

Asymptotic Analysis of the Jittering Kernel Density Estimator

T. Nagler^{1*}

¹*Dept. of Math. Technical Univ. of Munich, Munich, Germany*

Received September 4, 2017; in final form, November 7, 2017

Abstract—Jittering estimators are nonparametric function estimators for mixed data. They extend arbitrary estimators from the continuous setting by adding random noise to discrete variables. We give an in-depth analysis of the jittering kernel density estimator, which reveals several appealing properties. The estimator is strongly consistent, asymptotically normal, and unbiased for discrete variables. It converges at minimax-optimal rates, which are established as a by-product of our analysis. To understand the effect of adding noise, we further study its asymptotic efficiency and finite sample bias in the univariate discrete case. Simulations show that the estimator is competitive on finite samples. The analysis suggests that similar properties can be expected for other jittering estimators.

Keywords: density, discrete, jittering, kernel, minimax, mixed data.

2000 Mathematics Subject Classification: 62G07.

DOI: 10.3103/S1066530718010027

1. INTRODUCTION

Multivariate density estimation is a central field in nonparametric statistics. Yet many popular methods have a significant drawback in applications: they can only be applied to continuous data. Some estimators have been specifically designed to allow for mixed continuous and discrete data [1, 2, 6, 11], but the number is small compared to the methods available in a purely continuous framework.

A common trick among practitioners is to make the discrete variables continuous by adding a small amount of noise. The noisy data is continuous and the usual nonparametric estimators apply. But the addition of random noise can introduce bias, so this procedure generally lacks justification. Nagler [14] showed that adding noise still allows for valid estimates when the noise comes from a certain class of distributions. Then any nonparametric density estimator can be used in the mixed data setting. The resulting estimators are called *jittering estimators*.

Jittering estimators have so far been neglected in academic research, likely due to the widespread concern that jittering causes a loss in efficiency. The main objective of this article is to demonstrate that this concern is usually unjustified. To this end, we give an in-depth analysis of a simple instance from the class of jittering estimators: the *jittering kernel density estimator*, which is the jittering analog of the classical kernel density estimator [15, 17, 23]. We shall show that it maintains all the properties expected from a good nonparametric density estimator:

1. It is asymptotically normal and asymptotically unbiased for discrete variables (Theorem 1).
2. It is strongly and uniformly consistent (Theorem 2).
3. It is relatively efficient, even fully efficient in specific cases (Section 4.1).
4. It converges at minimax-optimal rates for a large class of target densities (Theorems 3 and 4). To the best of the author's knowledge, these are the first results on minimax-optimality of nonparametric density estimators for mixed data.

*E-mail: thomas.nagler@tum.de

Although focus is on only one instance of the class of jittering estimators, we can expect that others have similar properties.

The remainder of this article is organized as follows. Section 2 introduces the jittering estimator and some assumptions. Section 3 gives a comprehensive asymptotic analysis which is complemented by a study of the asymptotic efficiency and finite sample bias in the univariate discrete setting (Section 4). Section 5 establishes minimax-optimal rates for density estimation in a nonparametric mixed data model. Section 6 demonstrates that the estimator is also competitive on finite samples; Section 7 offers conclusions. Proofs of all theorems are deferred to Appendix A.

2. THE ESTIMATOR

Suppose that (\mathbf{Z}, \mathbf{X}) is a random vector with discrete component $\mathbf{Z} \in \mathbb{Z}^p$ and continuous component $\mathbf{X} \in \mathbb{R}^q$. We explicitly allow for the cases where $p \geq 1, q = 0$ (all variables are discrete) and $p = 0, q \geq 1$ (all variables are continuous). Our goal is to estimate the density f of (\mathbf{Z}, \mathbf{X}) based on ‘observations’ $(\mathbf{Z}_i, \mathbf{X}_i), i = 1, \dots, n$, which are *iid* random vectors having the same distribution as (\mathbf{Z}, \mathbf{X}) . In this context, f is the density with respect to the product of the counting and Lebesgue measures, i.e.,

$$f_{\mathbf{Z}, \mathbf{X}}(\mathbf{z}, \mathbf{x}) = \frac{\partial^q}{\partial x_1 \cdots \partial x_q} \Pr(\mathbf{Z} = \mathbf{z}, \mathbf{X} \leq \mathbf{x}).$$

Let K be a real-valued function, called *kernel*, and abbreviate $K(\mathbf{w}) = \prod_{j=1}^k K(w_j)$ for any $\mathbf{w} \in \mathbb{R}^k, k \in \mathbb{N}$. The classical kernel density estimator is defined as

$$\hat{f}(\mathbf{z}, \mathbf{x}) = \frac{1}{nh_n^p b_n^q} \sum_{i=1}^n K\left(\frac{\mathbf{Z}_i - \mathbf{z}}{h_n}\right) K\left(\frac{\mathbf{X}_i - \mathbf{x}}{b_n}\right), \quad (1)$$

where $h_n, b_n > 0$ are called *bandwidth parameters* and control the amount of smoothing. The above definition of the estimator is simplified to ease our exposition: we use only one parameter (h_n) for smoothing all components of \mathbf{Z} and one parameter (b_n) for smoothing the components of \mathbf{X} . In practice, one would use a single parameter for each variable or even a bandwidth matrix (see, e.g., [18]).

The estimator \hat{f} only works for continuous random vectors. To make it applicable to mixed data, we make all discrete variables continuous by adding noise. Let $\mathbf{E}_i \in \mathbb{R}^p, i = 1, \dots, n$, be *iid* random vectors independent of $(\mathbf{Z}_i, \mathbf{X}_i), i = 1, \dots, n$. Suppose further that the p components of \mathbf{E}_i are *iid* with density η . The jittering kernel density estimator is defined as the classical kernel density estimator applied to $(\mathbf{Z}_i + \mathbf{E}_i, \mathbf{X}_i), i = 1, \dots, n$:

$$\tilde{f}(\mathbf{z}, \mathbf{x}) = \frac{1}{nh_n^p b_n^q} \sum_{i=1}^n K\left(\frac{\mathbf{Z}_i + \mathbf{E}_i - \mathbf{z}}{h_n}\right) K\left(\frac{\mathbf{X}_i - \mathbf{x}}{b_n}\right). \quad (2)$$

To facilitate our analysis, the following conditions are imposed on the kernel function:

Assumptions.

K1: $K: [-1, 1] \rightarrow \mathbb{R}_{\geq 0}$ is a continuous function satisfying $\int K(t) dt = 1$.

K2: There is $\ell \in \mathbb{N}, \ell \geq 2$, such that for $k = 1, \dots, \ell - 1$,

$$\int_{[0,1]} t^k K(t) dt = 0, \quad \int_{[0,1]} t^\ell K(t) dt > 0.$$

Remark 1. A kernel function satisfying K2 is called ℓ th order kernel (see, e.g., [12]).

We further assume that the noise density η belongs to the class $\mathcal{E}_{\gamma_1, \gamma_2}$, as defined in [14].

Definition 1. We say that $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$ for some $0 < \gamma_1 \leq 0.5 \leq \gamma_2 < 1$ if

(i) η is an absolutely continuous probability density function,

- (ii) $\eta(x) = 1$ for all $x \in [-\gamma_1, \gamma_1]$,
- (iii) $\eta(x) = 0$ for all $x \in \mathbb{R} \setminus (-\gamma_2, \gamma_2)$.

The density of $(\mathbf{Z} + \mathbf{E}, \mathbf{X})$ is given by

$$f_\eta(\mathbf{z}, \mathbf{x}) = \sum_{\mathbf{z}' \in \mathbb{Z}^p} f(\mathbf{z}', \mathbf{x}) \prod_{j=1}^p \eta(z_j - z'_j), \quad (\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^p \times \mathbb{R}^q.$$

The class $\mathcal{E}_{\gamma_1, \gamma_2}$ ensures that f_η is well behaved. The most important properties are summarized in the following result, see Propositions 1 and 2 in [14].

Lemma 1. *Suppose the components of \mathbf{E} are iid with density $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$. Then the joint density f_η of $(\mathbf{Z} + \mathbf{E}, \mathbf{X})$ satisfies for all $(\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^p \times \mathbb{R}^q$, and $\mathbf{m} \in \mathbb{N}^p$ such that $\sum_{k=1}^p m_k = \bar{m}$,*

$$f_\eta(\mathbf{z}, \mathbf{x}) = f(\mathbf{z}, \mathbf{x}), \quad \frac{\partial^{\bar{m}} f_\eta(\mathbf{z}, \mathbf{x})}{\partial z_1^{m_1} \dots \partial z_p^{m_p}} = 0.$$

The first equality implies that we can equivalently estimate f_η instead of f . This is convenient because f_η is the density of a purely continuous random vector. The second equality states that all derivatives w.r.t. \mathbf{z} vanish, which makes estimation even easier.

Remark 2. The estimator \tilde{f} is similar to the estimators in [1] and [11]. The difference lies in the kernel function for discrete data. The estimators in [1] and [11] use a deterministic kernel function which is defined on the integers. In contrast, the jittering kernel density estimator (2) uses a random kernel $K\{(\cdot + \mathbf{E}_i)/b_n\}$ defined on a compact subset of \mathbb{R}^p where randomness is induced by \mathbf{E}_i .

3. ASYMPTOTIC ANALYSIS IN THE GENERAL SETTING

3.1. Asymptotic Distribution

We first study the asymptotic distribution of the jittering kernel density estimator. To motivate our first theorem, we recall a classical result from kernel density estimation in the purely continuous setting (see, e.g., [23]). If f is the density of a continuous random vector (\mathbf{Z}, \mathbf{X}) , sufficiently smooth, $\ell = 2$, and $h_n, b_n \rightarrow 0, nh_n^p b_n^q \rightarrow \infty$, then

$$\begin{aligned} \mathbb{E}\{\hat{f}(\mathbf{z}, \mathbf{x})\} &= f(\mathbf{z}, \mathbf{x}) + \frac{h_n^2 \sigma_2}{2} \sum_{k=1}^p \frac{\partial^2 f(\mathbf{z}, \mathbf{x})}{\partial z_k^2} + \frac{b_n^2 \sigma_2}{2} \sum_{j=1}^q \frac{\partial^2 f(\mathbf{z}, \mathbf{x})}{\partial x_j^2} + o(h_n^2 + b_n^2), \\ \text{Var}\{\hat{f}(\mathbf{z}, \mathbf{x})\} &= \frac{\kappa^{p+q} f(\mathbf{z}, \mathbf{x})}{nh_n^p b_n^q} + o\left(\frac{1}{nh_n^p b_n^q}\right). \end{aligned} \quad (3)$$

Recall that \tilde{f} is nothing else than \hat{f} applied to $(\mathbf{Z}_i + \mathbf{E}_i, \mathbf{X}_i)$, $i = 1, \dots, n$. Lemma 1 shows that $f_\eta(\mathbf{z}, \mathbf{x})$, the density of $(\mathbf{Z}_i + \mathbf{E}_i, \mathbf{X}_i)$, has vanishing derivatives with respect to \mathbf{z} . We can thus expect the first sum in the bias term in (3) to vanish asymptotically. In fact, it becomes exactly zero when $h_n \leq \min\{\gamma_1, 1 - \gamma_2\}$. The following result improves upon the properties implied by (3) by taking these considerations into account.

Assumptions.

- A1: $f(\mathbf{z}, \mathbf{x})$ is $\ell + 1$ times continuously differentiable with respect to \mathbf{x} .
- A2: $K1$ and $K2$ hold with $\ell \geq 2$.
- A3: $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$.
- A4: $b_n \rightarrow 0$ and $nh_n^p b_n^q \rightarrow \infty$ as $n \rightarrow \infty$.
- A5: There is $n_0 \in \mathbb{N}$ such that $h_n \leq \min\{\gamma_1, 1 - \gamma_2\}$ for all $n \geq n_0$.

Theorem 1. Under assumptions A1–A5, it holds for any $(\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^p \times \mathbb{R}^q$,

$$\begin{aligned} \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= f(\mathbf{z}, \mathbf{x}) + \frac{b_n^\ell \sigma_\ell}{\ell!} \sum_{j=1}^q \frac{\partial^\ell f(\mathbf{z}, \mathbf{x})}{\partial x_j^\ell} + o(b_n^\ell), \\ \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= \frac{f(\mathbf{z}, \mathbf{x})}{nb_n^q} \{h_n^{-p} \kappa^{p+q} - b_n^q f(\mathbf{z}, \mathbf{x})\} + o\left(\frac{1}{nh_n^p b_n^q}\right), \end{aligned}$$

where $\sigma_\ell = \int_{-1}^1 s^\ell K(s) ds$ and $\kappa = \int_{-1}^1 K^2(s) ds$. If further $nh_n^p b_n^{q+2\ell} = O(1)$,

$$\frac{\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}}{\text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}} \xrightarrow{d} \mathcal{N}(0, 1).$$

Remark 3. The assumptions in Theorem 1 differ from those usually made in the continuous framework. There are no assumptions on the smoothness of $\hat{f}_{\mathbf{z}+\mathbf{E}, \mathbf{X}}(\mathbf{z}, \mathbf{x})$ with respect to \mathbf{z} , because its local behavior is controlled by $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$. Further, h_n is not required to vanish asymptotically, but should be less than $\min\{\gamma_1, 1 - \gamma_2\}$ for large n . This is sufficient to ensure that there is no bias with respect to \mathbf{z} . Further decreasing h_n does not change the bias, but inflates the variance.

Remark 4. The asymptotic variance does not involve η or its class parameters γ_1 and γ_2 (and neither does the asymptotic bias). Intuitively, we would expect an increase in the estimator's variance because we are adding random noise. Apparently this effect is dominated by the sampling variability in the original data and asymptotically negligible. So there should be no benefit from averaging over multiple jitters (at least asymptotically). This is in contrast to empirical processes of jittered data [4].

3.2. Asymptotically Optimal Bandwidths

A standard tool for studying optimal bandwidths is the *asymptotic mean squared error*,

$$\text{AMSE}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} = [\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x})]^2 + \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}.$$

Under the assumptions of Theorem 1, we get

$$\text{AMSE}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} \approx \frac{b_n^{2\ell} \sigma_\ell^2}{(\ell!)^2} \left(\sum_{j=1}^q \frac{\partial^\ell f(\mathbf{z}, \mathbf{x})}{\partial x_j^\ell} \right)^2 + \frac{f(\mathbf{z}, \mathbf{x})}{nb_n^q} \{h_n^{-p} \kappa^{p+q} - b_n^q f(\mathbf{z}, \mathbf{x})\}.$$

For $h_n = O(1)$, it is easy to check that the bandwidth b_n minimizing the AMSE satisfies $b_n \sim n^{-1/(2\ell+q)}$. This is well known as the optimal rate for the classical kernel density estimator when $p = 0$. The AMSE further suggests that it is optimal to choose h_n as large as possible. The largest h_n allowed by A5 is $h_n = \min\{\gamma_1, 1 - \gamma_2\}$. Asymptotically, this is the optimal bandwidth. We shall see shortly that this choice means that we are not smoothing the discrete variables at all. This is not unreasonable: in contrast to the continuous case, smoothing discrete variables is not necessary for consistent nonparametric estimation (for a discussion, see [19]).

On finite samples $h_n = \min\{\gamma_1, 1 - \gamma_2\}$ can be too small. If $h_n \leq \min\{\gamma_1, 1 - \gamma_2\}$, the estimator can be written as

$$\tilde{f}(\mathbf{z}, \mathbf{x}) = \frac{1}{nh_n^p b_n^q} \sum_{i: \mathbf{Z}_i = \mathbf{z}} K\left(\frac{\mathbf{E}_i}{h_n}\right) K\left(\frac{\mathbf{X}_i - \mathbf{x}}{b_n}\right).$$

Indeed, the estimator neglects all observations where $\mathbf{Z}_i \neq \mathbf{z}$ and, thus, does not smooth with respect to the discrete variables. This also means that $\tilde{f}(\mathbf{z}, \mathbf{x}) = 0$ if $\mathbf{Z}_i \neq \mathbf{z}$ for all $i = 1, \dots, n$. Theorem 1 implicitly assumes that n is large enough to provide sufficiently many observations with $\mathbf{Z}_i = \mathbf{z}$. This is guaranteed asymptotically whenever $P(\mathbf{Z} = \mathbf{z}) > 0$, but often demands sample sizes much larger than what is common.

We conclude that Theorem 1 is not useful for bandwidth selection on samples of small or moderate size. Cross-validation techniques are more appropriate tools in the mixed data setting (see, e.g., [5, 10, 16]).

3.3. Consistency

Theorem 1 implies pointwise consistency of the jittering kernel density estimator, but Assumption A1 is more strict than necessary. The following result weakens this assumption and additionally establishes strong uniform consistency.

Assumption.

A1': The $(\ell - 1)$ th derivative of $f(\mathbf{z}, \mathbf{x})$ exists and is uniformly Lipschitz on $S \subseteq \mathbb{Z}^p \times \mathbb{R}^q$.

Theorem 2. Suppose that Assumptions A1', A2–A5 hold. Then, for all $(\mathbf{z}, \mathbf{x}) \in S$,

$$\tilde{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x}) = O_p\{b_n^\ell + (nh_n^p b_n^q)^{-1/2}\}, \quad (4)$$

$$\sup_S |\tilde{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})| = O_{a.s.}\left\{b_n^\ell + \left(\frac{\max\{\log \log n, \log h_n^{-1}, \log b_n^{-1}\}}{nh_n^p b_n^q}\right)^{1/2}\right\}. \quad (5)$$

Remark 5. If there are $h_0 > 0$, $n_0 \in \mathbb{N}$ such that $h_n \in (h_0, \min\{\gamma_1, 1 - \gamma_2\}]$ for all $n \geq n_0$, the rates of convergence in Theorem 2 do not involve p , the dimension of the discrete variables. So adding more discrete variables does not change the convergence rate of the estimator. In particular, there is no cost for recoding unordered categorical variables into several binary variables.

Remark 6. (i) The best rate in (4) is $n^{-\ell/(2\ell+q)}$ and is achieved when $h_n \sim 1$ and $b_n \sim n^{-1/(2\ell+q)}$.

(ii) For $q > 0$, the best rate in (5) is $(n/\log n)^{-\ell/(2\ell+q)}$ and is achieved when $h_n \sim 1$, $b_n \sim (n/\log n)^{-1/(2\ell+q)}$.

(iii) For $q = 0$, the best rate in (5) is $(n/\log \log n)^{-1/2}$ and is achieved when $h_n \sim 1$.

4. A CLOSER LOOK AT THE UNIVARIATE DISCRETE SETTING

The jittering kernel density estimator \tilde{f} handles continuous variables just like the classical kernel density estimator. How it smooths discrete variables is less obvious. To gain a better understanding, we study its asymptotic efficiency and finite sample bias when there is only one discrete variable ($p = 1$, $q = 0$).

4.1. Asymptotic Efficiency

For convenience, set $h_n \equiv \min\{\gamma_1, 1 - \gamma_2\}$. The expectation and variance in Theorem 1 become

$$\mathbb{E}\{\tilde{f}(z)\} = f(z), \quad \text{Var}\{\tilde{f}(z)\} = \frac{f(z)}{n} [\min(\gamma_1, 1 - \gamma_2)^{-1} \kappa - f(z)] + o(n^{-1}),$$

The most efficient point estimator for a discrete probability $f(z)$ is the sample frequency $f_n(z) = n^{-1} \sum_{i=1}^n \mathbf{1}(Z_i = z)$. It satisfies

$$\mathbb{E}\{f_n(z)\} = f(z), \quad \text{Var}\{f_n(z)\} = \frac{f(z)}{n} \{1 - f(z)\}.$$

The *asymptotic relative efficiency* (ARE) of \tilde{f} relative to f_n is defined as

$$\text{ARE}\{\tilde{f}(z) : f_n(z)\} = \frac{\text{AVar}\{f_n(z)\}}{\text{AVar}\{\tilde{f}(z)\}},$$

where AVar denotes the leading term of an asymptotic expansion of the variance. The ARE is interpreted as follows: If the estimator \tilde{f} is used with n observations, then one needs $\text{ARE} \times n$ observations to obtain the same accuracy with f_n . If the ARE is less than one, then f_n needs less observations, i.e., f_n is more efficient than \tilde{f} . If the ARE is greater than one, it is the other way around. If it is exactly one, the two estimators are equally efficient.

Straightforward calculations yield

$$\begin{aligned} \text{ARE}\{\tilde{f}(z) : f_n(z)\} &= \frac{1 - f(z)}{\min\{\gamma_1, 1 - \gamma_2\}^{-1}\kappa - f(z)} \\ &= \left(1 + \frac{\min\{\gamma_1, 1 - \gamma_2\}^{-1}\kappa - 1}{1 - f(z)}\right)^{-1} \leq 1. \end{aligned}$$

The relative efficiency depends on three quantities:

- It is increasing in $\min\{\gamma_1, 1 - \gamma_2\}$ and the most efficient choice is $\gamma_1 = \gamma_2 = 1/2$, which corresponds to the uniform error density on $(-1/2, 1/2)$. On the other hand, the relative efficiency approaches 0 for $\gamma_1 \rightarrow 0$ or $\gamma_2 \rightarrow 1$.
- It is decreasing in κ , which is the roughness of the kernel K . The ‘least rough’ kernel is the uniform kernel, i.e., $K(x) = 2^{-1}\mathbf{1}(|x| \leq 1)$, for which $\kappa = 1/2$. But this kernel is rather unpopular in practice. A more widely used kernel is the Epanechnikov kernel,

$$K(x) = 3/4(1 - x^2)\mathbf{1}(|x| \leq 1) \quad \text{for which } \kappa = 0.6.$$

- It is decreasing in $f(z)$. The worst case is that $f(z) = 1$, for which the ARE is zero. For a Bernoulli(1/2) variable, Uniform($-1/2, 1/2$) noise, and the Epanechnikov kernel, we get $\text{ARE} \approx 0.71$.

Remark 7. Suppose η is the uniform density on $(-1/2, 1/2)$ (for which $\gamma_1 = \gamma_2 = 1/2$), $h_n = 1/2$, and K is the uniform kernel (for which $\kappa = 1/2$). Then the two estimators are equally efficient. In fact, the estimator \tilde{f} becomes

$$\begin{aligned} \tilde{f}(z) &= \frac{1}{nh_n} \sum_{i=1}^n 2^{-1}\mathbf{1}(|Z_i + E_i - z| \leq h_n) \\ &= \frac{2}{n} \sum_{i=1}^n 2^{-1}\mathbf{1}(|Z_i + E_i - z| \leq 1/2) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(Z_i = z), \end{aligned}$$

which is exactly the sample frequency estimator f_n .

4.2. Finite sample bias

Assuming $h_n \leq \min\{\gamma_1, 1 - \gamma_2\}$, Theorem 1 shows that \tilde{f} is unbiased in a purely discrete setting. On small samples, it is often necessary to choose a larger bandwidth (see Section 3.2). When $h_n > \min\{\gamma_1, 1 - \gamma_2\}$, the estimator \tilde{f} is usually biased.

Lemma 2. Suppose that $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$ and K satisfies K1–K2. Then

$$\mathbb{E}\{\tilde{f}(z)\} - f(z) = \sum_{k=1}^{\lceil h_n^{-1/2} \rceil} \frac{\rho_k^\eta(h_n)f(z+k) - \{\rho_k^\eta(h_n) + \rho_{-k}^\eta(h_n)\}f(z) + \rho_{-k}^\eta(h_n)f(z-k)}{k^2},$$

where

$$\rho_k^\eta(h_n) = k^2 \int_{A_k^\eta(h_n)} K(t)\eta(k - h_nt) dx$$

and

$$A_k^\eta(h_n) = [(1 - \gamma_2 - k)h_n^{-1}, (-1 + \gamma_2 - k)h_n^{-1}] \cap [-1, 1].$$

To interpret the bias, it is helpful to focus on a simple case first.

Corollary 1. Suppose that $\eta(x) = \mathbf{1}(|x| \leq 1/2)$ and K is a symmetric function satisfying $K1-K2$. Then for all $z \in \mathbb{Z}$,

$$\mathbb{E}\{\tilde{f}(z)\} = f(z) + \sum_{k=1}^{\lceil h_n - 1/2 \rceil} \rho_k(h_n) \Delta_k^2 f(z),$$

where

$$\Delta_k^2 f(z) = \frac{f(z+k) - 2f(z) + f(z-k)}{k^2},$$

and

$$\rho_k(h_n) = k^2 \int_{A_k(h_n)} K(t) dt$$

with

$$A_k(h_n) = [(1/2 - k)h_n^{-1}, (-1/2 - k)h_n^{-1}] \cap [-1, 1].$$

The operator Δ_k^2 is known as the *second-order central difference operator* (see, e.g., [13]). It is commonly used as numerical approximation of second-order derivative of real-valued functions, which is

$$\frac{d^2 f(x)}{dx^2} = \lim_{s \rightarrow 0} \frac{f(x+s) - 2f(x) + f(x-s)}{s^2}.$$

We can interpret $\Delta_k^2 f$ as a discrete analogue to the second order derivative of a real-valued function. In this aspect, the discrete setting is similar to the continuous one (where the bias of \tilde{f} is proportional to the second order derivative).

The parameter k is called the *step size* and determines how local the derivative approximation is. The bias of \tilde{f} is a weighted sum of such ‘derivatives’ for several values of k . The bandwidth h_n limits the maximal step size and thereby controls the locality of the bias. Although not universally true, smaller values of h_n typically correspond to a smaller bias. A simple counterexample is when $f(z+k) = f(z-k) = f(z)$ for all $k \leq \lceil h_n - 1/2 \rceil$, where the bias is zero for all $h'_n \leq h_n$. There are also situations where decreasing h_n leads to a larger bias. This phenomenon also exists in the continuous setting, but is disguised by asymptotic approximations. When $h_n \leq 1/2$ as in Theorem 1, the estimator is unbiased.

The bias in Lemma 2 can be interpreted similarly. But Δ_k^2 is replaced by a weighted approximation of the derivative. If η or K are asymmetric, different weights will be assigned to the ‘forward derivative’ $k^{-1}\{f(z+k) - f(z)\}$ and the ‘backward derivative’ $k^{-1}\{f(z-k) - f(z)\}$.

5. MINIMAX RATE OPTIMALITY

The *maximum risk* associated with a class of densities \mathcal{F} and a (semi-) distance d is defined as

$$\mathcal{R}_n(\hat{f}, \mathcal{F}, d) = \sup_{f \in \mathcal{F}} \mathbb{E}_f \{d^2(\hat{f}, f)\}. \quad (6)$$

We consider two semi-distances related to pointwise and uniform consistency of \hat{f} , respectively:

$$d_{(\mathbf{z}, \mathbf{x})}(\hat{f}, f) = |\hat{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})| \quad \text{for some } (\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^p \times \mathbb{R}^q,$$

$$d_{\infty, \mathcal{S}}(\hat{f}, f) = \sup_{\mathcal{S}} |\hat{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})| \quad \text{for some } \mathcal{S} \subset \mathbb{Z}^p \times \mathbb{R}^q.$$

For \mathcal{F} , we shall consider all bounded density functions whose continuous part belongs to a Hölder class. For $\mathbf{a} \in \mathbb{N}_0^q$, we use the multi-index notation $|\mathbf{a}| = \sum_{j=1}^q a_j$, $\mathbf{x}^{\mathbf{a}} = x_1^{a_1} \cdots x_q^{a_q}$, and denote the partial derivatives of f with respect to \mathbf{x} as

$$D_{\mathbf{x}}^{\mathbf{a}} f(\mathbf{z}, \mathbf{x}) = \frac{\partial^{|\mathbf{a}|} f(\mathbf{z}, \mathbf{x})}{\partial^{a_1} x_1 \cdots \partial^{a_q} x_q}. \quad (7)$$

Definition 2. For $\lambda < \infty$ and $\beta = r + \alpha$, $r \in \mathbb{N}_0$, $0 < \alpha \leq 1$, the class $\mathcal{H}(\beta, \lambda)$ is defined as all functions $f: \mathbb{Z}^p \times \mathbb{R}^q \rightarrow \mathbb{R}$ such that for all $\mathbf{a} \in \mathbb{N}_0$ with $|\mathbf{a}| \leq r$,

- (i) f is a probability density on $\mathbb{Z}^p \times \mathbb{R}^q$,
- (ii) $D_{\mathbf{x}}^{\mathbf{a}} f(\mathbf{z}, \mathbf{x})$ exists for all $(\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^p \times \mathbb{R}^q$ and

$$\sup_{\mathbf{z} \in \mathbb{Z}^p, \mathbf{x}, \mathbf{x}' \in \mathbb{R}^q} \left\{ \frac{|D_{\mathbf{x}}^{\mathbf{a}} f(\mathbf{z}, \mathbf{x}) - D_{\mathbf{x}}^{\mathbf{a}} f(\mathbf{z}, \mathbf{x}')|}{\|\mathbf{x} - \mathbf{x}'\|_2^\alpha} + f(\mathbf{z}, \mathbf{x}) \right\} \leq \lambda.$$

Remark 8. If $p \geq 1$ and $q = 0$, $\mathcal{H}(\beta, \lambda)$ contains all densities on \mathbb{Z}^p . If $p = 0$ and $q \geq 1$, it is a Hölder class on \mathbb{R}^q .

The following result establishes convergence rates of the jittering kernel density estimator with respect to the maximum risk.

Theorem 3. Denote by \tilde{f} the estimator defined by Eq. (2). Suppose $f \in \mathcal{H}(\beta, \lambda)$ and Assumptions A2–A4 of Theorem 1 hold with $\ell \geq r + 1$, $\beta = r + \alpha$, $0 < \alpha \leq 1$, $\lambda < \infty$. Assume further that there are $h_0 > 0$, $n_0 \in \mathbb{N}$ such that $h_n \in (h_0, \min\{\gamma_1, 1 - \gamma_2\}]$ for all $n \geq n_0$. Then there exists $\bar{c} > 0$ such that

$$\limsup_{n \rightarrow \infty} r_n^{-2} \mathcal{R}_n(\tilde{f}, \mathcal{F}, d) \leq \bar{c},$$

in each of the following cases:

- (i) $r_n = n^{-\beta/(2\beta+q)}$, $d = d_{(\mathbf{z}, \mathbf{x})}$,
- (ii) $r_n = (n/\log n)^{-\beta/(2\beta+q)}$, $d = d_{\infty, \mathcal{S}}$, $q \geq 1$,
- (iii) $r_n = n^{-1/2}$, $d = d_{\infty, \mathcal{S}}$, $q = 0$, $|\mathcal{S}| < \infty$,
- (iv) $r_n = (n/\log \log n)^{-1/2}$, $d = d_{\infty, \mathcal{S}}$, $q = 0$, $|\mathcal{S}| = \infty$,

for arbitrary $(\mathbf{x}, \mathbf{z}) \in \mathbb{Z}^p \times \mathbb{R}^q$ and $\mathcal{S} \subset \mathbb{Z}^p \times \mathbb{R}^q$.

We shall see that the rates in Theorem 3(i)–(iii) are optimal in a minimax sense. The minimax risk is defined as

$$\mathcal{R}_n^*(\mathcal{F}, d) = \inf_{\tilde{f}} \mathcal{R}_n(\tilde{f}, \mathcal{F}, d) = \inf_{\tilde{f}} \sup_{f \in \mathcal{F}} E_f \{d^2(\tilde{f}, f)\},$$

where the infimum is taken over all possible estimators \tilde{f} of f . In our context, an ‘estimator’ is any measurable function of $(\mathbf{Z}_i, \mathbf{X}_i)$, $i = 1, \dots, n$.

Definition 3. A sequence of positive real numbers r_n is called

- (i) an upper bound on the minimax rate if there is \bar{c} such that

$$\limsup_{n \rightarrow \infty} r_n^{-2} \mathcal{R}_n^*(\mathcal{F}, d) \leq \bar{c};$$

- (ii) a lower bound on the minimax rate if there is $\underline{c} > 0$ such that

$$\liminf_{n \rightarrow \infty} r_n^{-2} \mathcal{R}_n^*(\mathcal{F}, d) \geq \underline{c};$$

- (iii) a minimax-optimal rate of convergence if both (i) and (ii) hold.

In a purely continuous setting, optimal rates have long been established [9, 20, 21]. To the best of the author's knowledge, there are no results on optimal rates in the mixed data setting.

To show that a rate is minimax-optimal, we have to check that it is both an upper and lower bound on the minimax rate. Theorem 3 already gives us an upper bound, since, for any estimator \hat{f} ,

$$\mathcal{R}_n^*(\mathcal{F}, d) = \inf_{\hat{f}} \mathcal{R}_n(\hat{f}, \mathcal{F}, d) \leq \mathcal{R}_n(\hat{f}, \mathcal{F}, d).$$

Lower bounds on the minimax rate can be deduced easily by considering subsets of $\mathcal{H}(\beta, \lambda)$ for which lower bounds are known (see Section A.4).

Theorem 4. *Let $\mathcal{S} \subset \mathbb{Z}^p \times \mathbb{R}^q$ and $(\mathbf{z}, \mathbf{x}) \in \mathcal{S}$. The minimax-optimal rate of convergence r_n^* associated with the class $\mathcal{H}(\beta, \lambda)$ and distance d satisfies*

- (i) $r_n^* = n^{-\beta/(2\beta+q)}$ for $d = d_{(\mathbf{z}, \mathbf{x})}$,
- (ii) $r_n^* = (n/\log n)^{-\beta/(2\beta+q)}$ for $d = d_{\infty, \mathcal{S}}$, $q \geq 1$,
- (iii) $r_n^* = n^{-1/2}$ for $d = d_{\infty, \mathcal{S}}$, $q = 0$, $|\mathcal{S}| < \infty$,
- (iv) $r_n^* \in [n^{-1/2}, (n/\log \log n)^{-1/2}]$ for $d = d_{\infty, \mathcal{S}}$, $q = 0$, $|\mathcal{S}| = \infty$.

Remark 9. Theorems 3 and 4 imply that the jittering kernel density estimator converges at minimax-optimal rates for cases (i)–(iii).

Remark 10. Theorem 4 only provides an interval for the optimal rate in case (iv). Minimax analysis for this setting is surprisingly hard; see [7] for minimax rates with respect to the ℓ_1 distance. The interval is quite narrow, differing only by a factor of size $\log \log n$. The exact rate, however, remains an open problem.

6. SIMULATION EXPERIMENTS

The jittering kernel density estimator has appealing asymptotic properties. This may come as a surprise: since we are adding noise to the data, we could expect that the data become less informative and uncertainty increases. We complement our asymptotic arguments with a small numerical experiment that illustrates the small sample performance of the estimator. Because of its wide use and close resemblance to our approach, we will use the estimator of [11] as a benchmark.

We use the following setup:

- We compare three estimators
 1. **jkde**: the jittering kernel density estimator with noise density $\eta(x) = \mathbf{1}(|x| < 1/2)$, for which $\gamma_1 = \gamma_2 = 1/2$.
 2. **jkde2**: the jittering kernel density estimator with noise density $\eta(x) = f_{U_{1/4,5}}(x)$ (as in Example 3 in [14]), for which $\gamma_1 = 3/8$, $\gamma_2 = 5/8$.
 3. **liracine**: the estimator of [11] as implemented in the **np** package [8].

Contrary to (2), we use one bandwidth parameter for each variable. Both estimators use likelihood cross-validation for bandwidth selection.

- We estimate the density f of a vector $(\mathbf{Z}, \mathbf{X}) \in \mathbb{Z}^p \times \mathbb{R}^q$, where $Z_j \sim \text{Binomial}(m, 0.3)$ for all $j = 1, \dots, p$, $X_j \sim \mathcal{N}(0, 1)$ for all $j = 1, \dots, q$. For the sake of simplicity, all variables are simulated independently.
- Results are based on $N_{\text{sim}} = 1000$ simulated data sets with sample sizes $n = 50, 200$.

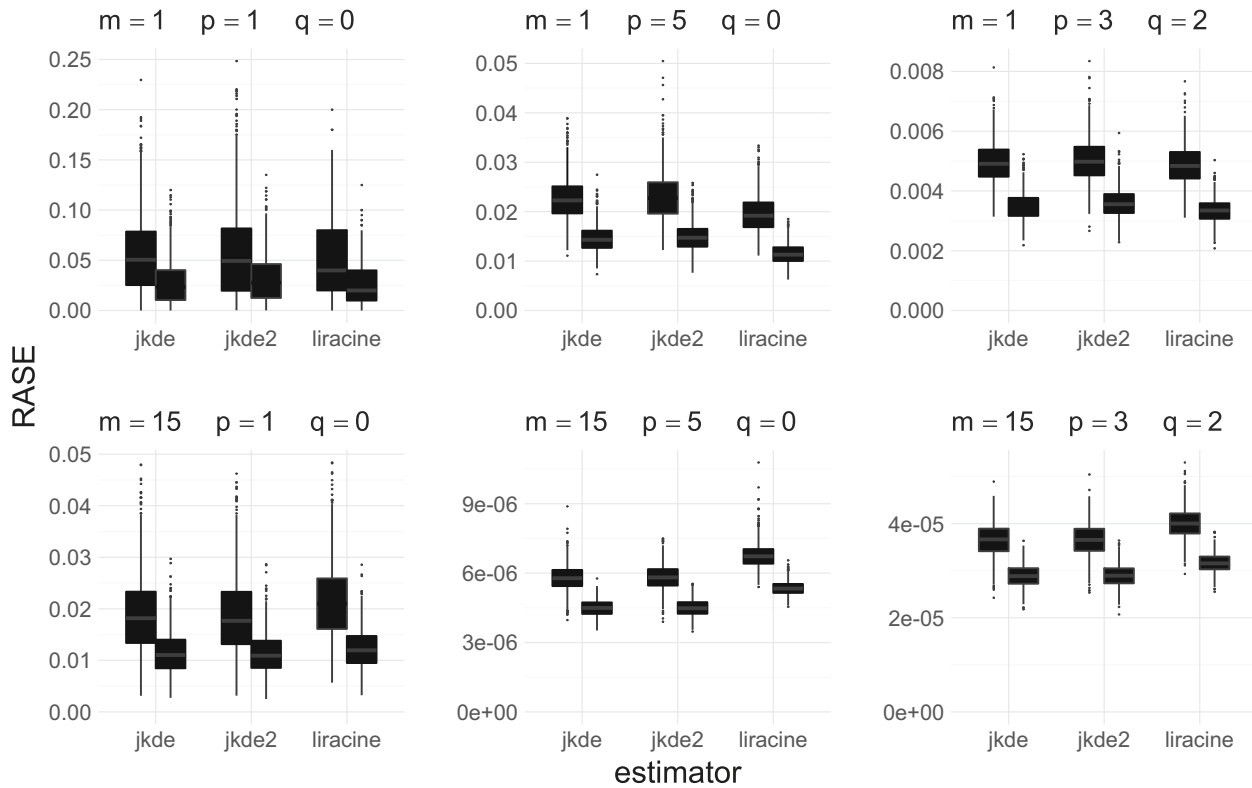


Fig. 1. RASE achieved by the two estimators for various choices of p , q , and m . Each estimator is represented by two boxes; the left box corresponds to $n = 50$, the right to $n = 200$.

- As a performance measure we use the *root average square error (RASE)* computed over a grid in $\mathbb{Z}^p \times \mathbb{R}^q$. More specifically, we use $\mathcal{Z} = \{0, \dots, m\}$, $\mathcal{X} = \{-2, -1.6, \dots, 2\}$, and

$$\text{RASE}(\hat{f}, f) = \sqrt{\sum_{z_1 \in \mathcal{Z}} \cdots \sum_{z_p \in \mathcal{Z}} \sum_{x_1 \in \mathcal{X}} \cdots \sum_{x_q \in \mathcal{X}} \{\hat{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})\}^2}.$$

Figure 1 shows the estimators' performance for various values of p , q and m . Each estimator is represented by two boxes, where the left box corresponds to $n = 50$ and the right box to $n = 200$. The choice of noise density seems to be of minor importance: `jkde` and `jkde2` give almost identical results. Compared to `liracine`, the two estimators show only subtle differences. The two jittering estimators are more accurate in all scenarios with $m = 15$, and less accurate when $m = 1$. This is related to our observation from Section 4.1 that the efficiency is worse when $f(z)$ is large. The relative performance of the three estimators is consistent across the two sample sizes under consideration. Overall, the jittering estimators are competitive with the benchmark estimator `liracine`. We found no evidence that adding artificial noise negatively affects the accuracy of the estimates. This confirms what was suggested by the estimator's asymptotic properties.

A. PROOFS

A.1. Proof of Theorem 1

We first calculate the bias term. Using a change of variables, we get

$$\begin{aligned} \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= \frac{1}{h_n^p b_n^q} \mathbb{E}\left\{K\left(\frac{\mathbf{Z} + \mathbf{E} - \mathbf{z}}{h_n}\right)K\left(\frac{\mathbf{X} - \mathbf{x}}{b_n}\right)\right\} \\ &= \frac{1}{h_n^p b_n^q} \int_{\mathbb{R}^{p+q}} K\left(\frac{\mathbf{s} - \mathbf{z}}{h_n}\right)K\left(\frac{\mathbf{t} - \mathbf{x}}{b_n}\right)f_\eta(\mathbf{s}, \mathbf{t}) \, ds \, dt \end{aligned}$$

$$= \int_{\mathbb{R}^{p+q}} K(\mathbf{u})K(\mathbf{v})f_\eta(\mathbf{z} + h_n\mathbf{u}, \mathbf{x} + b_n\mathbf{v}) d\mathbf{u} d\mathbf{v}.$$

Since $\eta \in \mathcal{E}_{\gamma_1, \gamma_2}$, it holds for all $(\mathbf{z}, \mathbf{x}) \in \mathbb{Z}^q \times \mathbb{R}^q$ and $0 \leq \epsilon \leq \min\{\gamma_1, 1 - \gamma_2\}$ that $f_\eta(\mathbf{z} + \epsilon, \mathbf{x}) = f(\mathbf{z}, \mathbf{x})$. Furthermore, K is zero outside $[-1, 1]$. Hence, for $h_n \leq \min\{\gamma_1, 1 - \gamma_2\}$,

$$\begin{aligned} \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= \int_{\mathbb{R}^{p+q}} K(\mathbf{u})K(\mathbf{v})f(\mathbf{z}, \mathbf{x} + b_n\mathbf{v}) d\mathbf{u} d\mathbf{v} \\ &= \int_{[-1, 1]^q} K(\mathbf{v})f(\mathbf{z}, \mathbf{x} + b_n\mathbf{v}) d\mathbf{v}. \end{aligned} \quad (8)$$

Recall the derivative notation in (7). An ℓ th order Taylor expansion of f yields that

$$\begin{aligned} \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x}) &= \sum_{1 \leq |\mathbf{a}| \leq \ell} \frac{b_n^{|\mathbf{a}|}}{|\mathbf{a}|!} \int_{[-1, 1]^q} K(\mathbf{v})\mathbf{v}^{\mathbf{a}} D_{\mathbf{x}}^{|\mathbf{a}|} f(\mathbf{z}, \mathbf{x}) d\mathbf{v} \\ &\quad + \sum_{|\mathbf{a}| = \ell + 1} \frac{b_n^{\ell + 1}}{(\ell + 1)!} \int_{[-1, 1]^q} K(\mathbf{v})\mathbf{v}^{\mathbf{a}} D_{\mathbf{x}}^{|\mathbf{a}|} f(\mathbf{z}, \mathbf{x} + \tau_{\mathbf{a}}\mathbf{v}) d\mathbf{v} \\ &= \frac{b_n^\ell}{\ell!} \sum_{j=1}^q \int_{[-1, 1]} K(v_j) v_j^\ell \frac{\partial^\ell f(\mathbf{z}, \mathbf{x})}{\partial x_j^\ell} dv_j, \\ &\quad + \sum_{|\mathbf{a}| = \ell + 1} \frac{b_n^{\ell + 1}}{(\ell + 1)!} \int_{[-1, 1]^q} K(\mathbf{v})\mathbf{v}^{\mathbf{a}} D_{\mathbf{x}}^{|\mathbf{a}|} f(\mathbf{z}, \mathbf{x} + \tau_{\mathbf{a}}\mathbf{v}) d\mathbf{v} \end{aligned}$$

for some $\tau_{\mathbf{a}} \in [0, 1]$, where the second equality is due to K2. The second sum is $o(b_n^\ell)$ because all terms are bounded by A1 and K1. In summary,

$$\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x}) = \frac{b_n^\ell \sigma_\ell}{\ell!} \sum_{j=1}^q \frac{\partial^\ell f(\mathbf{z}, \mathbf{x})}{\partial x_j^\ell} + o(b_n^\ell),$$

as claimed.

For the variance, we get

$$\begin{aligned} \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= \frac{1}{nh_n^{2p} b_n^{2q}} \text{Var}\left\{K\left(\frac{\mathbf{Z} + \mathbf{E} - \mathbf{z}}{h_n}\right) K\left(\frac{\mathbf{X} - \mathbf{x}}{b_n}\right)\right\} \\ &= \frac{1}{n} \left[\frac{1}{h_n^{2p} b_n^{2q}} \mathbb{E}\left\{K\left(\frac{\mathbf{Z} + \mathbf{E} - \mathbf{z}}{h_n}\right)^2 K\left(\frac{\mathbf{X} - \mathbf{x}}{b_n}\right)^2\right\} \right. \\ &\quad \left. - \frac{1}{h_n^{2p} b_n^{2q}} \mathbb{E}\left\{K\left(\frac{\mathbf{Z} + \mathbf{E} - \mathbf{z}}{h_n}\right) K\left(\frac{\mathbf{X} - \mathbf{x}}{b_n}\right)\right\}^2 \right]. \end{aligned}$$

The second term in square brackets has already been calculated for the bias. Using similar arguments, we can show

$$\begin{aligned} &\frac{1}{nh_n^p b_n^q} \int_{\mathbb{R}^{p+q}} K^2\left(\frac{\mathbf{s} - \mathbf{z}}{h_n}\right) K^2\left(\frac{\mathbf{t} - \mathbf{x}}{b_n}\right) f_\eta(\mathbf{s}, \mathbf{t}) ds dt \\ &= \kappa^p \int_{[-1, 1]^q} K^2(\mathbf{v}) f(\mathbf{z}, \mathbf{x} + b_n\mathbf{v}) d\mathbf{v} = \kappa^{p+q} f(\mathbf{z}, \mathbf{x}) + o(1). \end{aligned}$$

Together,

$$\begin{aligned} \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} &= \frac{\kappa^{p+q}}{nh_n^p b_n^q} f(\mathbf{z}, \mathbf{x}) + \frac{f^2(\mathbf{z}, \mathbf{x})}{n} + o\left(\frac{1}{nh_n^q b_n^q}\right) \\ &= \frac{f(\mathbf{z}, \mathbf{x})}{nb_n^q} \{h_n^{-p} \kappa^{p+q} - b_n^q f(\mathbf{z}, \mathbf{x})\} + o\left(\frac{1}{nh_n^p b_n^q}\right). \end{aligned}$$

To show that the estimator is asymptotically normal, define

$$Y_{i,n} = \frac{1}{nb_n^q} K\left(\frac{\mathbf{Z}_i + \mathbf{E}_i - \mathbf{z}}{h_n}\right) K\left(\frac{\mathbf{X}_i - \mathbf{x}}{b_n}\right).$$

Then $\tilde{f}(\mathbf{z}, \mathbf{x}) = \sum_{i=1}^n Y_{i,n}$, which is asymptotically normal if the Lyapunov condition,

$$\left\{ \sum_{i=1}^n \mathbb{E}(|Y_{i,n}|^3) \right\}^{1/3} \left\{ \sum_{i=1}^n \text{Var}(Y_{i,n}) \right\}^{-1/2} \rightarrow 0,$$

is fulfilled. With arguments similar to the derivation of $\text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}$, we get

$$\mathbb{E}(|Y_{i,n}|^3) = O(n^{-1}h_n^{-2p}b_n^{-2q}) \quad \text{and} \quad \text{Var}(Y_{i,n}) = O(h_n^{-p}b_n^{-q}).$$

Thus,

$$\left\{ \sum_{i=1}^n \mathbb{E}(|Y_{i,n}|^3) \right\}^{1/3} \left\{ \sum_{i=1}^n \text{Var}(Y_{i,n}) \right\}^{-1/2} = O\{(nh_n^p b_n^q)^{-1/6}\},$$

which is $o(1)$ due to Assumption A4.

A.2. Proof of Theorem 2

From the triangle inequality, we get the bound

$$|\tilde{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})| \leq |\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x})| + |\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}|. \quad (9)$$

We start as in the proof of Theorem 1, but expand (8) as a Taylor polynomial of order $\ell - 2$. We can then show that for some $\tau \in [0, 1]$,

$$\begin{aligned} \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x}) &= \frac{b_n^{\ell-1}}{(\ell-1)!} \sum_{j=1}^q \int_{[-1,1]} K(v_j) v_j^{\ell-1} \frac{\partial^{\ell-1} f(\mathbf{z}, \mathbf{x} + \tau b_n \mathbf{v})}{\partial x_j^{\ell-1}} dv_j \\ &= \frac{b_n^{\ell-1}}{(\ell-1)!} \sum_{j=1}^q \int_{[-1,1]} K(v_j) v_j^{\ell-1} \left\{ \frac{\partial^{\ell-1} f(\mathbf{z}, \mathbf{x} + \tau b_n \mathbf{v})}{\partial x_j^{\ell-1}} - \frac{\partial^{\ell-1} f(\mathbf{z}, \mathbf{x})}{\partial x_j^{\ell-1}} \right\} dv_j, \end{aligned}$$

where the second equality holds because of K2. Using A1', we get

$$\sup_{(\mathbf{z}, \mathbf{x}) \in \mathcal{S}} |\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x})| \leq \frac{b_n^\ell L \tau}{(\ell-1)!} \sum_{j=1}^q \int_{[-1,1]} |K(v_j)| |v_j|^\ell dv_j = O(b_n^\ell) \quad (10)$$

for a positive constant $L < \infty$. Furthermore,

$$\mathbb{E}\{|\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}|^2\} = \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} = O\{(nh_n^p b_n^q)^{-1}\},$$

as in Theorem 1. And since convergence in L^2 implies convergence in probability,

$$|\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}| = O_p\{(nh_n^p b_n^q)^{-1/2}\},$$

which, together with (10), proves (4).

Moreover, there is a positive constant $\bar{c}_1 < \infty$ such that almost surely

$$\lim_{n \rightarrow \infty} \sqrt{\frac{nh_n^p b_n^q}{\max\{\log \log n, \log h_n^{-1}, \log b_n^{-1}\}}} \sup_{\mathcal{S}} |\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}| \leq \bar{c}_1, \quad (11)$$

see Theorem 1 of [3]. Combining (10) and (11) proves (5).

A.3. Proof of Theorem 3

Note that we can write

$$\mathbb{E}\{d^2(\tilde{f}, f)\} = \mathbb{E}\left\{\sup_{\mathcal{S}'} |\tilde{f}(\mathbf{z}, \mathbf{x}) - f(\mathbf{z}, \mathbf{x})|^2\right\},$$

where $\mathcal{S}' = \{(\mathbf{z}, \mathbf{x})\}$ for $d_{(\mathbf{z}, \mathbf{x})}$ and $\mathcal{S}' = \mathcal{S}$ for $d_{\infty, \mathcal{S}}$. It holds

$$\frac{1}{2}\mathbb{E}\{d^2(\tilde{f}, f)\} \leq \sup_{\mathcal{S}'} |\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x})|^2 + \mathbb{E}\left[\sup_{\mathcal{S}'} |\tilde{f}(\mathbf{z}, \mathbf{x}) - \mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\}|^2\right] = a_1 + a_2. \quad (12)$$

Using arguments almost identical to (10), we obtain

$$\sup_{(\mathbf{z}, \mathbf{x}) \in \mathcal{S}} |\mathbb{E}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} - f(\mathbf{z}, \mathbf{x})| \leq \frac{b_n^\beta \lambda_T^{\beta-r}}{r!} \sum_{j=1}^q \int_{[-1,1]} |K(v_j)| |v_j|^\beta dv_j = b_n^\beta \bar{c}_2.$$

For bounding a_2 , we need to consider the characteristics of scenarios (i)–(iv).

(i) We proceed as in the proof of Theorem 1 to get

$$a_2 = \text{Var}\{\tilde{f}(\mathbf{z}, \mathbf{x})\} = \frac{\kappa^{p+q}}{nh_n^p b_n^q} f(\mathbf{z}, \mathbf{x}) + \frac{f^2(\mathbf{z}, \mathbf{x})}{n} + o\left(\frac{1}{nb_n^q}\right).$$

For $q \geq 1$, choosing $b_n \sim n^{-1/(2\beta+q)}$ yields

$$\limsup_{n \rightarrow \infty} n^{2\beta/(2\beta+q)} a_2 \leq \frac{\kappa^{p+q}}{h_0^p} f(\mathbf{z}, \mathbf{x}) = \bar{c}_3 < \infty.$$

If $q = 0$, it holds $f \leq 1$, and we get

$$\limsup_{n \rightarrow \infty} na_2 \leq \frac{\kappa^{p+q}}{h_0^p} f(\mathbf{z}) + f^2(\mathbf{z}) \leq \frac{\kappa^{p+q}}{h_0^p} + 1 = \bar{c}_4 < \infty. \quad (13)$$

(ii) With $h_n \sim 1$ and $b_n \sim (n/\log n)^{-1/(2\beta+q)}$ in (11), we get

$$\limsup_{n \rightarrow \infty} (n/\log n)^{2\beta/(2\beta+q)} a_2 \leq \bar{c}_1.$$

(iii) Using (13) yields

$$\limsup_{n \rightarrow \infty} na_2 \leq \sum_{\mathbf{z} \in \mathcal{S}'} \limsup_{n \rightarrow \infty} \mathbb{E}[|\tilde{f}_{\mathbf{z}}(\mathbf{z}) - \mathbb{E}\{\tilde{f}_{\mathbf{z}}(\mathbf{z})\}|^2] \leq |\mathcal{S}'| \bar{c}_4 = \bar{c}_5 < \infty.$$

(iv) With $h_n \sim 1$ and $b_n = 1$ in (11), we get

$$\limsup_{n \rightarrow \infty} (n/\log \log n) a_2 \leq \bar{c}_1.$$

Setting $\bar{c} = 2(\bar{c}_1 + \bar{c}_2 + \bar{c}_3 + \bar{c}_4 + \bar{c}_5)$ concludes the proof.

A.4. Proof of Theorem 4

We start with lower bounds for (i) and (iii). Fix $\mathbf{z} \in \mathbb{Z}^p$ and define

$$\mathcal{G}_1(\beta, \lambda) = \{f \in \mathcal{H}(\beta, \lambda) : f(\mathbf{z}', \mathbf{x}) = 0 \text{ for } \mathbf{z}' \neq \mathbf{z}\}.$$

This set contains all probability densities in $\mathcal{H}(\beta, \lambda)$ that correspond to a random vector (\mathbf{Z}, \mathbf{X}) with $\mathbf{Z} = \mathbf{z}$ almost surely. This is equivalent to the case where all variables are continuous. By definition, $\mathcal{G}_1(\beta, \lambda) \subset \mathcal{H}(\beta, \lambda)$ and, thus, $\mathcal{R}_n^*\{\mathcal{G}_1(\beta, \lambda), d\} \leq \mathcal{R}_n^*\{\mathcal{H}(\beta, \lambda), d\}$. The two rates in Theorem 4(i) and (iii) then follow from Theorem 9 in [9].

For (ii) and (iv), we can simply consider a parametric family of densities \mathcal{G}_2 . This yields the classical lower bound $n^{-1/2}$ for estimating a finite-dimensional parameter (see, e.g., Chapter 2 in [22]).

7. CONCLUSION

This article gives an in-depth analysis of the behavior of the jittering estimator. It was shown to have appealing large-sample properties and perform well on small samples.

Although our focus was on a particular instance of the class of jittering estimators, we also learned something about the class as a whole. Adding noise to discrete variables does not have a negative impact on estimation accuracy. This is true for both large samples (as confirmed by our asymptotic analysis) and small samples (as illustrated by simulations). More specifically, it allows for estimators that are optimal in terms of convergence rates and efficiency. It is likely that these findings generalize to more sophisticated density estimators or estimators of functionals of the density, such as regression functions.

SUPPLEMENTARY MATERIAL

- <https://github.com/tnagler/cctools>: an R package implementing the jittering kernel density estimator and likelihood cross-validation for the bandwidths.
- <https://gist.github.com/tnagler/786465cee2c774a844ff1846e7cdacd8>: code for the simulation study in Section 6.

ACKNOWLEDGEMENTS

This work was partially supported by the German Research Foundation (DFG grant CZ 86/5-1).

REFERENCES

1. I. A. Ahmad and P. B. Cerrito, “Nonparametric Estimation of Joint Discrete-Continuous Probability Densities with Applications”, *J. Statist. Planning and Inference* **41** (3), 349–364 (1994).
2. S. Efromovich, “Nonparametric Estimation of the Anisotropic Probability Density of Mixed Variables”, *J. Multivar. Anal.* **102** (3), 468–481 (2011).
3. U. Einmahl and D. M. Mason, “Uniform in Bandwidth Consistency of Kernel-Type Function Estimators,” *Ann. Statist.* **33** (3), 1380–1403, 06 (2005).
4. C. Genest, J. G. Nešlehová, and B. Rémillard, “Asymptotic Behavior of the Empirical Multilinear Copula Process under Broad Conditions”, *J. Multivar. Anal.* (2017).
5. P. Hall, J. Racine, and Q. Li, “Cross-Validation and the Estimation of Conditional Probability Densities”, *J. Amer. Statist. Assoc.* **99** (468), 1015–1026 (2004).
6. P. Hall et al., “Orthogonal Series Methods for Both Qualitative and Quantitative Data”, *Ann. Statist.* **11** (3), 1004–1007 (1983).
7. Y. Han, J. Jiao, and T. Weissman, “Minimax Estimation of Discrete Distributions under Loss”, *IEEE Trans. on Information Theory* **61** (11), 6343–6354 (2015).
8. T. Hayfield and J. S. Racine, “Nonparametric Econometrics: The `np` Package”, *J. Statist. Software* **27** (5), (2008).
9. I. Ibragimov and R. Khas'minskii, “Estimation of Distribution Density”, *J. Soviet Math.* **21** (1), 40–57 (1983).
10. C. G. G. Aitken and J. Aitchison, “Multivariate Binary Discrimination by the Kernel Method”, *Biometrika* **63** (3), 413–420 (1976).
11. Q. Li and J. Racine, “Nonparametric Estimation of Distributions with Categorical and Continuous Data”, *J. Multivar. Anal.* **86** (2), 266–292 (2003).
12. J. S. Marron, “Visual Understanding of Higher-Order Kernels”, *J. Comput. and Graph. Statist.* **3** (4), 447–458 (1994).
13. J. F. Monahan, *Numerical Methods of Statistics*. (Cambridge University Press, Cambridge, 2011).
14. T. Nagler, “A Generic Approach to Nonparametric Function Estimation with Mixed Data”, *Statist. Probab. Lett.*, DOI 10.1016/j.spl.2018.02.040 (in press).
15. E. Parzen, “On Estimation of a Probability Density Function and Mode”, *Ann. Math. Statist.* **33** (3), 1065–1076, 09 (1962).
16. J. Racine and Q. Li, “Nonparametric Estimation of Regression Functions with Both Categorical and Continuous Data”, *J. Econometrics* **119** (1), 99–130 (2004).
17. M. Rosenblatt, “Remarks on Some Nonparametric Estimates of a Density Function”, *Ann. Math. Statist.* **27** (3), 832–837, 09 (1956).
18. D. W. Scott, “The Curse of Dimensionality and Dimension Reduction”, in *Multivariate Density Estimation: Theory, Practice, and Visualization* (Wiley, New York, 2008) pp. 195–217.

19. L. Simar, V. Zelenyuk, et al., *To Smooth or not to Smooth? The Case of Discrete Variables in Nonparametric Regressions*, in *Working Paper Series* (Center for Efficiency and Productivity Analysis, 2011).
20. C. J. Stone, “Optimal Rates of Convergence for Nonparametric Estimators”, *Ann. Statist.* **8** (6), 1348–1360, 11 (1980).
21. C. J. Stone, “Optimal Uniform Rate of Convergence for Nonparametric Estimators of a Density Function or Its Derivatives”, *Recent Advances in Statist.* **5** (1983).
22. A. Tsybakov, *Introduction to Nonparametric Estimation*, in *Springer Series in Statist.* (Springer, New York, 2008).
23. M. Wand, “Error Analysis for General Multivariate Kernel Estimators”, *J. Nonparam. Statist.* **2** (1), 1–15 (1992).