



Robust Kernel Density Estimation with Median-of-Means principle

Pierre Humbert, Batiste Le Bars, Ludovic Minvielle

► To cite this version:

Pierre Humbert, Batiste Le Bars, Ludovic Minvielle. Robust Kernel Density Estimation with Median-of-Means principle. ICML 2022 - The 39th International Conference on Machine Learning (ICML), Jul 2022, Baltimore, United States. hal-02882092

HAL Id: hal-02882092

<https://hal.science/hal-02882092>

Submitted on 29 Jun 2020

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.

Robust Kernel Density Estimation with Median-of-Means principle

Pierre Humbert* Batiste Le Bars* Ludovic Minvielle*

Nicolas Vayatis

Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli
F-91190 Gif-sur-Yvette, France

June 5, 2020

Abstract

In this paper, we introduce a robust nonparametric density estimator combining the popular Kernel Density Estimation method and the Median-of-Means principle (MoM-KDE). This estimator is shown to achieve robustness to any kind of anomalous data, even in the case of adversarial contamination. In particular, while previous works only prove consistency results under known contamination model, this work provides finite-sample high-probability error-bounds without *a priori* knowledge on the outliers. Finally, when compared with other robust kernel estimators, we show that MoM-KDE achieves competitive results while having significant lower computational complexity.

1 Introduction

Over the past years, the task of learning in the presence of outliers has become an increasingly important objective in both statistics and machine learning. Indeed, in many situations, training data can be contaminated by undesired samples, which may badly affect the resulting learning task, especially in adversarial settings. Building robust estimators and algorithms that are resilient to outliers is therefore becoming crucial in many learning procedures. In particular, the inference of a probability density function from a contaminated random sample is of major concerns.

Density estimation methods are mostly divided into parametric and nonparametric techniques. Among the nonparametric family, the Kernel Density Estimator (KDE) is probably the most known and used for both univariate and multivariate densities [Parzen, 1962; Silverman, 1986; Scott, 2015], but it also known to be sensitive to dataset contaminated by outliers [Kim and Scott, 2011, 2012; Vandermeulen and Scott, 2014]. The construction of robust KDE is therefore an important area of research, that can have useful applications such as anomaly detection and resilience to adversarial data corruption. Yet, only few works have proposed such robust estimators.

Kim and Scott [2012] proposed to combine KDE with ideas from M-estimation to construct the so-called Robust Kernel Density Estimator (RKDE). However, no consistency results were provided and robustness was rather shown experimentally. Later, RKDE was proven to converge to the true density, however at the condition that the dataset remains uncorrupted [Vandermeulen and Scott, 2013]. More recently, Vandermeulen and Scott [2014] proposed another robust estimator, called Scaled and Projected KDE (SPKDE). Authors proved the L_1 -consistency of SPKDE under a variant of the Huber's ε -contamination model where two strong assumptions are made [Huber, 1992]. First, the contamination parameter ε is known, and second, the outliers are drawn from an uniform distribution when outside the support of the true density. Unfortunately, as they did not provided rates of convergence, it still remains unclear at which speed SPKDE converges to the true density. Finally, both RKDE and SPKDE require iterative algorithms to compute their estimators, increasing the overall complexity of their construction.

In statistical analysis, another idea to construct robust estimators is to use the Median-of-Means principle (MoM). Introduced by Nemirovsky and Yudin [1983], Jerrum et al. [1986], and Alon et al. [1999], the MoM was first designed to estimate the mean of a real random variable. It relies on the simple idea

*These authors contributed equally to this work.

that rather than taking the average of all the observations, the sample is split in several non-overlapping blocks over which the mean is computed. The MoM estimator is then defined as the median of these means. Easy to compute, the MoM properties have been studied by Minsker et al. [2015] and Devroye et al. [2016] to estimate the means of heavy-tailed distributions. Furthermore, due to its robustness to outliers, MoM-based estimators have recently gained a renewed of interest in the machine learning community [Lecué et al., 2020; Lecué and Lerasle, 2019].

Contributions. In this paper, we propose a new robust nonparametric density estimator based on the combination of the Kernel Density Estimation method and the Median-of-Means principle (MoM-KDE). We place ourselves in a more general framework than the classical Huber contamination model, called $\mathcal{O} \cup \mathcal{I}$, which gets rid of any assumption on the outliers. We demonstrate the statistical performance of the estimator through finite-sample high-confidence error bounds in the L_∞ -norm and show that MoM-KDE’s convergence rate is the same as KDE without outliers. Additionally, we prove the consistency in the L_1 -norm, which is known to reflect the global performance of the estimate. To the best of our knowledge, this is the first work that presents such results in the context of robust kernel density estimation, especially under the $\mathcal{O} \cup \mathcal{I}$ framework. Finally, we demonstrate the empirical performance of MoM-KDE on both synthetic and real data and show the practical interest of such estimator as it has a lower complexity than the baseline RKDE and SPKDE.

2 Median-of-Means Kernel Density Estimation

We first recall the classical kernel density estimator. Let X_1, \dots, X_n be independent and identically distributed (i.i.d.) random variables that have a probability density function (pdf) $f(\cdot)$ with respect to the Lebesgue measure on \mathbb{R}^d . The Kernel Density Estimate of f (KDE), also called the *Parzen–Rosenblatt estimator*, is a nonparametric estimator given by

$$\hat{f}_n(x) = \frac{1}{nh^d} \sum_{i=1}^n K\left(\frac{X_i - x}{h}\right), \quad (1)$$

where $h > 0$ and $K : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is an integrable function satisfying $\int K(u)du = 1$ [Tsybakov, 2008]. Such a function $K(\cdot)$ is called a *kernel* and the parameter h is called the *bandwidth* of the estimator. The bandwidth is a smoothing parameter that controls the bias-variance tradeoff of $\hat{f}_n(\cdot)$ with respect to the input data.

While this estimator is central in statistic, a major drawback is its weakness against outliers [Kim and Scott, 2008, 2011, 2012; Vandermeulen and Scott, 2014]. Indeed, as it assigns uniform weights $1/n$ to every points regardless of whether X_i is an outlier or not, inliers and outliers contribute equally in the construction of the KDE, which results in undesired “bumps” over outlier locations in the final estimated density (see Figure 1). In the following, we propose a KDE-based density estimator robust to the presence of outliers in the sample set. These outliers are considered in a general framework described in the next section.

2.1 Outlier setup

Throughout the paper, we consider the $\mathcal{O} \cup \mathcal{I}$ framework introduced by Lecué and Lerasle [2019]. This very general framework allows the presence of outliers in the dataset and relax the standard i.i.d. assumption on each observation. We therefore assume that the n random variables are partitioned into two (unknown) groups: a subset $\{X_i \mid i \in \mathcal{I}\}$ made of inliers, and another subset $\{X_i \mid i \in \mathcal{O}\}$ made of outliers such that $\mathcal{O} \cap \mathcal{I} = \emptyset$ and $\mathcal{O} \cup \mathcal{I} = \{1, \dots, n\}$. While we suppose the $X_{i \in \mathcal{I}}$ are i.i.d. from a distribution that admits a density f with respect to the Lebesgue measure, no assumption is made on the outliers $X_{i \in \mathcal{O}}$. Hence, these outlying points can be dependent, adversarial, or not even drawn from a proper probability distribution.

The $\mathcal{O} \cup \mathcal{I}$ framework is related to the well-known Huber’s ε -contamination model [Huber, 1992] where it is assumed that data are i.i.d. with distribution $g = \varepsilon f_{\mathcal{I}} + (1 - \varepsilon)f_{\mathcal{O}}$, and $\varepsilon \in [0, 1]$; the distribution $f_{\mathcal{I}}$ being related to the inliers and $f_{\mathcal{O}}$ to the outliers. However, there are several important differences. First, in the $\mathcal{O} \cup \mathcal{I}$ the proportion of outliers is fixed and equals to $|\mathcal{O}|/n$, whereas it is random in the Huber’s

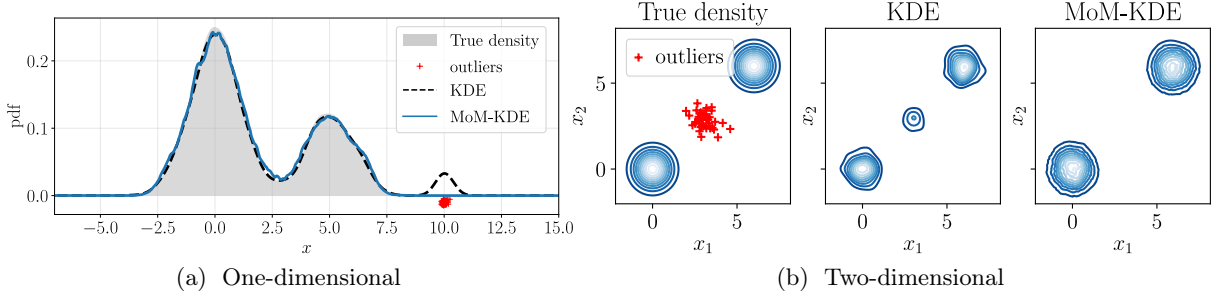


Figure 1: True density, outliers, KDE, and MoM-KDE. (a) Estimates from a 1-D true density and outliers from a normal density centered in $\mu_O = 10$ with variance $\sigma_O^2 = 0.1$. (b) Estimates from a 2-D true density and outliers from a normal density centered in $\mu_O = (3, 3)$ with variance $\sigma_O^2 = 0.5I_2$.

ε -contamination model [Lerasle, 2019]. Second, the $\mathcal{O} \cup \mathcal{I}$ is less restrictive. Indeed, contrary to Huber’s model which considers that inliers and outliers are respectively i.i.d from the same distributions, $\mathcal{O} \cup \mathcal{I}$ does not make a single assumption on the outliers.

2.2 MoM-KDE

We now present our main contribution, a robust kernel density estimator based on the MoM. This estimator is essentially motivated by the fact that the classical kernel density estimation at one point corresponds to an empirical average (see Equation (1)). Therefore, the MoM principle appears to be an intuitive solution to build a robust version of the KDE. A formal definition of MoM-KDE is given below.

Definition 1. (*MoM Kernel Density Estimator*) Let $1 \leq S \leq n$, and let B_1, \dots, B_S be a random partition of $\{1, \dots, n\}$ into S non-overlapping blocks B_s of equal size $n_s \triangleq n/S$. The MoM Kernel Density Estimator (MoM-KDE) of f at x_0 is given by

$$\hat{f}_{MoM}(x_0) \propto \text{Median} \left(\hat{f}_{n_1}(x_0), \dots, \hat{f}_{n_S}(x_0) \right), \quad (2)$$

where $\hat{f}_{n_s}(x_0)$ is the value of the standard kernel density estimator at x_0 obtained via the samples of the s -th block B_s . Note that $\hat{f}_{MoM}(\cdot)$ is not necessarily a density as its integral may not be equal to 1. When needed, we thus normalize it by its integral the same way it is proposed by Devroye and Lugosi [2012].

Broadly speaking, MoM estimators appear to be a good tradeoff between the unbiased but non robust empirical mean and the robust but biased median [Lecué et al., 2020]. A visual example of the robustness of MoM-KDE is displayed in Figure 1. We now give a simple example highlighting the robustness of MoM-KDE.

Example 1. (*MoM-KDE v.s. Uniform KDE*) Let the inliers be i.i.d. samples from a uniform distribution on the interval $[-1, 1]$ and the outliers be i.i.d. samples from another uniform distribution on $[-3, 3]$. Let the kernel function be the uniform kernel, $x_0 = 2$ and $h \in (0, 1)$. Then if $S > 2|\mathcal{O}|$, we obtain

$$|\hat{f}_{MoM}(x_0) - f(x_0)| = 0 \quad \text{a.s.} \quad \text{and} \quad \mathbb{P} \left(|\hat{f}_n(x_0) - f(x_0)| = 0 \right) = (1 - h/3)^{|\mathcal{O}|} \neq 1.$$

This result shows that the MoM-KDE makes (almost surely) no error at the point x_0 . On the contrary, the KDE here has a non-negligible probability to make an error.

2.3 Time complexity

The complexity of MoM-KDE to evaluate one point is the same as the standard KDE, $\mathcal{O}(n)$; $\mathcal{O}(S \cdot \frac{n}{S})$ for the block-wise evaluation and $\mathcal{O}(n)$ to compute the median with the *median-of-medians algorithm* [Blum et al., 1973]. Since RKDE and SPKDE are KDEs with modified weights, they also perform the evaluation step in $\mathcal{O}(n)$ time. However, these weights need to be learnt, thus requiring an additional non-negligible computing capacity. Indeed, each one of them rely on an iterative method – the iteratively reweighted least squares algorithm and the projected gradient descent algorithm, that both have a complexity of $\mathcal{O}(n_{iter} \cdot n^2)$, where n_{iter} is the number of needed iterations to reach a reasonable accuracy. MoM-KDE on the other hand does not require any learning procedure. Note that the evaluation step can be accelerated

Table 1: Computational complexity

Method	Learning	Evaluation	Iterative method
KDE [Parzen, 1962]	–	$\mathcal{O}(n)$	no
RKDE [Kim and Scott, 2012]	$\mathcal{O}(n_{iter} \cdot n^2)$	$\mathcal{O}(n)$	yes
SPKDE [Vandermeulen and Scott, 2014]	$\mathcal{O}(n_{iter} \cdot n^2)$	$\mathcal{O}(n)$	yes
MoM-KDE	–	$\mathcal{O}(n)$	no

through several ways, hence potentially reducing computational time of all these competing methods [Gray and Moore, 2003a,b; Wang and Scott, 2019; Backurs et al., 2019]. Theoretical time complexities are gathered in Table 1.

3 Theoretical analysis

In this section, we give a finite-sample high-probability error bound in the L_∞ -norm for MoM-KDE under the $\mathcal{O} \cup \mathcal{I}$ framework. To our knowledge, we are the first to provide such error bounds in robust kernel density estimation under this framework. In particular, our objective is to prove that even with a contaminated dataset, MoM-KDE achieves a similar convergence rate than KDE without outliers [Sriperumbudur and Steinwart, 2012; Jiang, 2017; Wang et al., 2019]. In order to build this high-probability error bound, it is assumed, among other standard hypotheses, that the true density is Hölder-continuous, a smoothness property usually considered in KDE analysis [Tsybakov, 2008; Jiang, 2017; Wang et al., 2019]. In addition, we show the consistency in the L_1 -norm. In this last result, we will see that the aforementioned assumptions are not necessary to obtain the consistency. In the following, we give the necessary definitions and assumptions to perform our non-asymptotic analysis.

3.1 Setup and assumptions

Let us first list the usual assumptions, notably on the considered kernel function, that will allow us to derive our results. They are standard in KDE analysis, and are chosen for their simplicity of comprehension [Tsybakov, 2008; Jiang, 2017]. More general hypotheses could be made in order to obtain the same results, notably assuming kernel of order ℓ (see for example the works of Tsybakov [2008] and Wang et al. [2019]).

Assumption 1. (Bounded density) $\|f\|_\infty < \infty$.

We make the following assumptions on the kernel K .

Assumption 2. (Density kernel) $\forall u \in \mathbb{R}^d, K(u) \geq 0$, and $\int K(u)du = 1$.

Assumption 3. (Spherically symmetric and non-increasing) There exists a non-increasing function $k : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that $K(u) = k(\|u\|)$ for all $u \in \mathbb{R}^d$, where $\|\cdot\|$ is any norm of \mathbb{R}^d .

Assumption 4. (Exponentially decaying tail) There exists positive constants $\rho, C_\rho, t_0 > 0$ such that for all $t > t_0$

$$k(t) \leq C_\rho \cdot \exp(-t^\rho) .$$

All the above assumptions are respected by most of the popular kernels, in particular the Gaussian, Exponential, Uniform, Triangular, Cosine kernel, etc. Furthermore, the last assumption implies that for any $m > 0$, we have $\int \|u\|^m K(u)du < \infty$ (finite norm moment) [Jiang, 2017]. Finally, when taken together, these assumptions imply that the kernel satisfies the VC property [Wang et al., 2019]. These are key properties to provide the bounds presented in the next section.

Before stating our main results, we recall the definition of the Hölder class of functions.

Definition 2. (Hölder class) Let T be an interval of \mathbb{R}^d , and $0 < \alpha \leq 1$ and $L > 0$ be two constants. We say that a function $f : T \rightarrow \mathbb{R}$ belongs to the Hölder class $\Sigma(L, \alpha)$ if it satisfies

$$\forall x, x' \in T, \quad |f(x) - f(x')| \leq L\|x - x'\|^\alpha . \quad (3)$$

This definition implies a smoothness regularization on the function f , and is a convenient property to bound the bias of KDE-based estimators.

3.2 L_∞ and L_1 consistencies of MoM-KDE

This section states our central finding, a L_∞ finite-sample error bound for MoM-KDE that proves its consistency and yields the same convergence rate as KDE with uncontaminated data. The latter is given by the following Lemma partly proven by Sriperumbudur and Steinwart [2012] and verified several times in the literature [Giné and Guillou, 2002; Jiang, 2017; Wang et al., 2019].

Lemma 1. (*L_∞ error-bound of the KDE without anomalies*) Suppose that f belongs to the class of densities $\mathcal{P}(\alpha, L)$ defined as

$$\mathcal{P}(\alpha, L) \triangleq \left\{ f \mid f \geq 0, \int f(x)dx = 1, \text{ and } f \in \Sigma(\alpha, L) \right\}, \quad (4)$$

where $\Sigma(\alpha, L)$ is the Hölder class of function on \mathbb{R}^d (Definition 2). Grant assumptions 1 to 4 and let $n > 1$, $h \in (0, 1)$ and $S \geq 1$ such that $nh^d \geq S$ and $nh^d \geq |\log(h)|$. Then with probability at least $1 - \exp(-S)$, we have

$$\|\hat{f}_n - f\|_\infty \leq C_1 \sqrt{\frac{S|\log(h)|}{nh^d}} + C_2 h^\alpha, \quad (5)$$

where $C_2 = L \int \|u\|^\alpha K(u)du < \infty$ and C_1 is a constant that only depends on $\|f\|_\infty$, the dimension d , and the kernel properties.

This Lemma comes from the well-known bias-variance decomposition, where we separately bound the variance (see Theorem 3.1 of Sriperumbudur and Steinwart [2012]) and the bias (see e.g. [Tsybakov, 2008] or [Rigollet et al., 2009]). It shows the consistency of KDE without anomalies, as soon as $h \rightarrow 0$ and $nh^d \rightarrow \infty$.

We now present our main result. Its objective is to show that even under the $\mathcal{O} \cup \mathcal{I}$ framework, we do not need any additional hypothesis – besides the ones of the previous lemma – to show that MoM-KDE achieves the same convergence rate as KDE when used with uncontaminated data.

Proposition 1. (*L_∞ error-bound of the MoM-KDE under the $\mathcal{O} \cup \mathcal{I}$*) Suppose that f belongs to the class of densities $\mathcal{P}(\alpha, L)$ and grant assumptions 1 to 4. Let S be the number of blocks, $\delta > 0$ such that $S > (2 + \delta)|\mathcal{O}|$, and $\Delta = (1/(2 + \delta) - |\mathcal{O}|/S)$. Then, for any $h \in (0, 1)$, δ sufficiently small, and $n \geq 1$ such that $nh^d \geq S \log(2(2 + \delta)/\delta)$, and $nh^d \geq S|\log(h)|$, we have with probability at least $1 - \exp(-2\Delta^2 S)$,

$$\|\hat{f}_{MoM} - f\|_\infty \leq C_1 \sqrt{\frac{S \log\left(\frac{2(2+\delta)}{\delta}\right) |\log(h)|}{nh^d}} + C_2 h^\alpha, \quad (6)$$

where $C_2 = L \int \|u\|^\alpha K(u)du < \infty$ and C_1 is a constant that only depends on $\|f\|_\infty$, the dimension d , and the kernel properties.

The proof is given in the supplementary material. From equation (6), the optimal choice of the bandwidth is $h \asymp \left(\frac{S \log(n)}{n}\right)^{1/(2\alpha+d)}$ leading to the final rate of $\left(\frac{S \log(n)}{n}\right)^{\alpha/(2\alpha+d)}$. This convergence rate is the same (up to a constant) to the one of KDE without anomalies, with the same exponential control (Lemma 1). Note that when there is no outlier, i.e. $|\mathcal{O}| = 0$, the bound holds for $S = 1$, and we recover the classical KDE minimax optimal rate [Wang et al., 2019]. In addition, the previous proposition states that the convergence of the MoM-KDE only depends on the number of outliers in the dataset, and not on their “type”. This estimator is therefore robust in a wide range of scenarios, including the adversarial one.

We now give a L_1 -consistency result under mild hypotheses, which is known to reflect the global performance of the estimate. Indeed, small L_1 error leads to accurate probability estimation [Devroye and Györfi, 1985].

Proposition 2. (*L_1 -consistency in probability*) If $n/S \rightarrow \infty$, $h \rightarrow 0$, $nh^d \rightarrow \infty$, and $S > 2|\mathcal{O}|$, then

$$\|\hat{f}_{MoM} - f\|_1 \xrightarrow[n \rightarrow \infty]{\mathcal{P}} 0. \quad (7)$$

This result is obtained by bounding the left-hand part by the errors in the healthy blocks only, i.e. those without anomalies. Under the hypothesis of the proposition, these errors are known to converge to 0 in probability [Wang et al., 2019]. The complete proof is given in supplementary material. Contrary to SPKDE [Vandermeulen and Scott, 2014], no assumption on the outliers generation process is necessary to obtain this consistency result. Moreover, while we need to assume that the proportion of outliers is perfectly known to prove the convergence of SPKDE, the MoM-KDE converges whenever the number of outliers is overestimated.

4 Numerical experiments

In this section, we display numerical results supporting the relevance of MoM-KDE. All experiences were run over a personal laptop computer using `Python`. The code of MoM-KDE is made available online¹.

Comparative methods. In the following experiments, we propose to compare MoM-KDE to the classical KDE and two robust versions of KDE, called RKDE [Kim and Scott, 2012] and SPKDE [Vandermeulen and Scott, 2014].

As previously explained, RKDE takes the ideas of robust M-estimation and translate it to kernel density estimation. Authors point out that classical KDE estimator can be seen as the minimizer of a squared error loss in the Reproducing Kernel Hilbert Space \mathcal{H} corresponding to the chosen kernel. Instead of minimizing this loss, they propose to minimize a robust version of it, $\sum_i \rho(\|\phi(X_i) - g\|_{\mathcal{H}})$, with respect to $g \in \mathcal{H}$. Here ϕ is the canonical feature map and $\rho(\cdot)$ is either the robust Huber or Hampel function. The solution of the newly expressed problem is then found using the iteratively reweighted least squares algorithm.

SPKDE proposes to scale the standard KDE in a way that it decontaminates the dataset. This is done by minimizing the function $\|\beta \hat{f}_n - g\|_2$ with respect to g , belonging to the convex hull of $\{k_h(\cdot, X_i)\}_{i=1}^n$. Here, β is an hyperparameter that controls the robustness and \hat{f}_n is the KDE estimator. The minimization is shown to be equivalent to a quadratic program over the simplex, solved via projected gradient descent.

Metrics. The performance of the MoM-KDE is measured through three metrics, two are used to measure the similarity between the estimated and the true density, and one describes performances of an anomaly detector based on the estimated density. The first one is the Kullback-Leibler divergence [Kullback and Leibler, 1951] which is the most used in robust KDE [Kim and Scott, 2008, 2011, 2012; Vandermeulen and Scott, 2014]. Used to measure the similarity between distributions, it is defined as

$$D_{\text{KL}}(\hat{f} \| f) = \int \hat{f}(x) \log \left(\frac{\hat{f}(x)}{f(x)} \right) dx .$$

As the Kullback-Leibler divergence is non-symmetric and may have infinite values when distributions do not share the same support, we also consider the Jensen-Shannon divergence [Endres and Schindelin, 2003; Liese and Vajda, 2006]. It is a symmetrized version of D_{KL} , with positive values, bounded by 1 (when the logarithm is used in base 2), and has found applications in many fields, such as deep learning [Goodfellow et al., 2014] or transfer learning [Segev et al., 2017]. It is defined as

$$D_{\text{JS}}(\hat{f} \| f) = \frac{1}{2} \left(D_{\text{KL}}(\hat{f} \| g) + D_{\text{KL}}(f \| g) \right), \quad \text{with} \quad g = \frac{1}{2}(\hat{f} + f) .$$

Motivated by real-world application, the third metric is not related to the true density, which is usually not available in practical cases. Instead, we quantify the capacity of the learnt density to detect anomalies using the well-known Area Under the ROC Curve criterion (AUC). An input point x_0 is considered abnormal whenever $\hat{f}(x_0)$ is below a given threshold.

Hyperparameters. All estimators are built using the Gaussian kernel. The number of blocks in MoM-KDE is selected on a regular grid of 20 values between 1 and $2|\mathcal{O}| + 1$ in order to obtain the lowest D_{JS} . The bandwidth h is chosen for KDE via the pseudo-likelihood k -cross-validation method [Turlach, 1993], and used for all estimators. The construction of RKDE follows exactly the indications of its authors [Kim and Scott, 2012] and $\rho(\cdot)$ is taken to be the Hampel function as they empirically showed that it is the most robust. For SPKDE, the true ratio of anomalies is given as an input parameter.

¹<https://github.com/lminvielle/mom-kde> For the sake of comparison, we also implemented RKDE and SPKDE.

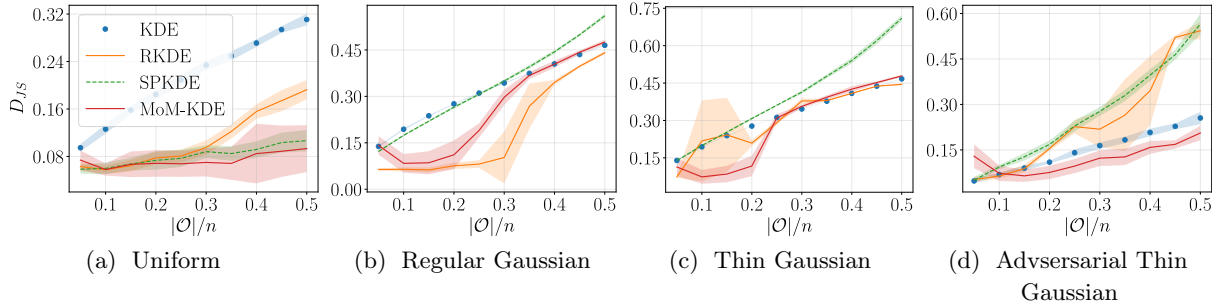


Figure 2: Density estimation with synthetic data. The displayed metric is the Jensen-Shannon divergence. A lower score means a better estimation of the true density.

4.1 Results on synthetic data.

To evaluate the efficiency of the MoM-KDE against KDE and its robust competitors, we set up several outlier situations. In all these situations, we draw $N = 1000$ inliers from an equally distributed mixture of two normal distribution $\mathcal{N}(\mu_1, \sigma_1)$ and $\mathcal{N}(\mu_2, \sigma_2)$ with $\mu_1 = 0$, $\mu_2 = 6$, and $\sigma_1 = \sigma_2 = 0.5$. The outliers however are sampled through various schemes:

- (a) **Uniform.** A uniform distribution $U([\mu_1 - 3, \mu_2 + 3])$ which is the classical setting used for outlier simulation.
- (b) **Regular Gaussian.** A *similar*-variance normal distribution $\mathcal{N}(3, 0.5)$ located between the two inlier clusters.
- (c) **Thin Gaussian.** A *low*-variance normal distribution $\mathcal{N}(3, 0.01)$ located between the two inliers clusters.
- (d) **Adversarial Thin Gaussian.** A low variance normal distribution $\mathcal{N}(0, 0.01)$ located on one of the inliers' Gaussian mode. This scenario can be seen as adversarial as an ill-intentioned agent may hide wrong points in region of high density. It is the most challenging setting for standard robust estimators as they are in general robust to outliers located outside the support of the density we wish to estimate.

For all situations, we consider several ratios of contamination and set the number of outliers $|\mathcal{O}|$ in order to obtain a ratio $|\mathcal{O}|/n$ ranging from 0.05 to 0.5 with 0.05-wide steps. Finally, to evaluate the pertinence of our results, for each set of parameters, data are generated 10 times.

We display in Figure 2 the results over synthetic data using the D_{JS} score. The average scores and standard deviations over the 10 experiments are represented for each outlier scheme and ratio. Overall, the results show the good performance of MoM-KDE in all the considered situations. Furthermore, they highlight the dependency of the two competitors to the type of outliers. Indeed, as SPKDE is designed to handle uniformly distributed outliers, the algorithm struggles when confronted with differently distributed outliers (see Figure 2 (b, c, d)). RKDE performs generally better, but fails against adversarial contamination, which may be explained by its tendency to down-weight points located in low-density regions, which for this particular case correspond to the inliers. Results over D_{KL} and AUC are reported in the supplementary materials. Generally, they show similar results and the same conclusions on the good performance of MoM-KDE can be made.

4.2 Results on real data.

Experiments are also conducted over six classification datasets: Banana, German, Titanic, Breast-cancer, Iris, and Digits. They contain respectively $n = 5300, 1000, 2201, 569, 150$ and 1797 data points having $d = 2, 20, 3, 30, 4$ and 64 input dimensions. They are all publicly available either from open repositories ² (for the first three) or directly from Scikit-learn package (for the last three) [Pedregosa et al., 2011]. We follow the approach of Kim and Scott [2012] that consists in setting the class labeled 0 as outliers and the rest as inliers. To artificially control the outlier proportion, we randomly downsample the abnormal

²<http://www.raetschlab.org/Members/raetsch/benchmark/>

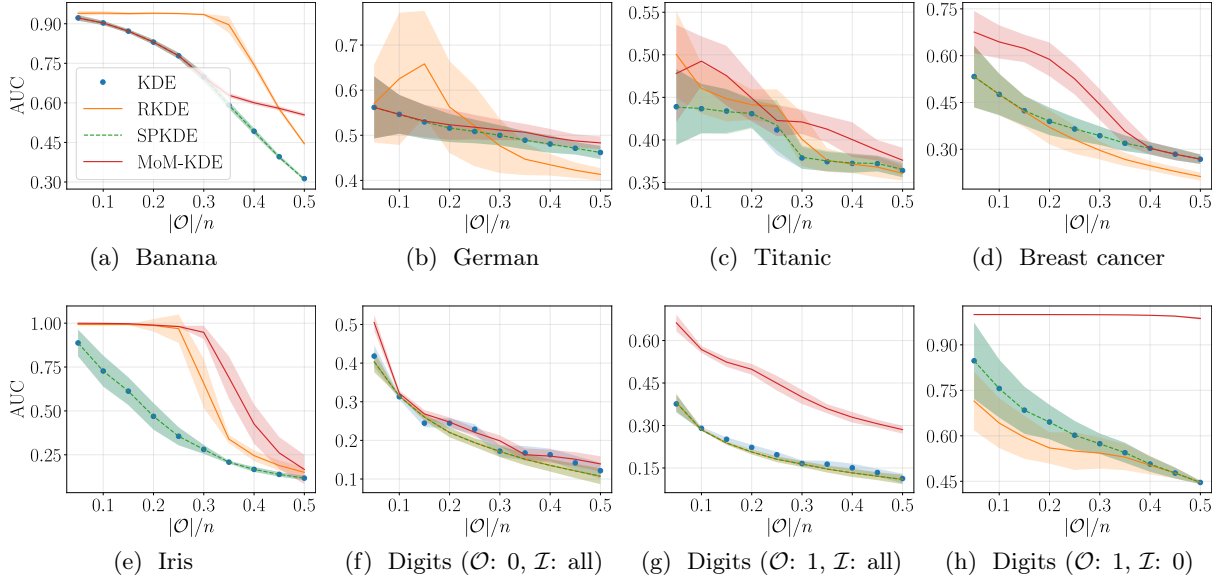


Figure 3: Anomaly detection with real datasets, measured with AUC over varying outlier proportion. A higher score means a better detection of the outliers. For Digits, we specify which classes are chosen to be inliers (\mathcal{I}) and outliers (\mathcal{O}).

class to reach a ratio $|\mathcal{O}|/n$ ranging from 0.05 to 0.5 with 0.05-wide steps. When a dataset does not contain enough outliers to reach a given ratio, we similarly downsample the inliers. For each dataset and each ratio, the experiments are performed 50 times, the random downsampling resulting in different learning datasets. The empirical performance is evaluated through the capacity of each estimator to detect anomalies, which we measure with the AUC.

Results are displayed in Figure 3. With the Digits dataset, we also explore additional scenarios with changing inlier and outlier classes (specified in figure titles). Overall, results are in line with performances observed over synthetic experiments, achieving good results in comparison to its competitors. Note that even in the highest dimensional scenarios, i.e. Digits and Breast cancer ($d = 64$ and $d = 30$), MoM-KDE still behaves well, outperforming its competitors. Additional results are reported in the supplementary materials.

5 Conclusion

The present paper introduced MoM-KDE, a new efficient way to perform robust kernel density estimation. The method has been shown to be consistent in both L_∞ and L_1 error-norm in presence of very generic outliers, enjoying a similar rate of convergence than the KDE without outliers. MoM-KDE achieved good empirical results in various situations while having a lower computational complexity than its competitors.

This work proposed to use the coordinate-wise median to construct its robust estimator. Future works will investigate the use of other generalization of the median in high dimension, e.g. the geometric median. In addition, further investigation will include a deeper statistical analysis under the hurdle contamination model in order to analyse the minimax optimality [Liu et al., 2019] of MoM-KDE.

References

- N. Alon, Y. Matias, and M. Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1):137–147, 1999.
- A. Backurs, P. Indyk, and T. Wagner. Space and time efficient kernel density estimation in high dimensions. In *Advances in Neural Information Processing Systems*, pages 15773–15782, 2019.
- M. Blum, R. W. Floyd, V. R. Pratt, R. L. Rivest, and R. E. Tarjan. Time bounds for selection. *Journal of Computer and System Sciences*, 7:448–461, 1973.

- L. Devroye and L. Györfi. *Nonparametric Density Estimation: The L1 View*. New York: John Wiley & Sons, 1985.
- L. Devroye and G. Lugosi. *Combinatorial methods in density estimation*. Springer Science & Business Media, 2012.
- L. Devroye, M. Lerasle, G. Lugosi, R. I. Oliveira, et al. Sub-gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.
- D. M. Endres and J. E. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information theory*, 49(7):1858–1860, 2003.
- E. Giné and A. Guillaou. Rates of strong uniform consistency for multivariate kernel density estimators. In *Annales de l’Institut Henri Poincaré (B) Probability and Statistics*, volume 38, pages 907–921. Elsevier, 2002.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- A. G. Gray and A. W. Moore. Nonparametric density estimation: Toward computational tractability. In *Proceedings of the 2003 SIAM International Conference on Data Mining*, pages 203–211. SIAM, 2003a.
- A. G. Gray and A. W. Moore. Rapid evaluation of multiple density models. In *AISTATS*, 2003b.
- P. J. Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992.
- M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- H. Jiang. Uniform convergence rates for kernel density estimation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1694–1703. JMLR. org, 2017.
- J. Kim and C. Scott. Robust kernel density estimation. In *2008 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3381–3384. IEEE, 2008.
- J. Kim and C. Scott. On the robustness of kernel density m-estimators. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 697–704. Citeseer, 2011.
- J. Kim and C. D. Scott. Robust kernel density estimation. *Journal of Machine Learning Research*, 13 (Sep):2529–2565, 2012.
- S. Kullback and R. A. Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22 (1):79–86, 1951.
- G. Lecué and M. Lerasle. Learning from mom’s principles: Le cam’s approach. *Stochastic Processes and their applications*, 129(11):4385–4410, 2019.
- G. Lecué, M. Lerasle, and T. Mathieu. Robust classification via mom minimization. *Machine Learning*, 2020.
- M. Lerasle. Lecture notes: Selected topics on robust statistical learning theory. *arXiv preprint arXiv:1908.10761*, 2019.
- F. Liese and I. Vajda. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, 2006.
- H. Liu, C. Gao, et al. Density estimation with contamination: minimax rates and theory of adaptation. *Electronic Journal of Statistics*, 13(2):3613–3653, 2019.
- S. Minsker et al. Geometric median and robust estimation in banach spaces. *Bernoulli*, 21(4):2308–2335, 2015.
- A. S. Nemirovsky and D. B. Yudin. *Problem complexity and method efficiency in optimization*. Wiley Interscience, New-York, 1983.

- E. Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830, 2011.
- P. Rigollet, R. Vert, et al. Optimal rates for plug-in estimators of density level sets. *Bernoulli*, 15(4): 1154–1178, 2009.
- D. W. Scott. *Multivariate density estimation: theory, practice, and visualization*. John Wiley & Sons, 2015.
- N. Segev, M. Harel, S. Mannor, K. Crammer, and R. El-Yaniv. Learn on source, refine on target: A model transfer learning framework with random forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1811–1824, Sep. 2017.
- B. W. Silverman. *Density estimation for statistics and data analysis*, volume 26. CRC press, 1986.
- B. Sriperumbudur and I. Steinwart. Consistency and rates for clustering with dbscan. In *Artificial Intelligence and Statistics*, pages 1090–1098, 2012.
- A. B. Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- B. A. Turlach. Bandwidth selection in kernel density estimation: A review. In *CORE and Institut de Statistique*. Citeseer, 1993.
- R. Vandermeulen and C. Scott. Consistency of robust kernel density estimators. In *Conference on Learning Theory*, pages 568–591, 2013.
- R. A. Vandermeulen and C. Scott. Robust kernel density estimation by scaling and projection in hilbert space. In *Advances in Neural Information Processing Systems*, pages 433–441, 2014.
- D. Wang, X. Lu, and A. Rinaldo. Dbscan: Optimal rates for density-based cluster estimation. *Journal of Machine Learning Research*, 20(170):1–50, 2019.
- Z. Wang and D. W. Scott. Nonparametric density estimation for high-dimensional data—algorithms and applications. *Wiley Interdisciplinary Reviews: Computational Statistics*, 11(4):e1461, 2019.

APPENDIX

A Technical proofs

Lemma 1. (*L_∞ error-bound of the KDE without anomalies*) Suppose that f belongs to the class of densities $\mathcal{P}(\alpha, L)$ defined as

$$\mathcal{P}(\alpha, L) \triangleq \left\{ f \mid f \geq 0, \int f(x)dx = 1, \text{ and } f \in \Sigma(\alpha, L) \right\},$$

where $\Sigma(\alpha, L)$ is the Hölder class of function on \mathbb{R}^d . Grant assumptions 1 to 4 and let $n > 1$, $h \in (0, 1)$ and $S \geq 1$ such that $nh^d \geq S$ and $nh^d \geq |\log(h)|$. Then with probability at least $1 - \exp(-S)$, we have

$$\|\hat{f}_n - f\|_\infty \leq C_1 \sqrt{\frac{S|\log(h)|}{nh^d}} + C_2 h^\alpha,$$

where $C_2 = L \int \|u\|^\alpha K(u)du < \infty$ and C_1 is a constant that only depends on $\|f\|_\infty$, the dimension d , and the kernel properties.

Proposition 1. (*L_∞ error-bound of the MoM-KDE under the $\mathcal{O} \cup \mathcal{I}$*) Suppose that f belongs to the class of densities $\mathcal{P}(\alpha, L)$ and grant assumptions 1 to 4. Let S be the number of blocks, $\delta > 0$ such that $S > (2 + \delta)|\mathcal{O}|$, and $\Delta = (1/(2 + \delta) - |\mathcal{O}|/S)$. Then, for any $h \in (0, 1)$, δ sufficiently small, and $n \geq 1$ such that $nh^d \geq S \log(2(2 + \delta)/\delta)$, and $nh^d \geq S|\log(h)|$, we have with probability at least $1 - \exp(-2\Delta^2 S)$,

$$\|\hat{f}_{MoM} - f\|_\infty \leq C_1 \sqrt{\frac{S \log\left(\frac{2(2+\delta)}{\delta}\right) |\log(h)|}{nh^d}} + C_2 h^\alpha,$$

where $C_2 = L \int \|u\|^\alpha K(u)du < \infty$ and C_1 is a constant that only depends on $\|f\|_\infty$, the dimension d , and the kernel properties.

Proof. From the definition of the MoM-KDE, we have the following implication [Lecué et al., 2020]

$$\left\{ \sup_x \left| \hat{f}_{MoM}(x) - f(x) \right| \geq \varepsilon \right\} \implies \left\{ \sup_x \sum_{k=1}^S I\left(\left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \geq S/2 \right\}.$$

Thus to upper-bound the probability of the left-hand event, it suffices to upper-bound the probability of the right-hand event. Moreover, we have

$$\begin{aligned} & \left| \hat{f}_{n_s}(x) - f(x) \right| \leq \sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| \\ \implies & I\left(\left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \leq I\left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \\ \implies & \sum_{k=1}^S I\left(\left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \leq \sum_{s=1}^S I\left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \\ \implies & \sup_x \sum_{s=1}^S I\left(\left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \leq \sum_{s=1}^S I\left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right), \end{aligned}$$

which implies that

$$\mathbb{P}\left(\sup_x \sum_{s=1}^S I\left(\left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \geq S/2\right) \leq \mathbb{P}\left(\sum_{s=1}^S I\left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right) \geq S/2\right).$$

Let $Z_s = I\left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon\right)$ and let $\mathcal{S} = \{s \in \{1, \dots, S\} \mid B_s \cap \mathcal{O} = \emptyset\}$ i.e. the set of indices s such that the block B_s does not contain any outliers. Since $\sum_{s \in \mathcal{S}^c} I(\cdot)$ is bounded by $|\mathcal{O}|$, almost surely, the following holds.

$$\begin{aligned}
\sum_{s=1}^S I \left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon \right) &= \sum_{s=1}^S Z_s = \sum_{s \in \mathcal{S}} Z_s + \sum_{s \in \mathcal{S}^c} Z_s \\
&\leq \sum_{s \in \mathcal{S}} Z_s + |\mathcal{O}| \\
&= \sum_{s \in \mathcal{S}} [Z_s - \mathbb{E}(Z_s) + \mathbb{E}(Z_s)] + |\mathcal{O}| \\
&= \sum_{s \in \mathcal{S}} [Z_s - \mathbb{E}(Z_s)] + \sum_{s \in \mathcal{S}} \mathbb{E}(Z_s) + |\mathcal{O}| \\
&\leq \sum_{s=1}^S [Z_s - \mathbb{E}(Z_s)] + S \cdot \mathbb{E}(Z_1) + |\mathcal{O}| \\
&\leq \sum_{s=1}^S [Z_s - \mathbb{E}(Z_s)] + S \cdot \mathbb{P} \left(\sup_x \left| \hat{f}_{n_1}(x) - f(x) \right| > \varepsilon \right) + |\mathcal{O}|, \tag{8}
\end{aligned}$$

where Z_1 is assumed, without loss of generality, to be associated to a block not containing outliers. This block always exists thanks to the hypothesis $S > (2 + \delta)|\mathcal{O}|$.

Let $\varepsilon = C_1 \sqrt{\frac{S \log(\frac{2(2+\delta)}{\delta}) |\log(h)|}{nh^d}} + C_2 h^\alpha$, then using Lemma 1 with $S = \log(\frac{2(2+\delta)}{\delta})$, we have

$$\mathbb{P} \left(\sup_x \left| \hat{f}_{n_1}(x) - f(x) \right| > \varepsilon \right) \leq \frac{\delta}{2(2 + \delta)}.$$

Combining this last inequality with equation (8) leads to

$$\begin{aligned}
\mathbb{P} \left(\sum_{s=1}^S I \left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon \right) \geq S/2 \right) &\leq \mathbb{P} \left(\sum_{s=1}^S [Z_s - \mathbb{E}(Z_s)] + S \cdot \frac{\delta}{2(2 + \delta)} + |\mathcal{O}| \geq S/2 \right) \\
&\leq \mathbb{P} \left(\sum_{s=1}^S [Z_s - \mathbb{E}(Z_s)] \geq S \left(\frac{1}{2} - \frac{\delta}{2(2 + \delta)} - \frac{|\mathcal{O}|}{S} \right) \right) \\
&\leq \mathbb{P} \left(\sum_{s=1}^S [Z_s - \mathbb{E}(Z_s)] \geq S \left(\frac{1}{2 + \delta} - \frac{|\mathcal{O}|}{S} \right) \right)
\end{aligned}$$

Tacking $\Delta = \left(\frac{1}{2 + \delta} - \frac{|\mathcal{O}|}{S} \right) > 0$ and applying Hoeffding's inequality to the right-hand side of the previous equation gives

$$\mathbb{P} \left(\sum_{s=1}^S I \left(\sup_x \left| \hat{f}_{n_s}(x) - f(x) \right| > \varepsilon \right) \geq S/2 \right) \leq e^{-2S\Delta^2},$$

which concludes the proof. \square

Proposition 2. (*L_1 -consistency in probability*) If $n/S \rightarrow \infty$, $h \rightarrow 0$, $nh^d \rightarrow \infty$, and $S > 2|\mathcal{O}|$, then

$$\|\hat{f}_{MoM} - f\|_1 \xrightarrow[n \rightarrow \infty]{\mathcal{P}} 0.$$

Proof. We first rewrite the MoM-KDE as

$$\hat{f}_{MoM}(x) = \sum_{s=1}^S \hat{f}_{n_s}(x) I_{A_s}(x),$$

where $A_s = \{x \mid \hat{f}_{MoM}(x) = \hat{f}_{n_s}(x)\}$. Without loss of generality, we assume that

$$A_k \bigcap_{s \neq \ell} A_\ell = \emptyset, \quad \bigcup_{s=1}^S A_s = \mathbb{R}^d, \quad \text{and} \quad \sum_{s=1}^S I_{A_s}(x) = 1.$$

$$\begin{aligned} \int \left| \hat{f}_{MoM}(x) - f(x) \right| dx &= \int \left| \sum_{s=1}^S \hat{f}_{n_s}(x) I_{A_s}(x) - f(x) \right| dx \\ &= \int \left| \sum_{s=1}^S (\hat{f}_{n_s}(x) - f(x)) I_{A_s}(x) \right| dx \\ &\leq \int \sum_{s=1}^S \left| \hat{f}_{n_s}(x) - f(x) \right| I_{A_s}(x) dx \\ &= \sum_{s=1}^S \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx \\ &= \sum_{s \in \mathcal{S}} \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx + \sum_{s \in \mathcal{S}^C} \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx. \end{aligned} \quad (9)$$

From the L_1 -consistency of the KDE in probability, if the number of anomalies grows at a small enough speed [Devroye and Gyorfi, 1985], the left part is bounded, i.e.

$$\sum_{s \in \mathcal{S}} \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx \leq \sum_{s \in \mathcal{S}} \int \left| \hat{f}_{n_s}(x) - f(x) \right| dx \xrightarrow[n \rightarrow \infty]{\mathcal{P}} 0. \quad (10)$$

We now upper-bound the right part of equation (9). Let consider a particular block A_s where $s \in \mathcal{S}^C$. In this block, the estimator \hat{f}_{n_s} is selected and is calculated with samples containing anomalies. As $\forall x \in A_s$, $\hat{f}_{n_s}(x)$ is the median (by definition), if $S > 2|\mathcal{O}|$, we can always find a $s' \in \mathcal{S}$ such that $\hat{f}_{n_s}(x) \leq \hat{f}_{n_{s'}}(x)$ or $\hat{f}_{n_s}(x) \geq \hat{f}_{n_{s'}}(x)$.

Now let denote by $A_s^+ = \{x \in A_s \mid \hat{f}_{n_s}(x) \geq f(x)\}$ and $A_s^- = \{x \in A_s \mid \hat{f}_{n_s}(x) < f(x)\}$. We have $A_s^+ \cup A_s^- = A_s$ and each one of these blocks can be decomposed respectively into $|\mathcal{S}|$ sub-blocks (not necessarily disjoint) $\{A_s^{s',+}\}_{s' \in \mathcal{S}}$ and $\{A_s^{s',-}\}_{s' \in \mathcal{S}}$ such that $\forall s' \in \mathcal{S}$, $A_s^{s',+} = \{x \in A_s \mid \hat{f}_{n_{s'}}(x) \geq \hat{f}_{n_s}(x) \geq f(x)\}$ and $A_s^{s',-} = \{x \in A_s \mid \hat{f}_{n_{s'}}(x) \leq \hat{f}_{n_s}(x) < f(x)\}$. Finally, the right-hand term of equation (9) can be upper-bounded by

$$\begin{aligned} \sum_{s \in \mathcal{S}^C} \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx &\leq \sum_{s \in \mathcal{S}^C} \int_{A_s^+} \left| \hat{f}_{n_s}(x) - f(x) \right| dx + \int_{A_s^-} \left| \hat{f}_{n_s}(x) - f(x) \right| dx \\ &\leq \sum_{s \in \mathcal{S}^C} \sum_{s' \in \mathcal{S}} \int_{A_s^{s',+}} \left| \hat{f}_{n_s}(x) - f(x) \right| dx + \int_{A_s^{s',-}} \left| \hat{f}_{n_s}(x) - f(x) \right| dx \\ &\leq \sum_{s \in \mathcal{S}^C} \sum_{s' \in \mathcal{S}} \int_{A_s^{s',+}} \left| \hat{f}_{n_{s'}}(x) - f(x) \right| dx + \int_{A_s^{s',-}} \left| \hat{f}_{n_{s'}}(x) - f(x) \right| dx \\ &\leq \sum_{s \in \mathcal{S}^C} \sum_{s' \in \mathcal{S}} \int \left| \hat{f}_{n_{s'}}(x) - f(x) \right| dx + \int \left| \hat{f}_{n_{s'}}(x) - f(x) \right| dx. \end{aligned}$$

Since $\forall s' \in \mathcal{S}$ we have $\int \left| \hat{f}_{n_{s'}}(x) - f(x) \right| dx \xrightarrow[n \rightarrow \infty]{\mathcal{P}} 0$, we can conclude using similar arguments as those used for (10) that $\sum_{s \in \mathcal{S}^C} \int_{A_s} \left| \hat{f}_{n_s}(x) - f(x) \right| dx \xrightarrow[n \rightarrow \infty]{\mathcal{P}} 0$, which concludes the proof. \square

B Additional results

As stated in the main paper, we display here the additional results containing:

- For synthetic data, the Kullback-Leibler divergence in both directions, i.e. $D_{\text{KL}}(\hat{f}, f)$ and $D_{\text{KL}}(f, \hat{f})$, and the ROC AUC measuring the performance of an anomaly detector based on \hat{f} . Results are displayed on Figure 4.
- For Digits data, the ROC AUC measuring the performance of an anomaly detector based on \hat{f} . As stated in the main paper, this is done under multiple scenarios, where outliers and inliers can be chosen among the nine available classes. Here we show the AUC when the outliers are set as one class (class 2 to class 9), and inliers are set as “the rest” of all classes. Results are displayed on Figure 5.

Synthetic data. When considering the Kullback-Leibler divergence, results lead to a very similar conclusion as previously stated, that is, an overall good performance of MoM-KDE while its competitors, notably SPKDE, are more data-dependent. When the density estimate \hat{f} is used in a simple anomaly detector, results are quite different. Indeed, when outliers are uniformly distributed, even if MoM-KDE seems to better estimate the true density (according to D_{JS} and D_{KL}), this doesn’t make \hat{f}_{MoM} a better anomaly detector. It seems that in this case, the outliers are either easily detected because distant from the density estimate, or located in dense regions, thus making them impossible to identify, and this for all density estimates provided by competitors. In the case of adversarial contamination, the conclusion is quite similar. Although MoM-KDE better fits the true density, the situation is extremely difficult for anomaly detection, hence making all competitors yield very poor results. In the two other cases – Gaussian outlier, anomaly detection results follow the density estimation.

Digits data. Results over Digits scenarios are inline with main conclusions over real data. Although from one scenario to another, all methods have varied results, the overall observation is that MoM-KDE is either similar or better than its competitors.

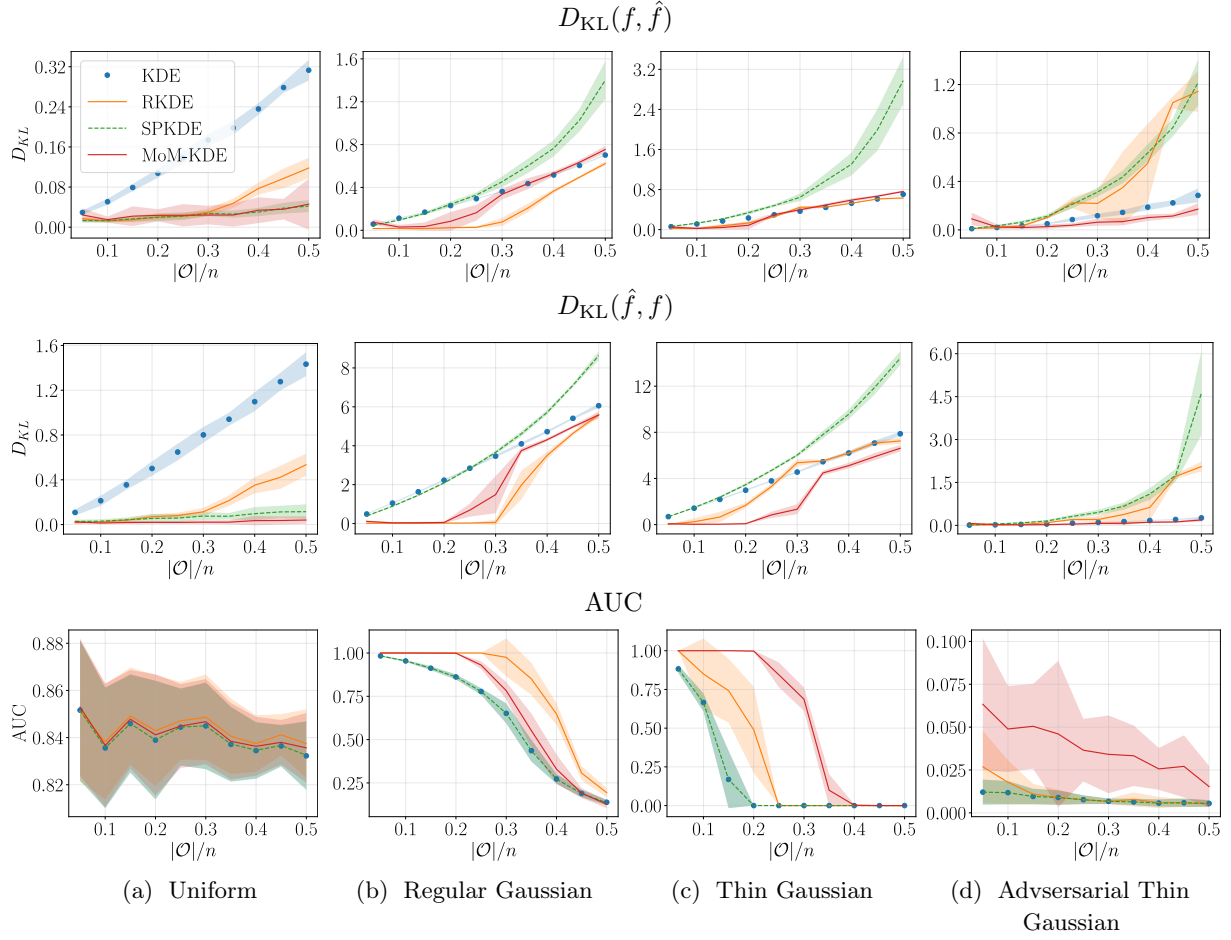


Figure 4: Density estimation with synthetic data. The displayed metrics are the Kullback-Leibler divergence (a lower score means a better estimation of the true density) and the AUC (a higher score means a better detection of the outliers).

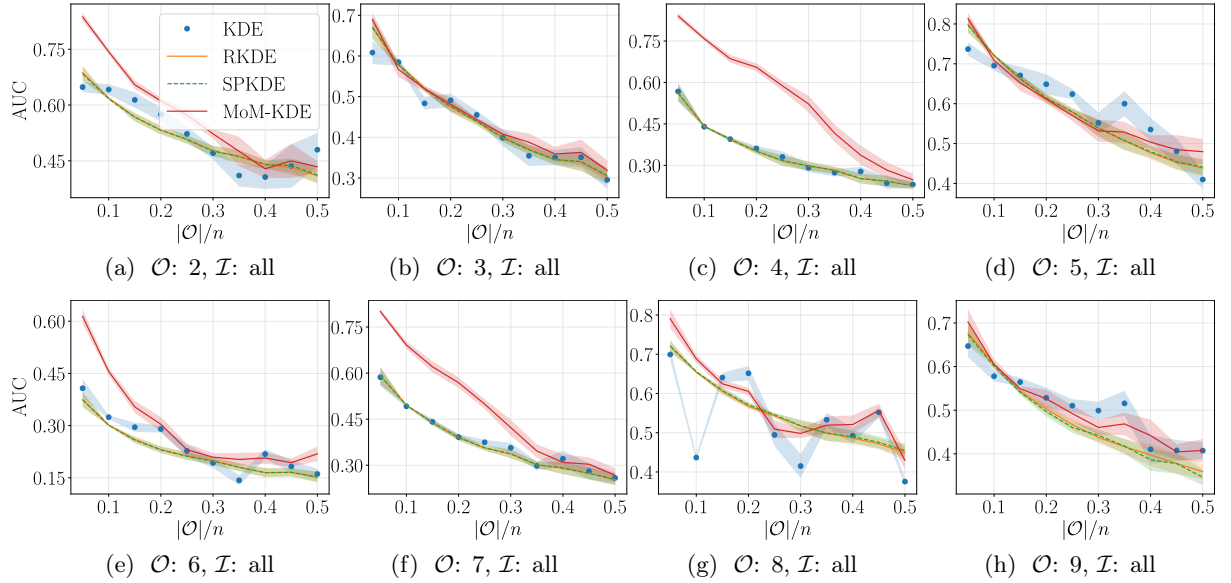


Figure 5: Anomaly detection with Digits data, measured with AUC over varying outlier proportion. A higher score means a better detection of the outliers. We specify which classes are chosen to be inliers (\mathcal{I}) and outliers (\mathcal{O}).