

(iv)

Lemma E.6 gives that with probability $1 - \delta_m$,

$$|Q_m(\text{supp}(Q)) - 1| \leq \rho_m + \sqrt{\frac{\log(2/\delta)}{2m}}.$$

Hence combining this with (ii) gives the desired result.

(v)

Note that Lemma D.3 implies that for all $x \in \text{supp}(P)$, $\liminf_{n \rightarrow \infty} p_{h_n}(x) > 0$, so $p_{h_n}(x) > 2c_P$ for large enough n . And hence as $n \rightarrow \infty$,

$$\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty) \rightarrow \emptyset. \quad (24)$$

And similar argument holds for $\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty)$, so as $m \rightarrow \infty$,

$$\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty) \rightarrow \emptyset. \quad (25)$$

Also, since K has compact support, for any $x \notin \text{supp}(P)$, $x \notin \text{supp}(P_{h_n})$ once $h_n < d(x, \text{supp}(P))$. Hence as $n \rightarrow \infty$,

$$\text{supp}(P_{h_n}) \setminus \text{supp}(P) \rightarrow \emptyset. \quad (26)$$

And similarly, as $m \rightarrow \infty$,

$$\text{supp}(Q_{h_m}) \setminus \text{supp}(Q) \rightarrow \emptyset. \quad (27)$$

Hence by applying (24), (25), (26), (27) to the definition of $B_{n,m}$ in (15) gives that as $n, m \rightarrow \infty$,

$$B_{n,m} \rightarrow \emptyset.$$

And in particular,

$$Q(B_{n,m}) \rightarrow 0.$$

□

Below we state a more formal version of Proposition 4.1.

Proposition E.8. *Suppose Assumption 1,2,A1,A2 hold. Suppose $\alpha \rightarrow 0$, $\delta_n \rightarrow 0$, $\delta_n^{-1} = O\left(\left(\frac{\log n}{nh_n^d}\right)^{\frac{4+d}{4+2d}}\right)$, $h_n \rightarrow 0$, $nh_n^d \rightarrow \infty$, and $nh_n^{-d}\rho_n^2\delta_n^{-2} \rightarrow 0$, and similar relations hold for h_m, ρ_m . Then there exists some constant $C > 0$ not depending on anything else such that*

$$\begin{aligned} |\text{TopP}_{\mathcal{X}}(\mathcal{Y}) - \text{precision}_P(\mathcal{Y})| &\leq C \left(Q(B_{n,m}) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right), \\ |\text{TopR}_{\mathcal{Y}}(\mathcal{X}) - \text{recall}_Q(\mathcal{X})| &\leq C \left(P(A_{n,m}) + \rho_n + \sqrt{\frac{\log(1/\delta)}{n}} \right), \end{aligned}$$

where

$$\begin{aligned} A_{n,m} &:= (\text{supp}(Q) \setminus q_{h_m}^{-1}[2c_Q, \infty)) \cup p_{h_n}^{-1}(0, 2c_P) \cup \text{supp}(Q_{h_m}) \setminus \text{supp}(Q), \\ B_{n,m} &:= (\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \cup q_{h_m}^{-1}(0, 2c_Q) \cup \text{supp}(P_{h_n}) \setminus \text{supp}(P). \end{aligned}$$

And as $n, m \rightarrow \infty$, $P(A_{n,m}) \rightarrow 0$ and $Q(B_{n,m}) \rightarrow 0$ hold.

Proof of Proposition E.8. Now this is an application of Claim E.7 (i) (ii) (v) to the definitions of $\text{precision}_P(\mathcal{Y})$ and $\text{recall}_Q(\mathcal{X})$ in (1) and (2) and the definitions of $\text{TopP}_{\mathcal{X}}(\mathcal{Y})$ and $\text{TopR}_{\mathcal{Y}}(\mathcal{X})$ in (3) and (4).

□

Similarly, we state a more formal version of Theorem 4.2.

Theorem E.9. *Suppose Assumption 1,2,A1,A2 hold. Suppose $\alpha \rightarrow 0$, $\delta_n \rightarrow 0$, $\delta_n^{-1} = O\left(\left(\frac{\log n}{nh_n^d}\right)^{\frac{4+d}{4+2d}}\right)$, $h_n \rightarrow 0$, $nh_n^d \rightarrow \infty$, and $nh_n^{-d}\rho_n^2\delta_n^{-2} \rightarrow 0$, and similar relations hold for h_m, ρ_m . Then there exists some constant $C > 0$ not depending on anything else such that*

$$\begin{aligned} |\text{TopP}_{\mathcal{X}}(\mathcal{Y}) - \text{precision}_P(Q)| &\leq C \left(Q(B_{n,m}) + \rho_m + \sqrt{\frac{\log(1/\delta)}{m}} \right), \\ |\text{TopR}_{\mathcal{Y}}(\mathcal{X}) - \text{recall}_Q(P)| &\leq C \left(P(A_{n,m}) + \rho_n + \sqrt{\frac{\log(1/\delta)}{n}} \right), \end{aligned}$$

where $A_{n,m}$ and $B_{n,m}$ are the same as in Proposition E.8. Again as $n, m \rightarrow \infty$, $P(A_{n,m}) \rightarrow 0$ and $Q(B_{n,m}) \rightarrow 0$ hold.

Proof of Theorem E.9. Now this is an application of Claim E.7 (iii) (iv) (v) to the definitions of $\text{precision}_P(Q)$ and $\text{recall}_Q(P)$ and the definitions of $\text{TopP}_{\mathcal{X}}(\mathcal{Y})$ and $\text{TopR}_{\mathcal{Y}}(\mathcal{X})$ in (3) and (4). \square

Proof of Lemma 4.3. Recall that $B_{n,m}$ is defined in (15) as

$$B_{n,m} = (\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \cup q_{h_m}^{-1}(0, 2c_Q) \cup \text{supp}(P_{h_n}) \setminus \text{supp}(P),$$

and hence $Q(B_{n,m})$ can be upper bounded as

$$Q(B_{n,m}) \leq Q(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) + Q(q_{h_m}^{-1}(0, 2c_Q)) + Q(\text{supp}(P_{h_n}) \setminus \text{supp}(P)). \quad (28)$$

For the first term of RHS of (28), suppose $\text{supp}(K) \subset \mathcal{B}_{\mathbb{R}^d}(0, 1)$ for convenience, and let $A_{-h_n} := \{x \in \text{supp}(P) : d(x, \mathbb{R}^d \setminus \text{supp}(P)) \geq h_n\}$. Then from Assumption A3, for all $x \in A_{-h_n}$, $p_{h_n}(x) \geq p_{\min}$ holds. Hence for large enough N such that for $n \geq N$, $c_P < p_{\min}$, then $p_{h_n}(x) \geq p_{\min}$ for all $x \in A_{-h_n}$, and hence

$$Q(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) \leq Q(\text{supp}(P) \setminus A_{-h_n}). \quad (29)$$

Also, $\text{supp}(P)$ being bounded implies that all the coefficients a_0, \dots, a_{d-1} in Proposition D.5 for $\text{supp}(P)$ are in fact finite. Then $q \leq q_{\max}$ from Assumption A3 and Proposition D.5 implies

$$Q(\text{supp}(P) \setminus A_{-h_n}) = O(h_n). \quad (30)$$

And hence combining (29) and (30) gives that

$$Q(\text{supp}(P) \setminus p_{h_n}^{-1}[2c_P, \infty)) = O(h_n). \quad (31)$$

For the second term of RHS of (28), with a similar argument,

$$Q(q_{h_m}^{-1}(0, 2c_Q)) = O(h_m). \quad (32)$$

Finally, for the third term of RHS of (28), Proposition D.5 implies

$$Q(\text{supp}(P_{h_n}) \setminus \text{supp}(P)) = O(h_n). \quad (33)$$

Hence applying (31), (32), (33) to (28) gives that

$$Q(B_{n,m}) = O(h_n + h_m). \quad \square$$

F Related Work

Improved Precision & Recall (P&R). Existing metrics such as IS and FID assess the performance of generative models with a single score while showing usefulness in determining performance rankings between models and are still widely used. These evaluation metrics have problems that they cannot provide detailed interpretations of the evaluation in terms of fidelity and diversity. Sajjadi et al. [12] tried to solve this problem by introducing the original Precision and Recall, which however has a limitation as it is inaccurate due to the simultaneous approximation of real support and fake support through “ k -means clustering”. For example, if we evaluate the fake images having high fidelity and large value of k , many fake features can be classified in a small cluster with no real features, resulting in a low fidelity score ($\text{precision}_P := Q(\text{supp}(P))$). Kynkäänniemi et al. [13] focuses on the limitation from the support estimation and presents an P&R that allows us to more accurately assess fidelity and diversity by approximating real support and fake support separately:

$$\text{precision}(\mathcal{X}, \mathcal{Y}) := \frac{1}{M} \sum_i^M f(Y_i, \mathcal{X}), \quad \text{recall}(\mathcal{X}, \mathcal{Y}) := \frac{1}{N} \sum_j^N f(X_j, \mathcal{Y}),$$

$$\text{where } f(Y_i, \mathcal{X}) = \begin{cases} 1, & \text{if } Y_i \in \text{supp}(P), \\ 0, & \text{otherwise.} \end{cases}$$

Here, $\text{supp}(P)$ is defined as the union of spheres centered on each real feature X_i and whose radius is the distance between k th nearest real features of X_i

Density & Coverage (D&C). Naeem et al. [14] has reported that P&R cannot stably provide accurate fidelity and diversity when real or fake features are present through experiment. This limitation of P&R is that it has a problem with approximating overestimated supports by outliers due to the use of " k -nearest neighborhood" algorithm. For a metric that is robust to outlying features, Naeem et al. [14] proposes a new evaluation metric that only relies on the real support, based on the fact that a generative model often generates artifacts which possibly results in outlying features in the embedding space:

$$density(\mathcal{X}, \mathcal{Y}) := \frac{1}{M} \sum_j^M \sum_i^N 1_{Y_j \in f(X_i)}, \quad coverage(\mathcal{X}, \mathcal{Y}) := \frac{1}{N} \sum_i^N 1_{\exists j \text{ s.t. } Y_j \in f(X_i)},$$

$$\text{where } f(X_i) = \mathcal{B}(X_i, NND_k(X_i))$$

In the equation, $NND_k(X_i)$ means the k th-nearest neighborhood distance of X_i and $\mathcal{B}(X_i, NND_k(X_i))$ denotes the hyper-sphere centered at X_i with radius $NND_k(X_i)$. However, D&C is only a partial solution because it still gives an inaccurate evaluation when a real outlying feature exists. Unlike coverage, density is not upper-bounded which makes it unclear exactly which ideal score a generative model should achieve.

Geometric Evaluation of Data Representations (GCA). Poklukar et al. [27] proposes a metric called GCM that assesses the fidelity and diversity of fake images. GCM uses the geometric and topological properties of connected components formed by vertex \mathcal{V} (feature) and edge \mathcal{E} (connection). Briefly, when the pairwise distance between vertices is less than a certain threshold ϵ , a connected component is formed by connecting two vertices with an edge. This set of connected components can be thought of as a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and each connected component separated from each other in \mathcal{G} is defined as a subgraph \mathcal{G}_i (i.e., $\mathcal{G} = \cup_i \mathcal{G}_i$). GCA uses the consistency $c(\mathcal{G}_i)$ and the quality $q(\mathcal{G}_i)$ to select the important subgraph \mathcal{G}_i that is formed on both real vertices and fake vertices ($\mathcal{V} = \mathcal{X} \cup \mathcal{Y}$). $c(\mathcal{G}_i)$ evaluates the ratio of the number of real vertices and fake vertices and $q(\mathcal{G}_i)$ computes the ratio of the number of edges connecting real vertices and fake vertices to the total number of edges in \mathcal{G}_i . Given the consistency threshold η_c and quality threshold η_q , the set of important subgraphs is defined as $\mathcal{S}(\eta_c, \eta_q) = \cup_{q(\mathcal{G}_i) > \eta_q, c(\mathcal{G}_i) > \eta_c} \mathcal{G}_i$. Based on this, GCA precision and recall are defined as follows:

$$\text{GCA precision} = \frac{|\mathcal{S}^{\mathcal{Y}}|_{\mathcal{V}}}{|\mathcal{G}^{\mathcal{Y}}|_{\mathcal{V}}}, \quad \text{GCA recall} = \frac{|\mathcal{S}^{\mathcal{X}}|_{\mathcal{V}}}{|\mathcal{G}^{\mathcal{X}}|_{\mathcal{V}}},$$

In above, $\mathcal{S}^{\mathcal{Y}}$ and $\mathcal{G}^{\mathcal{Y}}$ represent a subset ($\mathcal{V} = \mathcal{Y}, \mathcal{E}$) of each graph \mathcal{S} and \mathcal{G} , respectively. One of the drawbacks is that the new hyper-parameters ϵ , η_c , and η_q need to be set arbitrarily according to the image dataset in order to evaluate fake images well. In addition, even if we use the hyper-parameters that seem most appropriate, it is difficult to verify that the actual evaluation results are correct because they do not approximate the underlying feature distribution.

Manifold Topology Divergence (MTD). Barannikov et al. [28] proposes a new metric that uses the life-length of the k -dimensional homology of connected components between real features and fake features formed through Vietoris-Rips filtration [41]. MTD is simply defined as the sum of the life-length of homologies and is characterized by an evaluation trend consistent with the FID. Specifically, MTD constructs the graph \mathcal{G} by connecting edges between features with a distance smaller than the threshold ϵ . Then the birth and death of the k -dimensional homologies in \mathcal{G} are recorded by adjusting the threshold ϵ from 0 to ∞ . The life-length $l_i(h)$ is obtained through $death_i - birth_i$ of any k -dimensional homology h , and let life-length set of all homologies as $L(h)$. MTD repeats the process of obtaining $L(h)$ on randomly sampled subsets $\mathcal{X}' \in \mathcal{X}$ and $\mathcal{Y}' \in \mathcal{Y}$ at each iteration, and defined as the following:

$$MTD(\mathcal{X}, \mathcal{Y}) = \frac{1}{n} \sum_i^n L_i(h), \quad \text{where } L_i(h) \text{ is the life-length set defined at } i\text{th iteration.}$$

However, MTD has a limitation of only considering a fixed dimensional homology for evaluation. In addition, MTD lacks interpretability as it evaluates the model with a single assessment score.

G Philosophy of our metric & Practical Scenarios

G.1 Philosophy of our metric

All evaluation metrics have different resolutions and properties. The philosophy of our metric is to propose a reliable evaluation metric based on statistically and topologically certain things. In a real situation, the data often contain outliers or noise, and these data points come from a variety of sources (e.g. human error, feature embedding network). These errors may play as outliers, which leads to an overestimation of the data distribution, which in turn leads to a false impression of improvement when developing generative models. As discussed in Section 2, P&R and its variants have different ways to estimate the supports, which overlook the possible presence and effect of noise. Naeem et al. [14] revealed that previous support estimation approaches may inaccurately estimate the true support under noisy conditions, and partially solved this problem by proposing to only use the real support. However, this change goes beyond the natural definition of precision and recall, and results in losing some beneficial properties like boundedness. Moreover, this is still a temporary solution since real data can also contain outliers. We propose a solution to the existing problem by minimizing the effects of noise, using topologically significant data points and retaining the natural definition of precision and recall.

G.2 Practical Scenarios

From this perspective, we present two examples of realistic situations where outliers exist in the data and filtering out them can have a significant impact on proper model analysis and evaluation. With real data, there are many cases where outliers are introduced into the data due to human error [17, 18]. Taking the simplest MNIST as an example, suppose our task of interest is to generate 4. Since image number 7 is included in data set number 4 due to incorrect labeling (see Figure 1 of [17]), the support of the real data in the feature space can be overestimated by such outliers, leading to an unfair evaluation of generative models (as in Section 5.1.2 and 5.2.2); That is, the sample generated with weird noise may be in the overestimated support, and existing metrics without taking into account the reliability of the support could not penalize this, giving a good score to a poorly performing generator.

A similar but different example is when noise or distortions in the captured data (unfortunately) behave adversely on the feature embedding network used by the current evaluation metrics (as in Section 5.1.2 and 5.2.2); e.g., visually it is the number 7, but it is mismatched near the feature space where there are usually 4 and becomes an outlier. Then the same problem as above may occur. Note that in these simple cases, where the definition of outliers is obvious with enough data, one could easily examine the data and exclude outliers a priori to train a generative model. In the case of more complicated problems such as the medical field [18], however, it is often not clear how outliers are to be defined. Moreover, because data are often scarce, even outliers are very useful and valuable in practice for training models and extracting features, making it difficult to filter outliers in advance and decide not to use them.

On the other hand, we also provide an example where it is very important to filter out outliers in the generator sample and then evaluate them. To evaluate the generator, samples are generated by sampling from the preset latent space (typically Gaussian). Even after training is complete and the generators' outputs are generally fine, there's a latent area where generators aren't fully trained. Note that latent space sampling may contain samples from regions that the generator does not cover well during training ("unfortunate outlier"). When unfortunate outliers are included, the existing evaluation metrics may underestimate or overestimate the generator's performance than its general performance. (To get around this, it is necessary to try this evaluation several times to statistically stabilize it, but this requires a lot of computation and becomes impractical, especially when the latent space dimension is high.)

Especially considering the evaluation scenario in the middle of training, the above situation is likely to occur due to frequent evaluation, which can interrupt training or lead to wrong conclusions. On the other hand, we can expect that our metric will be more robust against the above problem since it pays more attention to the core (samples that form topologically meaningful structures) generation performance of the model.

G.3 Details of the noise framework in the experiments

We have assessed the robustness of TopP&R against two types of non-IID noise perturbation through toy data experiments (in Section 5.1.3) and real data experiments (in Section 5.2.2). The scatter noise we employed consists of randomly extracted noise from a uniform distribution. This noise, even when extensively added to our data, does not form a data structure with any significant signal (topological feature). In other words, from the perspective of TopP&R, the formation of topologically significant data structures implies that data samples should possess meaningful probability values in the feature space. However, noise following a uniform distribution across all feature dimensions holds very small probability values, and thus, it does not constitute a topologically significant data structure. Consequently, such noise is filtered out by the bootstrap bands c_X or c_Y that we approximate (see Section 2).

The alternate noise we utilized in our experiments, swap noise, possesses distinct characteristics from scatter noise. Initially, given the real and fake data that constitute important data structures, we introduce swap noise by randomly selecting real and fake samples and exchanging their positions. In this process, the swapped fake samples follow the distribution of real data, and conversely, the real samples adhere to the distribution of fake data. This implies the addition of noise that follows the actual data distribution, thereby generating significant data structures. When the number of samples undergoing such positional swaps remains small, meaningful probability values cannot be established within the distribution. Statistically, an extremely small number of samples cannot form a probability distribution itself. As a result, our metric operates robustly in the presence of noise. However, as the count of these samples increases and statistically significant data structures emerge, TopP&R estimates the precise support of such noise as part of their distribution. By examining Figure 4 and 5, as well as Table A12, it becomes evident that conventional metrics struggle with accurate support estimation due to vulnerability to noise. This limitation results in an inability to appropriately evaluate aspects involving minority sets or data forming long distributions.

G.4 Limitations

Since the KDE filtration requires extensive computations in high dimensions, a random projection that preserves high-dimensional distance and topological properties is inevitable. Based on the topological structure of the features present in the embedded low dimensional space, we have shown that TopP&R theoretically and experimentally has several good properties such as robustness to the noise from various sources and sensitivity to small changes in distribution. However, matching the distortion from the random projection to the noise level allowed by the confidence band is practically infeasible: The bootstrap confidence band is of order $(nh_n^d)^{-\frac{1}{2}}$, and hence for the distortion from the random projection to match this confidence band, the embedded dimension d should be of order $\Omega(nh_n^d \log n)$ from Remark (B.2). However, computing the KDE filtration in dimension $\Omega(nh_n^d \log n)$ is practically infeasible. Hence in practice, the topological distortion from the random projection is not guaranteed to be filtered out by the confidence band, and there is a possibility of obtaining a less accurate evaluation score compared to calculations in the original dimensions.

Exploring the avenue of localizing uncertainty at individual data points separately presents an intriguing direction of research. Such an approach could potentially be more sample-efficient and disregard less data points. However, considering our emphasis on preserving topological signals, achieving localized uncertainty estimation in this manner is challenging within the current state of the art. For instance, localizing uncertainty in the kernel density estimate \hat{p} is feasible, as functional variability is inherently local. To control uncertainty at a specific point x , we primarily need to analyze the function value $\hat{p}(x)$ at that point. In contrast, localizing uncertainty for homological features situated at point x requires the analysis of all points connected by the homological feature. Consequently, even if our intention is to confine uncertainty to a local point, it demands estimating the uncertainty at the global structural level. This makes localizing the uncertainty for topological features difficult given the tools in topology and statistics we currently have.

H Experimental Details

H.1 Implementation details of embedding

We summarize the detailed information of our embedding networks implemented for the experiments. In Figure 2, 3, 4, A10, A4, and 5, P&R and D&C are computed from the features of ImageNet pre-trained VGG16 (fc2 layer), and TopP&R is computed from features placed in \mathbb{R}^{32} with additional random linear projection. In the experiment in Figure A10, the SwAV embedder is additionally considered. We implement ImageNet pre-trained InceptionV3 (fc layer), VGG16 (fc2 layer), and SwAV as embedding networks with random linear projection to 32 dimensional feature space to compare the ranking of GANs in Table 1.

H.2 Implementation details of confidence band estimator

Algorithm 1 Confidence Band Estimator

<pre> # KDE: kernel density estimator # h: kernel bandwidth parameter # k: number of repeats # $\hat{\theta}$: set of difference Given $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ $\hat{p}_h = KDE(\mathcal{X})$ for iteration = 1, 2, ..., k do # Define \mathcal{X}^* with bootstrap sampling \mathcal{X}^* = random sample n times with repeat from \mathcal{X} # \hat{p}_h^* replaces population density $\hat{p}_h^* = KDE(\mathcal{X}^*)$ # compute $\hat{\theta}$ with bootstrap samples Append $\hat{\theta}$ with $\sqrt{n} \ \hat{p}_h - \hat{p}_h^*\ _\infty$ end for </pre>	<pre> # grid search for the confidence band for $q \in [\min(\hat{\theta}), \max(\hat{\theta})]$ do count = 0 for element $\in \hat{\theta}$ do if element > q then count = count + 1 end if end for # define the band threshold if count/k $\approx \alpha$ then $q_\alpha = q$ end if end for # define estimated confidence band $c_\alpha = q_\alpha / \sqrt{n}$ </pre>
---	---

H.3 Choice of confidence level

For the confidence level α , we would like to point out that α is not the usual hyperparameter to be tuned: It has a statistical interpretation of the probability or the level of confidence to allow error, noise, etc. The most popular choices are $\alpha = 0.1, 0.05, 0.01$, leading to 90%, 95%, 99% confidence. We used $\alpha = 0.1$ throughout our experiments.

H.4 Estimation of Bandwidth parameter

As we discussed in Section 2, since TopP&R estimates the manifold through KDE with kernel bandwidth parameter h , we need to approximate it. The estimation techniques for h are as follows: (a) a method of selecting h that maximizes the survival time ($S(h)$) or the number of significant homological features ($N(h)$) based on information obtained about persistent homology using the filtration method, (b) a method using the median of the k-nearest neighboring distances between features obtained by the balloon estimator (for more details, please refer to [42], [43], and [44]). Note that, the bandwidths h for all the experiments in this paper are estimated via Balloon Bandwidth Estimator.

For (a), following the notation in Section A, let the i th homological feature of persistent diagram be (b_i, d_i) , then we define its life length as $l_i(h) = d_i - b_i$ at kernel bandwidth h . With confidence band $c_\alpha(h)$, we select h that maximizes one of the following two quantities:

$$N(h) = \#\{i : l_i(h) > c_\alpha(h)\}, \quad S(h) = \sum_i [l_i(h) - c_\alpha(h)]_+.$$

Note that, we denote the confidence band c_α as $c_\alpha(h)$ considering the kernel bandwidth parameter h of KDE in Algorithm 1.

For (b), the balloon bandwidth estimator is defined as below:

Algorithm 2 Balloon Bandwidth Estimator

```

# h: kernel bandwidth
# KND: kth nearest distance
# idx: index
# sort: sort in ascending order

Given  $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ 
for  $idx = 1, 2, \dots, n$  do
  # Compute L2 distance between  $X_{idx}$  and  $\mathcal{X}$ 
   $d(X_{idx}, \mathcal{X}) = \|X_{idx} - \mathcal{X}\|_2^2$ ,
  i.e.,  $\{d(X_{idx}, X_1), d(X_{idx}, X_2), \dots\}$ 
  # Calculate kth nearest neighbor distance
   $KND_{idx} = \text{sort}(d(X_{idx}, \mathcal{X}))[k]$ 
end for
Given  $KND = \{KND_1, KND_2, \dots, KND_n\}$ 
# Define the estimated bandwidth  $\hat{h}$ 
 $\hat{h} = \text{median}(KND)$ 

```

H.5 Computational complexity

TopP&R is computed using the confidence band c_α and KDE bandwidth h . The computational cost of the confidence band calculation (Algorithm 1) and the balloon bandwidth calculation (Algorithm 2) are $O(k * n^2 * d)$ and $O(n^2 * d)$, respectively, where n represents the data size, d denotes the data dimension, and k stands for the number of repeats ($k = 10$ in our implementation). The resulting computational cost of TopP&R is $O(k * n^2 * d)$. In our experiments based on real data, calculating TopP&R once using real and fake features, each consisting of 10k samples in a 4096-dimensional space, takes approximately 3-4 minutes. This computation speed is notably comparable to that of P&R and D&C, utilizing CPU-based computations. Note that, P&R and D&C that we primarily compare in our experiments are algorithms based on k -nearest neighborhood for estimating data distribution, and due to the computation of pairwise distances, the computational cost of these algorithm is approximately $O(n^2 * d)$.

Table A1: Wall clock, CPU time and Time complexity for evaluation metrics for 10k real and 10k fake features in 4096 dimension. For TopP&R, r is random projection dimension ($r \ll d$). The results are measured by the time module in python.

	Time	PR	DC	TopPR (Ours)
Wall Clock		1min 43s	1min 40s	1min 58s
CPU time		2min 44s	2min 34s	2min 21s
Time complexity		$O(n^2 * d)$	$O(n^2 * d)$	$O(k * n^2 * r)$

H.6 Mean hamming distance

Hamming distance (HD) [45] counts the number of items with different ranks between A and B, then measures how much proportion differs in the overall order, i.e. for $A_i \in A$ and $B_i \in B$, $HD(A, B) := \sum_{i=1}^n 1\{k : A_i \neq B_i\}$ where k is the number of differently ordered elements. The **mean HD** is calculated as follows to measure the average distances of three ordered lists: Given three ordered lists A, B , and C , $\bar{HD} = (HD(A, B) + HD(A, C) + HD(B, C))/3$.

H.7 Explicit values of bandwidth parameter

Since our metric adaptively reacts to the given samples of P and Q , we have two h 's per experiment. For example, in the translation experiment (Figure 2), there are 13 steps in total, and each time we estimate h for P and Q , resulting in a total of 26 h 's. To show them all at a glance, we have listed all values in one place. We will also provide the code that can reproduce the results in our experiments upon acceptance.

Table A2: Bandwidth parameters h of distribution shift in Section 5.1.1.

μ	← shift												shift →
real h	8.79	8.60	8.58	8.67	8.63	8.71	8.66	8.73	8.54	8.63	8.54	8.75	8.78
fake h	8.46	8.86	8.60	8.49	8.44	8.80	8.55	8.63	8.74	8.61	8.58	8.83	8.60

Table A3: Bandwidth parameters h of sequential mode drop in Section 5.1.2.

Steps	0	1	2	3	4	5	6
real h	10.5	10.6	10.5	10.7	10.5	10.8	10.6
fake h	11.0	10.2	10.0	9.82	9.57	9.29	8.78

Table A4: Bandwidth parameters h of simultaneous mode drop in Section 5.1.2.

Steps	0	1	2	3	4	5	6	7	8	9	10
real h	10.5	10.6	10.6	10.8	10.3	10.6	10.7	10.6	10.4	10.4	10.6
fake h	10.6	10.5	10.7	10.3	10.4	10.0	9.88	9.31	9.19	9.07	8.79

Table A5: Bandwidth parameters h of non-IID noise perturbation in Section 5.1.3.

Steps		0	1	2	3	4	5	6	7	8	9
Scatter	real h	8.76	8.71	8.78	8.59	8.85	8.71	8.61	8.97	9.00	8.95
	fake h	8.61	8.76	8.63	8.93	8.46	8.72	8.67	8.68	9.07	8.88
Swap	real h	8.65	8.44	8.76	8.62	8.55	8.75	8.96	8.59	8.53	8.74
	fake h	8.52	8.55	8.68	8.97	8.87	8.94	8.84	9.05	8.94	8.90

Table A6: Bandwidth parameters h of FFHQ truncation trick in Section I.11.

Ψ	$\Psi \downarrow$										$\Psi \uparrow$
real h	6.26	6.41	6.31	6.23	6.06	6.24	6.20	6.61	6.39	6.28	6.26
fake h	1.04	1.67	2.36	2.90	3.41	3.93	4.31	4.90	5.51	5.60	6.46

Table A7: Bandwidth parameters h of CIFAR10 sequential mode drop in Section 5.2.1.

Steps	0	1	2	3	4	5	6	7	8	9
real h	7.09	6.95	6.96	6.92	7.02	6.81	6.86	6.88	6.78	6.67
fake h	6.87	6.75	6.88	6.64	6.40	6.64	6.41	6.44	6.62	6.04

Table A8: Bandwidth parameters h of CIFAR10 simultaneous mode drop in Section 5.2.1.

Steps	0	1	2	3	4	5	6	7	8	9	10
real h	6.78	6.95	7.05	6.62	6.85	6.51	6.94	6.91	6.80	7.00	6.78
fake h	6.75	6.69	7.14	6.68	6.75	6.84	6.61	6.68	6.50	6.32	6.17

Table A9: Bandwidth parameters h of FFHQ non-IID noise perturbation in Section 5.2.2.

Steps		0	1	2	3	4	5	6	7	8	9
Scatter	real h	7.24	7.03	7.01	7.45	7.38	7.66	7.46	7.36	7.79	7.66
	fake h	5.41	5.77	5.67	5.49	5.71	5.65	5.82	6.08	5.81	6.02
Swap	real h	7.25	7.20	6.99	6.67	7.07	6.96	7.03	6.94	6.83	6.86
	fake h	6.87	6.63	6.60	6.91	6.74	6.77	6.92	6.83	6.83	6.82

I Additional Experiments

I.1 Survivability of minority sets in the long-tailed distribution

An important point to check in our proposed metric is the possibility that a small part of the total data (i.e., minority sets), but containing important information, can be ignored by the confidence band. We emphasize that since our metric takes topological features into account, even minority sets are not filtered conditioned that they have topologically significant structures. We assume that signals or data

Table A10: Bandwidth parameters h of FFHQ noise addition in Section 5.2.3.

		Noise intensity	0	1	2	3
Gaussian noise	real h		6.52	6.27	6.41	6.49
	fake h		6.32	6.15	4.06	2.64
Gaussian blur	real h		6.08	6.26	6.26	6.47
	fake h		6.12	5.13	4.60	3.40
Black rectangle	real h		6.06	6.28	6.49	6.44
	fake h		6.36	6.66	6.31	5.93

Table A11: Bandwidth parameters h of GAN ranking in Section 5.2.4.

		StyleGAN2	ReACGAN	BigGAN	PDGAN	ACGAN	WGAN-GP
InceptionV3	real h	2.605	2.629	2.581	2.587	2.615	2.606
	fake h	2.226	2.494	2.226	2.639	2.573	2.063
VGG16	real h	7.732	8.021	8.175	7.792	7.845	8.036
	fake h	7.880	6.839	7.178	6.174	6.839	3.615
SwAV	real h	0.774	0.765	0.785	0.780	0.780	0.822
	fake h	0.740	0.667	0.6746	0.619	0.627	0.502

that are minority sets have topological structures, but outliers exist far apart and lack a topological structure in general.

To test this, we experimented with CIFAR10 [46], which has 5,000 samples per class. We simulate a dataset with the majority set of six classes (2,000 samples per class, 12,000 total) and the minority set of four classes (500 samples per class, 2,000 total), and an ideal generator that exactly mimics the full data distribution. As shown in the Table A12, the samples in the minority set remained after the filtering process, meaning that the samples were sufficient to form a significant structure. Both D&C and TopP&R successfully evaluate the distribution for the ideal generator. To check whether our metric reacts to the change in the distribution even with this harsh setting, we also carried out the mode decay experiment. We dropped the samples of the minority set from 500 to 100 per class, which can be interpreted as an 11.3% decrease in diversity relative to the full distribution (Given (a) ratio of the number of samples between majority and minority sets = 12,000 : 2,000 = 6 : 1 and (b) 80% decrease in samples per minority class, the true decay in the diversity is calculated as $\frac{1}{(1+6)} \times 0.8 = 11.3\%$ with respect to the entire samples). Here, recall and coverage react somewhat less sensitively with their reduced diversities as 3 p.p. and 2 p.p., respectively, while TopR reacted most similarly (9 p.p.) to the ideal value. In summary, TopP&R shows much more sensitiveness to the changes in data distribution like mode decay. Thus, once the minority set has survived the filtering process, our metric is likely to be much more responsive than existing methods.

Table A12: Proportion of surviving minority samples in the long-tailed distribution after the noise exclusion with confidence band $c_{\mathcal{X}}$. The p.p. indicates the percentage points.

	Before mode drop	After mode drop
Proportion of survival (before / after filtering)	100% \rightarrow 59%	100% \rightarrow 57%
TopP&R (Fidelity / Diversity)	0.99 / 0.96	0.99 / 0.87 (9 p.p. \downarrow)
P&R (Fidelity / Diversity)	0.73 / 0.73	0.74 / 0.70 (3 p.p. \downarrow)
D&C (Fidelity / Diversity)	0.99 / 0.96	1.02 / 0.94 (2 p.p. \downarrow)

I.2 Experiment on Non-IID perturbation with outlier removal methods

We compared the performance of metrics on a denoised dataset using Local Outlier Factor (LOF) [47] and Isolation Forest (IF) [48]. The experimental setup is identical to that of Figure 4 (see Section G.3 for details of our noise framework). In this experiment, we applied the outlier removal method before calculating P&R and D&C. From the result, P&R and D&C still do not provide a stable evaluation, while TopP&R shows the most consistent evaluation for the two types of noise without changing its trend.

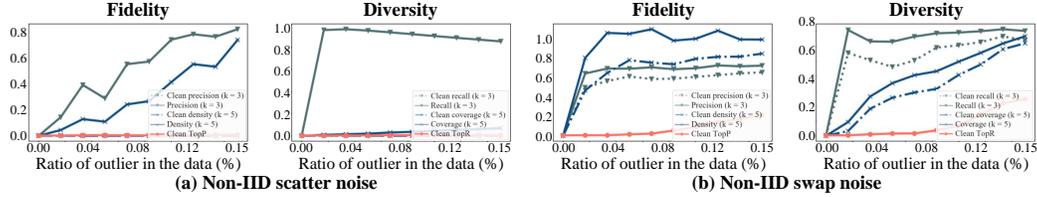


Figure A2: Behaviors of evaluation metrics on Non-IID perturbations. The dotted line in the graph shows the performance of metrics after the removal of outliers using the IF. We use "Clean" as a prefix to denote the evaluation after IF.

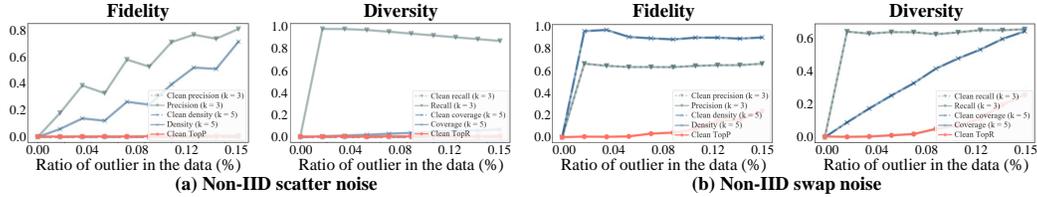


Figure A3: Behaviors of evaluation metrics on Non-IID perturbations. The dotted line in the graph shows the performance of metrics after the removal of outliers using the LOF. We use "Clean" as a prefix to denote the evaluation after LOF.

Since TopP&R can discriminate topologically significant signals by estimating the confidence band (using KDE filtration function), there is no need to arbitrarily set cutoff thresholds for each dataset. On the other hand, in order to distinguish inliers from specific datasets using LOF and IF methods, a threshold specific to the dataset must be arbitrarily set each time, so there is a clear constraint that different results are expected each time for each parameter setting (which is set by the user). Through this, it is confirmed that our evaluation metric guarantees the most consistent scoring based on the topologically significant signals. Since the removal of outliers in the TopP&R is part of the support estimation process, it is not appropriate to replace our unified process with the LOF or IF. In detail, TopP&R removes outliers through a clear threshold called confidence band which is defined by KDE, and this process simultaneously defines KDE's super-level set as the estimated support. If this process is separated, the topological properties and interpretations of estimated support also disappear. Therefore, it is not practical to remove outliers by other methods when calculating TopP&R.

I.3 Sequential and simultaneous mode dropping with real dataset

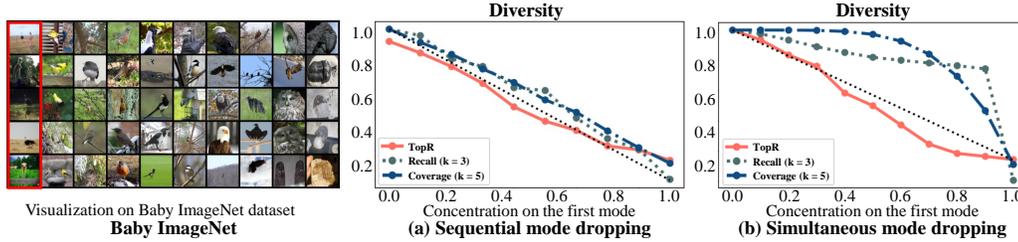


Figure A4: Comparison of evaluation metrics under sequential and simultaneous mode dropping scenario with Baby ImageNet[15].

I.4 Sensitiveness to noise intensity

The purpose of this experiment (Figure A5) is to closely observe the noise sensitivity of metrics in Figure 6. As the noise intensity is incrementally increased until all the metrics converge to 0, the most ideal outcome for the metrics is to exhibit a linear trend to capture these changes most effectively. From the results, it is able to observe that both TopP&R and P&R exhibit the most linear trend in their evaluation outcomes, while D&C demonstrates that least capability to reflect differences based on noise intensity.

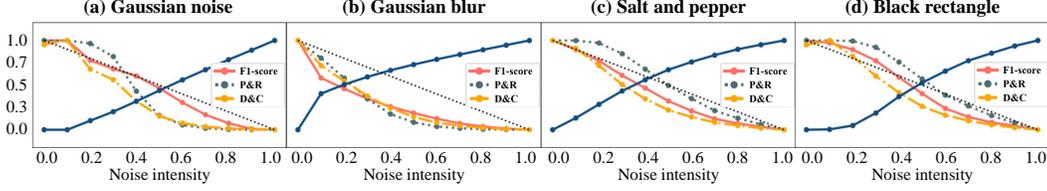


Figure A5: Verification of whether TopP&R can make an accurate quantitative assessment of noise image features. Gaussian noise, gaussian blur, salt and pepper, and black rectangle noise are added to the FFHQ images and emedded with T4096. The noise intensity is linearly increased until all metrics converge to zero.

I.5 Evaluating state-of-the-art generative models on the ImageNet

We conducted our evaluation using the ImageNet dataset, which covers a wide range of classes. We employed commonly accepted metrics in the community, including FID and KID, to rank the considered generative models effectively. The purpose of this experiment is not only to establish the applicability of our ranking metric for generative models, akin to existing single-score metrics, but also to emphasize its interpretability. In Table A13, we calculated the HD between P&R variants and the more reliable KID metric. This comparison aimed to assess the alignment of our metric with established single-score metrics. The results consistently showed that TopP&R’s evaluation results are the most consistent with single-score metrics across various embedding networks. Conversely, P&R exhibited an inconsistent evaluation trend compared to KID, which struggled to effectively differentiate between ReACGAN and BigGAN. Note that, Kang et al. [15] previously demonstrated that BigGAN performs worse than ReACGAN. However, surprisingly, P&R consistently rated BigGAN as better than ReACGAN across all embedding networks. These findings imply that the P&R’s unreliable support estimation under noisy conditions might hinder its ability to accurately distinguish between generative models.

Table A13: Ranking results based on FID, KID, MTD, TopP&R, P&R and D&C for ImageNet (128×128) trained generative models. SwAV, VGG16, and InceptionV3 represent the embedders, and the numbers in parentheses indicate the ranks assigned by each metric to the evaluated models. The HD is the hamming distance between KID and metrics.

	Model	ADM	StyleGAN-XL	ReACGAN	BigGAN	HD
InceptionV3	FID (↓)	9.8715 (1)	10.943 (2)	24.409 (3)	33.935 (4)	0
	KID (↓)	0.0010 (1)	0.0025 (2)	0.0062 (3)	0.0205 (4)	-
	TopP&R (↑)	0.9168 (1)	0.8919 (2)	0.8886 (3)	0.7757 (4)	0
	D&C (↑)	1.0333 (2)	1.0587 (1)	0.8468 (3)	0.5400 (4)	2
	P&R (↑)	0.6988 (1)	0.6713 (2)	0.4462 (4)	0.5079 (3)	2
	MTD (↓)	1.7308 (2)	1.5518 (1)	2.2122 (3)	3.0891 (4)	2
VGG16	TopP&R (↑)	0.9778 (1)	0.8896 (2)	0.8376 (3)	0.5599 (4)	0
	D&C (↑)	1.0022 (3)	0.9159 (4)	1.1322 (1)	1.1020 (2)	4
	P&R (↑)	0.7724 (2)	0.7785 (1)	0.4092 (4)	0.5237 (3)	4
	MTD (↓)	15.748 (3)	22.313 (4)	10.268 (1)	13.992 (2)	4
SwAV	TopP&R (↑)	0.8851 (2)	0.9102 (1)	0.6457 (3)	0.4698 (4)	2
	D&C (↑)	1.0717 (1)	0.8898 (4)	1.0654 (2)	1.0578 (3)	3
	P&R (↑)	0.6096 (1)	0.5596 (2)	0.1335 (4)	0.1544 (3)	2
	MTD (↓)	0.4165 (3)	0.7789 (4)	0.3601 (2)	0.3200 (1)	4

Table A14: Ranking results based on TopP&R, P&R, and D&C for ImageNet (128×128) trained generative models. SwAV, VGG16, and InceptionV3 represent the embedders, and the numbers in parentheses indicate the ranks assigned by each metric to the evaluated models. Note that, the random projection is applied to TopP&R, P&R, and D&C.

	Model	ADM	StyleGAN-XL	ReACGAN	BigGAN
InceptionV3	TopP&R (\uparrow)	0.9168 (1)	0.8919 (2)	0.8886 (3)	0.7757 (4)
	D&C w/ rand proj (\uparrow)	1.0221 (2)	1.0249 (1)	0.9218 (3)	0.7099 (4)
	P&R w/ rand proj (\uparrow)	0.7371 (1)	0.7346 (2)	0.6117 (4)	0.6305 (3)
VGG16	TopP&R (\uparrow)	0.9778 (1)	0.8896 (2)	0.8376 (3)	0.5599 (4)
	D&C w/ rand proj (\uparrow)	1.0086 (3)	0.9187 (4)	1.0895 (1)	1.0803 (2)
	P&R w/ rand proj (\uparrow)	0.7995 (2)	0.8056 (1)	0.5917 (4)	0.6304 (3)
SwAV	TopP&R (\uparrow)	0.8851 (2)	0.9102 (1)	0.6457 (3)	0.4698 (4)
	D&C w/ rand proj (\uparrow)	1.0526 (3)	0.9479 (4)	1.1037 (2)	1.1127 (1)
	P&R w/ rand proj (\uparrow)	0.6873 (1)	0.6818 (2)	0.4179 (3)	0.4104 (4)

I.6 Verification of random projection effect in generative model ranking

We also applied the same random projection used by TopP&R to P&R and D&C in Section I.6. From the results of the Table A14, it is evident that the application of random projection to P&R leads to different rankings than its original evaluation tendencies. Similarly, the rankings of D&C with random projection still exhibit significant variations based on the embedding. Thus, we have demonstrated that random projection does not provide consistent evaluation tendencies across different embeddings, while also highlighting that this characteristic is unique to TopP&R. To quantify these differences among metrics, we computed the mean Hamming Distance (MHD) between the results in the embedding spaces of each metric. The calculated MHD scores were 1.33 for TopP&R, 2.67 for P&R, and 3.33 for D&C, respectively.

I.7 Verifying the effect of random projection to the noisy data

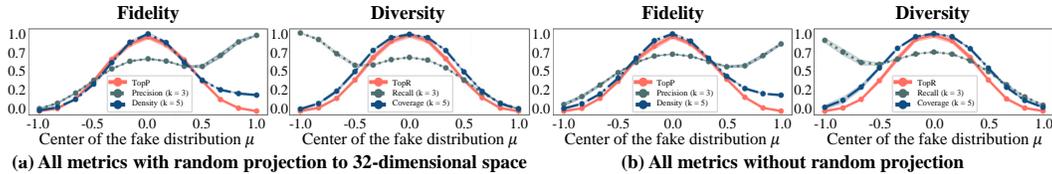


Figure A6: Behaviors of evaluation metrics for outliers on real and fake distribution. For both real and fake data, the outliers are fixed at $3 \in \mathbb{R}^d$ where d is a dimension, and the parameter μ is shifted from -1 to 1. (a) depicts the comparison of all metrics using a random projection from 64 dimensions to 32 dimensions. (b) illustrates the difference between metrics without using random projection, where real and fake data are considered in 32 dimensions.

We experiment to ascertain whether utilizing a random projection address the drawbacks of existing metrics that are susceptible to noise, and we also conduct tests to validate whether TopP&R possesses noise robustness without using random projection. Through Figure A6 (a), it is evident that P&R and D&C still retain vulnerability to noisy features, and it is apparent that random projection itself does not diminish the impact of noise. Furthermore, Figure A6 (b) reveals that TopP&R, even without utilizing a random projection, exhibits robustness to noise through approximating the data support based on statistically and topologically significant features.

Building upon the insight from the toy data experiment in Figure A6 that random projection does not confer noise robustness to P&R and D&C, we aim to further demonstrate this fact through a real data experiment. In the experiment depicted in Figure A7, we follow the same setup as the previous experiment in Figure 5, applying the same random projection to P&R and D&C only. The results of

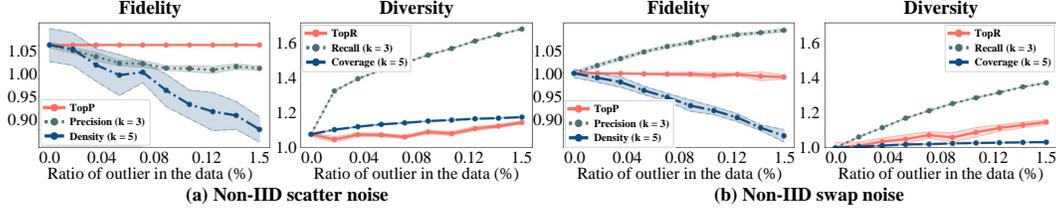


Figure A7: Comparison of evaluation metrics on Non-IID perturbations using FFHQ dataset. We replaced certain ratio of \mathcal{X} and \mathcal{Y} (a) with outliers and (b) by switching some of real and fake features. P&R and D&C were computed using the same 32-dimensional random projection.

this experiment affirm that random projection does not provide additional robustness against noise for P&R and D&C, while TopP&R continues to exhibit the most robust performance.

I.8 Robustness of TopP&R with respect to random projection dimension

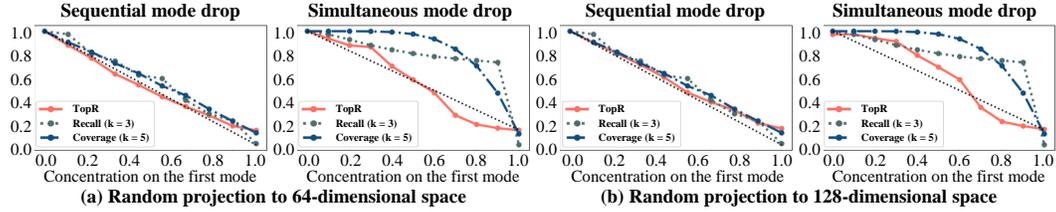


Figure A8: Comparison of evaluation trends of TopP&R with respect to random projection dimension in both sequential and simultaneous mode drop scenarios. Baby ImageNet dataset is used for the comparison, and TopP&R is computed using random projections to (a) 64 dimensions and (b) 128 dimensions.

To investigate whether the dimension of random projection influences the evaluation trend of TopP&R, we conduct mode dropping experiment using the Baby ImageNet dataset, similar to the experiment in Section I.3. Results in (a) and (b) of Figure A8 respectively compare the performance of TopP&R using random projection in dimensions of 64 and 128 with other metrics. From the experimental results, we observe that TopP&R is the most sensitive to the distributional changes, akin to the tendencies observed in the previous toy mode dropping experiment. Furthermore, upon examining the differences across random projection dimensions, TopP&R shows consistent evaluation trends regardless of dimensions.

I.9 Truncation trick

ψ is a parameter for the truncation trick and is first introduced in [4] and [1]. We followed the approach in [4] and [1]. GANs generate images using the noise input z , which follows the standard normal distribution $\mathcal{N}(0, I)$ or uniform distribution $\mathcal{U}(-1, 1)$. Suppose GAN inadvertently samples noise outside of distribution, then it is less likely to sample the image from the high density area of the image distribution $p(z)$ defined in the latent space of GAN, which leads to generate an image with artifacts. The truncation trick takes this into account and uses the following truncated distribution. Let f be the mapping from the input to the latent space. Let $w = f(z)$, and $\bar{w} = \mathbb{E}[f(z)]$, where z is either from $\mathcal{N}(0, I)$ or $\mathcal{U}(-1, 1)$. Then we use $w' = \bar{w} + \psi(w - \bar{w})$ as a truncated latent vector. If the value of ψ increases, then the degree of truncation decreases which makes images have greater diversity but possibly lower fidelity.

I.10 Toy experiment of trade-off between fidelity and diversity

We have designed a new toy experiment to mimic the truncation trick. With data distributions of 10k samples, $\mathcal{X} \sim \mathcal{N}(\mu = 0, I) \in \mathbb{R}^{32}$ and $\mathcal{Y} \sim \mathcal{N}(\mu = 0.6, \sigma^2) \in \mathbb{R}^{32}$, we measure the fidelity and diversity while incrementally increasing σ from 0.7 to 1.3. For smaller σ values, the evaluation metric should exhibit higher fidelity and lower diversity, while increasing σ should demonstrate decreasing

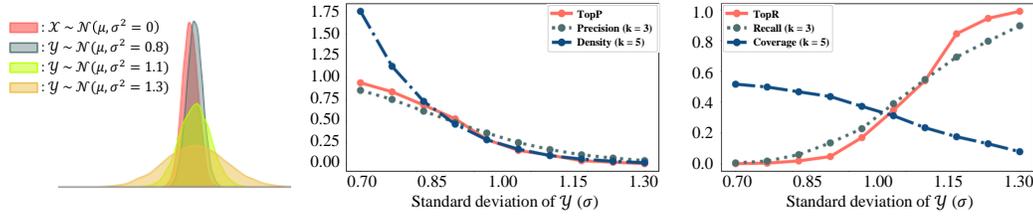


Figure A9: Trade-off between fidelity and diversity in a toy experiment. In this experiment, the μ of \mathcal{X} was set to 0, and the μ of \mathcal{Y} was set to 0.6. By incrementally increasing σ from 0.7 to 1.3, the evaluation trend based on changes in the \mathcal{Y} distribution is measured.

fidelity and higher diversity. From the results in Figure A9, both TopP&R and P&R exhibit a trade-off between fidelity and diversity.

I.11 Resolving fidelity and diversity

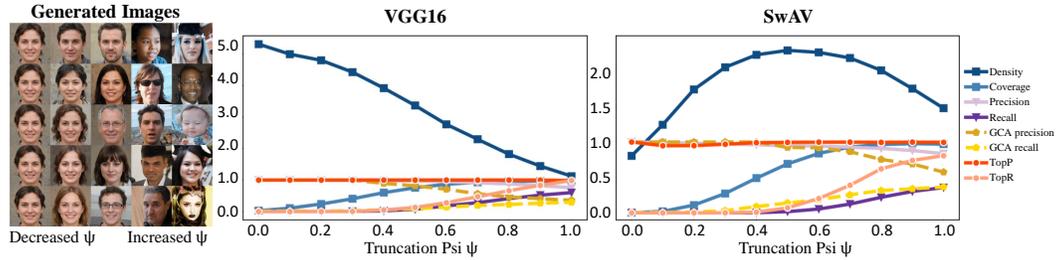


Figure A10: Behaviors of metrics with truncation trick. The horizontal axis corresponds to the value of ψ denoting the increased diversity. The images are generated via StyleGAN2 with FFHQ dataset [1].

To test whether TopP&R responds appropriately to the change in the underlying distributions in real scenarios, we test the metric on the generated images of StyleGAN2 [2] using the truncation trick [1]. StyleGAN2, trained on FFHQ as shown by Han et al. [49], tends to generate samples mainly belonging to the majority of the real data distribution, struggling to produce rare samples effectively. In other words, while StyleGAN2 excels in generating high-fidelity images, it lacks diversity. Therefore, the evaluation trend in this experiment should reflect a steady high fidelity score while gradually increasing diversity. Han et al. [49] demonstrates that the generated image quality of StyleGAN2 that lies within the support region exhibits sufficiently realistic visual quality. Therefore, particularly in the case of TopP&R, which evaluates fidelity excluding noise, the fidelity value should remain close to 1. As shown in Figure A10, every time the distribution is transformed by ψ , TopP&R responds well and shows consistent behavior across different embedders with bounded scores in $[0, 1]$, which are important virtues as an evaluation metric. On the other hand, *Density* gives unbounded scores (fidelity > 1) and shows inconsistent trend depending on the embedder. Because *Density* is not capped in value, it is difficult to interpret the score and know exactly which value denotes the best performance (*e.g.*, in our case, the best performance is when TopP&R = 1).