

Density estimation on an unknown submanifold

Clément Berenfeld and Marc Hoffmann

*Université Paris-Dauphine & PSL, CNRS
CEREMADE, 75016 Paris, France*

e-mail: berenfeld@ceremade.dauphine.fr; hoffmann@ceremade.dauphine.fr

Abstract: We investigate the estimation of a density f from a n -sample on an Euclidean space \mathbb{R}^D , when the data are supported by an unknown submanifold M of possibly unknown dimension $d < D$, under a *reach* condition. We investigate several nonparametric kernel methods, with data-driven bandwidths that incorporate some learning of the geometry via a local dimension estimator. When f has Hölder smoothness β and M has regularity α , our estimator achieves the rate $n^{-\alpha \wedge \beta / (2\alpha \wedge \beta + d)}$ for a pointwise loss. The rate does not depend on the ambient dimension D and we establish that our procedure is asymptotically minimax for $\alpha \geq \beta$. Following Lepski's principle, a bandwidth selection rule is shown to achieve smoothness adaptation. We also investigate the case $\alpha \leq \beta$: by estimating in some sense the underlying geometry of M , we establish in dimension $d = 1$ that the minimax rate is $n^{-\beta / (2\beta + 1)}$ proving in particular that it does not depend on the regularity of M . Finally, a numerical implementation is conducted on some case studies in order to confirm the practical feasibility of our estimators.

MSC2020 subject classifications: 62C20, 62G05, 62G07.

Keywords and phrases: Point clouds, manifold reconstruction, nonparametric estimation, adaptive density estimation, kernel methods, Lepski's method.

Received May 2020.

1. Introduction

1.1. Motivation

Suppose we observe an n -sample (X_1, \dots, X_n) of size n distributed on an Euclidean space \mathbb{R}^D according to some density function f . We wish to recover f at some arbitrary point $x_0 \in \mathbb{R}^D$ nonparametrically. If the smoothness of f at x_0 measured in a strong sense is of order β – for instance by a Hölder condition or with a prescribed number of derivatives – then the optimal (minimax) rate for recovering $f(x_0)$ is of order $n^{-\beta / (2\beta + D)}$ and is achieved by kernel or projection methods, see *e.g.* the classical textbooks by Silverman [57], Devroye and Györfi [16] or Tsybakov [62, Sec. 1.2-1.3]. Extension to data-driven bandwidths (Bowman [12], Chiu [14]) offers the possibility to adapt to unknown smoothness, see Goldenshluger and Lepski [28, 29, 30] for a modern mathematical formulation. More generally, recommended reference on adaptive estimation is the textbook

by Giné and Nickl [27]. In many situations however, the dimension D of the ambient space is large, hitherto disqualifying such methods for practical applications. Opposite to the curse of dimensionality, a broad guiding principle in practice is that the observations (X_1, \dots, X_n) actually live on smaller dimensional structures and that the effective dimension of the problem is smaller if one can take advantage of the geometry of the data (Fefferman et al. [23]). This classical paradigm probably goes back to a conjecture of Stone [59] that paved the way to the study of the celebrated single-index model in nonparametric regression, where a structural assumption is put in the form $f(x) = g(\langle \vartheta, x \rangle)$, where $\langle \cdot, \cdot \rangle$ is the scalar product on \mathbb{R}^D , for some unknown univariate function $g: \mathbb{R} \rightarrow \mathbb{R}$ and direction $\vartheta \in \mathbb{R}^D$. Under appropriate assumptions, the minimax rate of convergence for recovering $f(x)$ with smoothness β drops to $n^{-\beta/(2\beta+1)}$ and does not depend on the ambient dimension D , see *e.g.* Gaïffas and Lecué [24], Lepski and Serdyukova [45] and the references therein. Also, in the search for significant variables, one postulates that f only depends on $d < D$ coordinates, leading to the structural assumption $f(x_1, \dots, x_D) = F(x_{i_1}, \dots, x_{i_d})$ for some unknown function $F: \mathbb{R}^d \rightarrow \mathbb{R}$ and $\{i_1, \dots, i_d\} \subset \{1, \dots, D\}$. In an analogous setting, the minimax rate of convergence becomes $n^{-\beta/(2\beta+d)}$ and this is also of a smaller order of magnitude than $n^{-\beta/(2\beta+D)}$, see Hoffmann and Lepski [36] in the white noise model.

The next logical step is to assume that the data (X_1, \dots, X_n) live on a d -dimensional submanifold M of the ambient space \mathbb{R}^D . When the manifold is known prior to the experiment, nonparametric density estimation dates back to Devroye and Györfi [16] when M is the circle, and on a homogeneous Riemannian manifold by Hendriks [34], see also Pelletier [54]. Several results are known for specific geometric structures like the sphere or the torus involved in many applied situations: inverse problems for cosmological data (Kim and Koo [41], Kim et al. [42], Kerkycharian et al. [39], in geology Hall et al. [32] or flow calculation in fluid mechanics Eugeciouglu and Srinivasan [18]). For genuine compact homogeneous Riemannian manifolds, a general setting for smoothness adaptive density estimation and inference has recently been considered by Kerkycharian et al. [38], or even in more abstract metric spaces in Cleanthous et al. [15]. See also Baldi et al. [7], Castillo et al. [13] and the references therein. A common strategy adapts conventional nonparametric tools like projection or kernel methods to the underlying geometry, via the spectral analysis of the Beltrami-Laplace operator on M . Under appropriate assumptions, this leads to exact or approximate eigenbases (spherical harmonics for the sphere, needlets and so on) or properly modified kernel methods, according to the Riemannian metric on M .

If the submanifold M itself is unknown, getting closer in spirit to a dimension reduction approach, the situation becomes drastically different: M hence its geometry is unknown, and considered as a nuisance parameter. In order to recover the density f at a given point $x_0 \in \mathbb{R}^D$ of the ambient space, one has to understand the minimal geometry of M that must be learned from the data and how this geometry affects the optimal reconstruction of f . This is the topic of the paper.

We consider in the paper a seemingly unusual framework where the support of a distribution is unknown while the aim is to recover the density at a point $x_0 \in \mathbb{R}^D$ which is known to be on the support. As mentioned above, this actually covers at least two situations:

- The data are high-dimensional and it is reasonable to believe that they actually lie on a smaller dimensional subset of the ambient space \mathbb{R}^D , which can be assumed to be a submanifold. In that case, x_0 can be seen as an extraneous observation X from the density f (extracted for instance from the point cloud), and the analysis can be implicitly performed conditional on $X = x_0$;
- The data naturally lie on a submanifold, like a spheroid for geological application, or a cell membrane in microbiology (see for instance Klein et al. [43] who describe a technique that yields such a point cloud). In this case, x_0 can be seen as an observation X like above, but there is also the situation where the statistician can know whether or not a given point x_0 is within the support (for instance a point on a cell membrane, or a geographical location on the Earth surface) without knowing the geometric features of the latter and without needing to estimate them.

1.2. Main results

We construct a class of compact smooth submanifolds of dimension d of the Euclidean space \mathbb{R}^D , without boundaries, that constitute generic models for the unknown support of the target density f that we wish to reconstruct. We further need a *reach condition*, a somehow unavoidable notion in manifold reconstruction that goes back to Federer [21]: it is a geometric invariant that quantifies both local curvature conditions and how tightly the submanifold folds on itself. It is related to the scale at which the sampling rate n can effectively recover the geometry of the submanifold, see Section 2.3 below. We consider regular manifolds M with reach bounded below that satisfy the following property: M admits a local parametrisation at every point $x \in M$ by its tangent space $T_x M$, and this parametrisation is sufficiently regular. A natural candidate is given by the exponential map $\exp_x : T_x M \rightarrow M \subset \mathbb{R}^D$. More specifically, for some regularity parameter $\alpha \geq 0$, we require a certain uniform bound for the $(\alpha + 1)$ -fold differential of the exponential map to hold, quantifying in some sense the regularity of the parametrisation in a minimax spirit, see Definition 2.2 below. Our approach is close to that of Aamari and Levrard [3, Def. 1] that consider arbitrary parametrisations among those close to the inverse of the projection onto tangent spaces. Given a density function $f : M \rightarrow [0, \infty)$ with respect to the volume measure on M , we have a natural extension of smoothness spaces on M by requiring that $f \circ \exp_x : T_x M \rightarrow \mathbb{R}$ is a smooth map in any reasonable sense, see Section 2.2 below. This is line for instance with Triebel [61] for the characterisation of function spaces on a Riemannian manifold.

Our main result is that in order to reconstruct $f(x_0)$ efficiently at a point $x_0 \in \mathbb{R}^D$ when f has smoothness β and lives on an unknown submanifold

of smoothness α and unknown dimension $d < D$, it is sufficient to consider estimators of the form

$$\widehat{f}_h(x_0) = \frac{1}{nh^{\widehat{d}(x_0)}} \sum_{i=1}^n K\left(\frac{x_0 - X_i}{h}\right), \quad x \in \mathbb{R}^D, \quad (1)$$

where $K : \mathbb{R}^D \rightarrow \mathbb{R}$ is a certain kernel and $\widehat{d}(x_0) = \widehat{d}(x_0, X_1, \dots, X_n)$ is an estimator of the local dimension of the support of f in the vicinity of x based on a scaling estimator as introduced in Farahmand et al. [20]. We prove in Theorem 3.1 that following a classical bias-variance trade-off for the bandwidth h , the rate $n^{-\alpha \wedge \beta / (2\alpha \wedge \beta + d)}$ is achievable for pointwise and global loss when the dimension of M is d , irrespectively of the ambient dimension D . In particular, it is noteworthy that in terms of manifold learning, only the dimension of M needs to be estimated. When $\alpha \geq \beta$, we also have a lower bound (Theorem 3.2) showing that our result is asymptotically minimax optimal. Moreover, by implementing Lepski's principle (Lepskii [47]), we are able to construct a data driven bandwidth $\widehat{h} = \widehat{h}(x_0, X_1, \dots, X_n)$ that achieves in Theorem 3.4 the rate $n^{-\alpha \wedge \beta / (2\alpha \wedge \beta + d)}$ up to a logarithmic term — unavoidable in the case of pointwise loss due to the Lepski-Low phenomenon (Lepskii [48], Low [49]). When the dimension d is known, the estimator (1) has already been investigated in squared-error norm in Ozakin and Gray [52] for a fixed manifold M and smoothness $\beta = 2$.

A remaining issue at this stage is to understand how the regularity of M can affect the minimax rates of convergence for smooth functions, *i.e.* when $\alpha \leq \beta$. We only have a partial answer to that question, when we restrict our attention to the one-dimensional case $d = 1$. When M is known, Pelletier [54] studied estimators of the form

$$\frac{1}{nh^d} \sum_{i=1}^n \frac{1}{\vartheta_{x_0}(X_i)} K\left(\frac{d_M(x_0, X_i)}{h}\right), \quad (2)$$

where $K : \mathbb{R} \rightarrow \mathbb{R}$ is a radial kernel, d_M is the intrinsic Riemannian distance on M and the correction term $\vartheta_{x_0}(X_i)$ is the volume density function on M (Besse [10, p. 154]) that accounts for the value of the density of the volume measure at X_i in normal coordinates around x_0 , taking into account how the submanifold curves around X_i . By establishing in Lemma 3.9 that ϑ_x is constant (and identically equal to one) when $d = 1$, we have another estimator by simply learning the geometry of M via its intrinsic distance d_M in (2). This can be done by efficiently estimating d_M in dimension $d = 1$ thanks to the Isomap method as coined by Tenenbaum et al. [60]. Therefore, in the special case when the dimension d of M is known and equal to 1, we are able to construct an estimator that achieves in Theorem 3.3 the rate $n^{-\beta / (2\beta + 1)}$, therefore establishing that in dimension $d = 1$ at least, the regularity of the manifold M does not affect the minimax rate for estimating f even when M is unknown. However, the volume density function ϑ_{x_0} is not constant as soon as $d \geq 2$ and obtaining a global picture in higher dimensions remains an open and presumably challenging problem.

1.3. Organisation of the paper

In Section 2, we provide with all the necessary material and notation from classical geometry for the unfamiliar reader. Section 2.1 together with the construction of smoothness spaces – here Hölder spaces on a submanifold in Section 2.2. We elaborate in particular on the reach of a subset of the Euclidean space in Section 2.3 and construct a statistical model for sampling n data from a density f with regularity β living on an unknown submanifold M of unknown dimension d and smoothness α in an ambient space of dimension D in Section 2.4. In this setting, we establish in Section 2.5 that a reach condition, *i.e.* assuming that the reach of M is bounded below, is necessary in order to reconstruct d . This is stated precisely in Theorem 2.6.

We give our main results in Section 3 and more specifically in Section 3.1. When the dimension d and the smoothness parameters α of the unknown manifold M and the smoothness β of f are known, Theorem 3.1 states the existence of an estimator that achieves the rate $n^{-\alpha \wedge \beta / (2\alpha \wedge \beta + d)}$ in expected pointwise loss, and Theorem 3.2 establishes that a minimax lower bound is $n^{-\beta / (2\beta + d)}$. Theorem 3.3 shows the existence of estimators in dimension $d = 1$ that achieve the rate $n^{-\beta / (2\beta + 1)}$, which is therefore minimax in that case. Theorem 3.4 states the existence of smoothness and dimension adaptive estimators, when α, β and d are unknown. Section 3.2 elaborates on special kernels upon which the estimators that achieve the aforementioned results are constructed, and their properties with respect to bias and variance analysis. The underlying geometry of M makes the usual orthogonality to non-constant polynomials of a certain degree (the order of the kernel) irrelevant, and a specific construction must be undertaken. Section 3.3 focuses on the case of one-dimensional submanifolds M when $d = 1$, where we explicitly construct a kernel estimator that achieves the minimax rate of convergence, revisiting the estimator (2) of Pelletier [54] and relying on the Isomap algorithm. In Section 3.4, we implement Lepski's algorithm on the bandwidth of our kernel estimators, following Lepski et al. [46]; this achieves smoothness adaptation w.r.t. $\alpha \wedge \beta$. Finally, in Section 3.5, we build an estimator of the dimension d of M , following ideas of Farahmand et al. [20] and that enables us to obtain simultaneous adaptation w.r.t. $\alpha \wedge \beta$ and d by plug-in.

Finally, numerical examples are developed in Section 4: we elaborate on examples of non-isometric embeddings of the circle and the torus in dimension 1 and 2 and explore in particular rates of convergence on Monte-Carlo simulations, illustrating how effective Lepski's method can be in that context. The proof are delayed until Appendix A.

2. Manifold-supported probability distributions

2.1. Some material from geometry

We endow \mathbb{R}^D with its usual Euclidean product and norm, respectively denoted by $\langle \cdot, \cdot \rangle$ and $\| \cdot \|$. We denote by $B(x, r)$ the open ball of \mathbb{R}^D of center x and

radius r . For any subspace $H \subset \mathbb{R}^D$, we set $B_H(x, r) = H \cap B(x, r)$ for the open ball in H for the induced norm.

We recall some basic notions of geometry of submanifolds of the Euclidean space \mathbb{R}^D for the unfamiliar reader. We borrow material from the classical textbooks by Gallot et al. [25] and Lee [44]. In all the paper, we consider C^∞ Riemannian manifolds (M, g) that we informally call *smooth*, which is an abstract manifold M endowed with a C^∞ atlas (Gallot et al. [25, Def 1.6 p.6]) and a C^∞ metric g (Gallot et al. [25, Def 2.1 p.52]). Such a manifold can always be embedded isometrically into some Euclidean space (Gromov and Rokhlin [31]), meaning that the pull-back of the canonical Euclidean metric coincides with the metric of the manifold. When we consider a smooth submanifold $M \subset \mathbb{R}^D$, we mean that M is the image through such an embedding of a smooth abstract Riemannian manifold.

Since we quantitatively compare the smoothness of manifolds within a large class of models, we need to pick a canonical parametrisation. For this reason, we consider the *exponential map* (Gallot et al. [25, Def 2.86 p.85]); for any smooth submanifold $M \subset \mathbb{R}^D$ and any $x \in M$, it defines a smooth parametrisation

$$\exp_x : B_{T_x M}(0, \varepsilon) \rightarrow M$$

of M around x , provided that ε is chosen small enough (Gallot et al. [25, Cor 2.89 p.86]). The supremum of all such ε is called the injectivity radius at x and is denoted $\text{inj}_M(x)$. When M is a closed subset of \mathbb{R}^D , the exponential maps are well defined on the whole tangent spaces. This is (one side of) the Hopf-Rinow theorem Lee [44, Thm 6.13 p.108].

Given a submanifold M of dimension d , the *volume measure* of M , denoted by μ_M , is the restriction of the d -dimensional Hausdorff measure \mathcal{H}^d to M , see Federer [22, Sec 2.10.2 p.171] for a definition. It can be shown (Evans and Gariepy [19, Ex D p.102]) that this definition coincides with the usual one of volume measure of a Riemannian manifold: if $\psi : M \rightarrow \mathbb{R}$ is a continuous function with support in $\exp_x(B_{T_x M}(0, \varepsilon))$ for ε smaller than $\text{inj}_M(x)$, we have

$$\mu_M(\psi) = \int_{B_{T_x M}(0, \varepsilon)} \psi \circ \exp_x(v) \sqrt{\det g^x(v)} dv,$$

with $g_{ij}^x(v) = \langle d\exp_x(v)[e_i], d\exp_x(v)[e_j] \rangle$ and where (e_1, \dots, e_d) is an arbitrary orthonormal basis of $T_x M$. We refer to Gallot et al. [25, Sec 3.H.1 and Sec 3.H.2] for further details on the volume measure. The volume of M , denoted by $\text{vol } M$, is simply $\mu_M(\mathbb{1})$. It is finite when M is a compact submanifold of \mathbb{R}^D .

2.2. Hölder spaces on submanifolds of \mathbb{R}^D

Let M be a smooth submanifold of \mathbb{R}^D and let $\rho > 0$. We say that a vector-valued function $\varphi : M \rightarrow \mathbb{R}^m$ with $m \geq 1$ is γ -Hölder with $\gamma > 0$ if for all $x \in M$, the map

$$\varphi \circ \exp_x : B_{T_x M}(0, \rho_x) \rightarrow \mathbb{R}^m \quad \text{where} \quad \rho_x = \rho \wedge \text{inj}_M(x) \quad (3)$$

is γ -Hölder in the usual sense, namely

- (i) $\varphi \circ \exp_x$ is $k = \lceil \gamma - 1 \rceil$ -times differentiable;
- (ii) and verifies

$$\forall v, w \in B_{T_x M}(0, \rho), \quad \|d^k(\varphi \circ \exp_x)(v) - d^k(\varphi \circ \exp_x)(w)\|_{\text{op}} \leq R \|v - w\|^\delta$$

with $\delta = \gamma - k > 0$ and for some $R > 0$.

We denote by $\mathcal{H}^\gamma(M, \mathbb{R}^m)$ the space of all such functions, and define for $\varphi \in \mathcal{H}^\gamma(M, \mathbb{R}^m)$ the Hölder semi-norm

$$|\varphi|_\gamma = \sup_{x \in M} \sup_{v, w \in B_{T_x M}(0, \rho_x)} \frac{\|d^k(\varphi \circ \exp_x)(v) - d^k(\varphi \circ \exp_x)(w)\|_{\text{op}}}{\|v - w\|^\delta}.$$

The characterisation of the smoothness of a function $\varphi : M \rightarrow \mathbb{R}^m$ through the exponential maps is a classical way to define functional spaces over Riemannian manifolds, see for instance Triebel [61].

Remark 2.1. The functional set $\mathcal{H}^\gamma(M, \cdot)$ also depends on the parameter ρ , merely introduced for a technical reason and dropped from the notation. When the manifold M has a reach (defined in the next Section 2.3) bounded from below by $\tau > 0$, a natural choice for ρ is $\pi\tau$, in light of the results of Proposition 2.3 below.

Remark 2.2. This definition of Hölder smoothness is also intrinsic, meaning that it does not depend on the way M is embedded into \mathbb{R}^D . If indeed $\psi : N \rightarrow M$ is an isometry between N and M , then $|\varphi|_\gamma = |\varphi \circ \psi|_\gamma$ for any $\varphi : M \rightarrow \mathbb{R}^d$ and $\gamma > 0$.

2.3. The reach of a subset

One of the main concerns when dealing with observations sampled from a geometrically structured probability measure is to determine the suitable scale at which one should look at the data. Indeed, given finite-sized point cloud in \mathbb{R}^D , there are infinitely many submanifolds that interpolate the point cloud, see Figure 1 for an illustration. A popular notion of regularity for a subset of the Euclidean space is the *reach*, introduced by Federer [21].

Definition 2.1. Let K be a compact subset of \mathbb{R}^D . The reach τ_K of K is the supremum of all $r \geq 0$ such that the orthogonal projection pr_K on K is well-defined on the r -neighbourhood K^r of K , namely

$$\tau_K = \sup\{r \geq 0 \mid \forall x \in \mathbb{R}^D, d(x, K) \leq r \Rightarrow \exists! y \in K, d(x, K) = \|x - y\|\}.$$

When M is a compact submanifold of \mathbb{R}^D , the reach τ_M quantifies two geometric invariants: locally, it measures how curved the manifold is, and globally, it measures how close it is to intersect itself (the so-called bottleneck effect). See Figure 2 for an illustration of the phenomenon. A *reach condition*, meaning that




FIG 1. An arbitrary points cloud (Left) for $D = 2$, and two smooth one-dimensional submanifolds passing through all its points (Middle, Right). A reach condition tends to discard the Right manifold as a likely candidate among all possible submanifolds the point cloud is sampled from.

the reach is bounded below, is necessary in order to obtain minimax inference results in manifold learning. These include: homology inference (Niyogi et al. [51], Balakrishnan et al. [6]), curvature (Aamari and Levrard [3]) and reach estimation itself (Aamari et al. [2]) as well as manifold estimation (Genovese et al. [26], Aamari and Levrard [3]).




FIG 2. For the first manifold M (Left), the value of the reach τ_M comes from its curvature. For the second one (Right), the reach is equal to τ_M because it is close to self intersecting (a bottleneck effect). The blue area represents the tubular neighbourhood over which the orthonormal projection on each manifold is well-defined.

2.4. A statistical model for sampling on a unknown manifold

In the following, we fix a point $x_0 \in \mathbb{R}^D$ in the ambient space. See Section 1.1 for a discussion on such a setting. Our statistical model is characterized by two quantities: the regularity of its support and the regularity of the density defined on this support. The support belongs to a class of submanifolds M , for which we need to fix some kind of *canonical* parametrisation. This is what Aamari and Levrard [3] propose by asking the support M to admit a local parametrisation at all point $x \in M$ by $T_x M$, and that this parametrisation is close to being the inverse of the projection over this tangent space. We follow this idea by imposing a constraint on the exponential map. In the following, we take $\tau > 0$ and set $\rho = \pi\tau$ in the definition of the Hölder spaces of Section 2.2.

Definition 2.2. Let $1 \leq d < D$ be integers. We let $\mathcal{C}_d(\tau)$ define the set of submanifolds M of \mathbb{R}^D that contain x_0 and satisfying the following properties:

- (i) (Dimension) M is a smooth submanifold of dimension d without boundaries;
- (ii) (Compactness) M is compact;
- (iii) (Reach condition) We have $\tau_M \geq \tau$.

For $\alpha \geq 1$ and $L > 0$, we define $\mathcal{C}_{d,\alpha}(\tau, L)$ as the set of $M \in \mathcal{C}_d(\tau)$ that fulfill the additional condition:

- (iv) The inclusion map $\iota_M : M \hookrightarrow \mathbb{R}^D$ is $(\alpha + 1)$ -Hölder with $|\iota_M|_{\alpha+1} \leq L$.

Remark 2.3. The definitions above endow our model with *global* constraints, even though most of them can be stated in a *local* fashion, with properties of the support holding in a neighbourhood of our candidate point $x \in \mathbb{R}^D$. This meets two expectations:

- staying close to the existing manifold setting in statistics, like in Aamari and Levrard [3];
- allowing for further developments, like estimation in global losses, such as L^p -norms, Wasserstein norms, or the sup-norm (Wu and Wu [65]).

The reach condition $\tau_M \geq \tau > 0$ in (iii) is essential in estimating consistently a density at a point in our setting, as shown in Theorem 2.6 in Section 2.5. Furthermore, a reach constraint enables one to benefit from several interesting geometric properties.

Proposition 2.3. Let M be a compact smooth submanifold of \mathbb{R}^D with $\tau_M \geq \tau$. Then the injectivity radius inj_M is everywhere greater than $\pi\tau$.

This result is a corollary of Alexander and Bishop [4, Thm 1.3], as explained in Aamari and Levrard [3, Lem A.1]. Pick $M \in \mathcal{C}_d(\tau)$. For any $x \in M$, the map $v \mapsto \exp_x(v) - x$ is bounded above by $\pi\tau$ on $B_{T_x M}(0, \pi\tau)$, since for any $v \in B_{T_x M}(0, \pi\tau)$, we have $\|\exp_x(v) - x\| \leq d_M(\exp_x(v), x) = \|v\|$, where d_M is the intrinsic distance on M . This uniform bound along with the Hölder condition (iv) allows one to obtain a uniform bound on the first derivatives of the exponential map.

Lemma 2.4. For $M \in \mathcal{C}_{d,\alpha}(\tau, L)$, any $x \in M$, and any $1 \leq j \leq \lceil \alpha \rceil$, we have

$$\sup_{v \in B_{T_x M}(0, \pi\tau/2)} \|d^j \exp_x(v)\|_{\text{op}} \leq L_j,$$

with L_j depending on d, τ, L and α only.

See Lemma A.6 in the appendix for further details on the proof. In the light of this result, the model of Definition 2.2 is thus quite close to the one proposed by Aamari and Levrard [3].

Remark 2.4. There is no obvious equivalence between M having a reach greater than τ , and M being Hölder, in particular because the reach is a global quantity,

while the Hölder smoothness property, as defined in Section 2.2, is a local feature. However, having a reach greater than τ implies that $\|d^2 \exp_x(0)\| \leq 2/\tau$, in light of Proposition A.1. For this reason, we always have that the submanifold M is at least 2-Hölder.

We are ready to define the class of density functions that we study, built upon submanifolds in the class $\mathcal{C}_d(\tau)$.

Definition 2.5. *Let $1 \leq d \leq D - 1$, $\alpha \geq 1$, $\beta > 0$, $\tau > 0$, $L > 0$, $R > 0$ and $0 \leq f_{\min} < f_{\max}$. We define $\Sigma_{\alpha,\beta}^d(\tau, L, f_{\min}, f_{\max}, R)$, or $\Sigma_{\alpha,\beta}^d$ for short, as the set of probability measures P on \mathbb{R}^D (endowed with its Borel σ -field) such that*

- (i) *There exists $M_P \in \mathcal{C}_{d,\alpha}(\tau, L)$ such that $\text{supp } P = M_P$;*
- (ii) *There exists a version of the Radon-Nikodym derivative $\frac{dP}{d\mu_{M_P}}$, denoted by f_P , that belongs to $\mathcal{H}^\beta(M_P, \mathbb{R})$;*
- (iii) *This version satisfies $f_{\min} \leq f_P \leq f_{\max}$ and $|f_P|_\beta \leq R$.*

Some remarks: **1)** The support of any $P \in \Sigma_{\alpha,\beta}^d(\tau, L, f_{\min}, f_{\max}, R)$ contains the candidate point x_0 by construction, see Definition 2.2 where $\mathcal{C}_{d,\alpha}(\tau, L)$ is defined. **2)** Condition (i) discards the possibility that f_P is zero on non-null subset of M ; in particular f_P is non zero around x_0 (but can be zero at x_0 nonetheless). This ensures that x_0 does not lie too far from the data. An alternate definition is to impose a condition like $P \ll \mu_M$. This leads to the same results in the next sections, but with a slight ambiguity in the choice of M . **3)** The parameters in subscript or superscript (d, α, β) control the rate of convergence of the estimation, while the parameters ($\tau, L, f_{\min}, f_{\max}, R$) control the pre-factor in the rates of convergence. For notational simplicity, we sometimes omit them when no confusion can be made.

2.5. Choice of a loss function and the reach assumption

For $P \in \Sigma_{\alpha,\beta}^d$ and a n -sample (X_1, \dots, X_n) drawn from P , our goal is to recover the value of $f_P(x_0)$ thanks to an estimator $\widehat{f}(x_0)$ built on top of the data (X_1, \dots, X_n) . We measure the accuracy of estimation by the maximal expected risk or order p , for $p \geq 1$, defined by

$$\sup_{P \in \Sigma_{\alpha,\beta}^d} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}(x_0) - f_P(x_0)|^p]^{1/p}$$

We look for an estimator with the smallest possible maximal risk as the number of observations n goes to ∞ . We first show that if we let $\tau = 0$, *i.e.* if we do not impose a reach condition, then it is impossible to estimate $f_P(x_0)$ consistently as $n \rightarrow \infty$ for any estimator, thus establishing that the reach assumption $\tau > 0$ is unavoidable.

Theorem 2.6. *In the setting of Definition 2.5, if we let $\tau = 0$, the following*

lower bound holds

$$\inf_{\hat{f}(x)} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \geq \frac{1}{2}(f_{\max} - f_{\min}) > 0,$$

where the infimum is taken over all estimators $\hat{f}(x_0)$ of $f_P(x_0)$.

The proof is given in Appendix A.2. This result is in line with a reach condition $\tau > 0$, a customary necessary condition in a minimax reconstruction in geometric inference, when the manifold is unknown, see Niyogi et al. [51], Genovese et al. [26], Balakrishnan et al. [6], Kim et al. [40], Aamari and Levrard [3] and the references therein.

3. Density estimation at a fixed point

Recall that we fix a point $x_0 \in \mathbb{R}^D$ where we wish to estimate f_P . Throughout the section, the symbols \lesssim and \gtrsim denote inequalities up to a constant that, unless specified otherwise, depends on the parameters $d, \alpha, \beta, \tau, L, f_{\min}, f_{\max}, R$ and p . The expression *for n large enough* means for n bigger than a constant that depends on the same parameters.

3.1. Main results

Let $D \geq 2$, $\tau > 0$, $L > 0$, $R > 0$, $0 \leq f_{\min} < f_{\max}$ and $p \geq 1$. Recall that we write $\Sigma_{\alpha, \beta}^d$ for short for $\Sigma_{\alpha, \beta}^d(\tau, L, f_{\min}, f_{\max}, R)$ as defined in Definition 2.5. The main results of this section are the following.

Theorem 3.1 (Upper bound). *For any $1 \leq d \leq D - 1$, $\alpha \geq 1$ and $\beta > 0$, there exists an estimator $\hat{f}(x_0)$ – explicitly constructed in Section 3.2 below – depending on α, β and d , such that, for n large enough,*

$$\sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim n^{-\alpha \wedge \beta / (2\alpha \wedge \beta + d)}.$$

The estimator of Theorem 3.1 is a kernel density estimator that depends on α, β and d through the choice of the kernel and its order (in a certain sense specified below), together with its bandwidth. Its analysis is given in Section 3.2. The estimator is indeed optimal in a minimax sense, as soon as $\alpha \geq \beta$.

Theorem 3.2 (Lower bound). *Let $1 \leq d \leq D - 1$, $\alpha \geq 1$ and $\beta > 0$. If L and f_{\max} are large enough and if f_{\min} is small enough (depending on τ), then*

$$\liminf_{n \rightarrow \infty} n^{\beta / (2\beta + d)} \inf_{\hat{f}(x_0)} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \geq C_* > 0$$

where C_* only depends on τ and R .

See Appendix A.3 for a proof. The rates from Theorem 3.1 and Theorem 3.2 agree, provided the underlying manifold M is regular enough, namely that $\alpha \geq \beta$. This probably covers most cases of interest in practice. However, when $\alpha < \beta$ the question of optimality remains. We investigate in Section 3.3 below the simpler case $d = 1$ and show that it is then possible to achieve the rate $n^{-\beta/(2\beta+1)}$, at the extra cost of learning the geometry of M in a specific sense.

Theorem 3.3 (One-dimensional case). *Let $d = 1$ and $\beta > 0$. Assume that $f_{\min} > 0$. Then there exists an estimator $\hat{f}^{1D}(x_0)$ – explicitly constructed in Section 3.3 below – depending on β , such that, for any $\alpha \geq 1$ and for n large enough,*

$$\sup_{P \in \Sigma_{\alpha, \beta}^1} \mathbb{E}_{P^{\otimes n}} [|\hat{f}^{1D}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim n^{-\beta/(2\beta+1)}.$$

The estimator described in Theorem 3.1 and Theorem 3.3 requires the specification of α , β and d , that are usually unknown in practice. We can circumvent this impediment by building an adaptative procedure with respect to these parameters. In Section 3.4 we adapt to the smoothness parameters α and β by implementing Lepski’s method (Lepskii [47]); in Section 3.5, we adapt to d by plugging-in a dimension estimator. We obtain the following result:

Theorem 3.4 (Adaptation). *Let $\ell \geq 0$. Assume that $f_{\min} > 0$. Then, there exists an estimator $\hat{f}^{\text{adapt}}(x_0)$ – explicitly constructed in Section 3.5 below – depending on ℓ such that, for any $\alpha \geq 1$, β in $[0, \ell]$ and any $1 \leq d \leq D - 1$, we have, for n large enough,*

$$\sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}^{\text{adapt}}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \left(\frac{\log n}{n} \right)^{\frac{\alpha \wedge \beta}{2\alpha \wedge \beta + d}}.$$

We were unable to obtain oracle inequalities in the spirit of the Goldenshluger-Lepski method, see Goldenshluger and Lepski [28, 29, 30], due to the non-Euclidean character of the support of f_P : our route goes along the more classical approach of Lepski et al. [46]. Obtaining oracle inequalities in this framework remain an open problem.

3.2. Kernel estimation

Classical nonparametric density estimation methods are based on kernel smoothing Parzen [53], Silverman [57]. In this section, we combine kernel density estimation with the minimal geometric features needed in order to recover efficiently their density. Since the intrinsic dimension d is not prone to change in this section, we further drop d in (most of) the notation. The proofs of this section can be found in Appendix A.4.

Let $K : \mathbb{R}^D \rightarrow \mathbb{R}$ be a smooth function vanishing outside the unit ball $B(0, 1)$. Given an n -sample (X_1, \dots, X_n) drawn from a distribution P on \mathbb{R}^D , we are

interested in the behaviour of the kernel estimator

$$\widehat{f}_h(x_0) = \frac{1}{nh^d} \sum_{j=1}^n K\left(\frac{X_j - x_0}{h}\right), \quad h > 0. \quad (4)$$

Note that the normalisation here is h^d and not h^D as one would set for a classical kernel estimator in \mathbb{R}^D . Our main result is that $\widehat{f}_h(x_0)$ behaves well when P is supported on a d -dimensional submanifold of \mathbb{R}^D .

We need some notation. For $P \in \Sigma_{\alpha,\beta}^d$, define

$$f_h(P, x_0) = \mathbb{E}_{P^{\otimes n}}[\widehat{f}_h(x_0)],$$

$$\mathcal{B}_h(P, x_0) = f_h(P, x_0) - f_P(x_0),$$

and

$$\widehat{\xi}_h(P, x_0) = \widehat{f}_h(x_0) - f_h(P, x_0),$$

that correspond respectively to the mean, bias and stochastic deviation of the estimator $\widehat{f}_h(x_0)$. We also introduce the quantity

$$\Omega(h) = \sqrt{\frac{2\omega}{nh^d}} + \frac{\|K\|_\infty}{nh^d} \quad \text{with} \quad \omega = 4^d \zeta_d \|K\|_\infty^2 f_{\max},$$

where ζ_d is the volume of the unit ball in \mathbb{R}^d . The quantity $\Omega(h)$ will prove to be a good majorant of the stochastic deviations of $\widehat{f}_h(x_0)$. The usual bias-stochastic decomposition of $\widehat{f}_h(x)$ leads to

$$\mathbb{E}_{P^{\otimes n}}[|\widehat{f}_h(x_0) - f_P(x_0)|^p]^{1/p} \leq |\mathcal{B}_h(P, x_0)| + (\mathbb{E}_{P^{\otimes n}}[|\widehat{\xi}_h(P, x_0)|^p])^{1/p}. \quad (5)$$

We study each term separately. The stochastic term can readily be bounded from above.

Proposition 3.5. *Let $p \geq 1$. There exists a constant $c_p > 0$ depending on p only such that for any $P \in \Sigma_{\alpha,\beta}^d$ and any $h < \tau/2$:*

$$\left(\mathbb{E}_{P^{\otimes n}}[|\widehat{\xi}_h(P, x_0)|^p]\right)^{1/p} \leq c_p \Omega(h).$$

Now we turn to the bias term. We need certain properties for the kernel K . More precisely, we assume that

Assumption 3.6. *The kernel K verifies*

- (i) K is smooth and supported on the unit ball $B(0, 1)$;
- (ii) For any d -dimensional subspace H of \mathbb{R}^D , we have $\int_H K(v)dv = 1$.

One way to obtain Assumption 3.6 is to set $\Lambda(x) = \exp(-1/(1-\|x\|^2))$ for $x \in B(0, 1)$ and $\Lambda(x) = 0$ otherwise. Since Λ is rotationally invariant, its integral is the same over any d -subspace H of \mathbb{R}^D . Thus, with $\lambda_d = \int_{H_0} \Lambda(v)dv$ where $H_0 = \mathbb{R}^d \times \{0_{\mathbb{R}^{D-d}}\}$, the function $K(x) = \lambda_d^{-1} \Lambda(x)$ is a smooth kernel, supported on the unit ball of the ambient space \mathbb{R}^D that satisfies Assumption 3.6. In the following, we pick an arbitrary kernel K such that Assumption 3.6 is satisfied.

Lemma 3.7. For $P \in \Sigma_{\alpha,\beta}^d$ and any $h < \tau/2$, setting $k = \lceil \alpha \wedge \beta - 1 \rceil$, we have

$$f_h(P, x_0) = f(x_0) + \sum_{j=1}^k h^j G_j(P, x_0) + R_h(P, x_0), \quad (6)$$

with $|G_j(P, x_0)| \lesssim 1$ and $|R_h(P, x_0)| \lesssim h^{\alpha \wedge \beta}$.

The existence of such an expansion allows, by carefully choosing the kernel, to cancel the intermediate terms. Starting from a kernel K satisfying Assumption 3.6, we recursively define a sequence of smooth kernels $(K^{(d,\ell)})_{\ell \geq 1}$, simply denoted by $K^{(\ell)}$ in this section, with support in $B(0, 1)$ as follows (see Figure 3). For $z \in \mathbb{R}^D$, we put

$$\begin{cases} K^{(1)}(z) = K(z) \\ K^{(\ell+1)}(z) = 2^{1+d/\ell} K^{(\ell)}(2^{1/\ell} z) - K^{(\ell)}(z) \quad \forall \ell \geq 1. \end{cases} \quad (7)$$

A few remarks can be made: **1)** In a classical kernel density estimation framework, the integer $\ell-1$ plays the role of the order of the kernel. **2)** The assumption that K is compactly supported is seemingly quite strong. This is a way to make sure that the support of $x \mapsto K((x-x_0)/h)$ is within the injectivity ball of the map \exp_{x_0} for any $h < \pi\tau$. **3)** The construction of K is simply an example of a Richardson's extrapolation as coined by Richardson [55]. **4)** This construction somewhat differs from the classical constructions than can be found in textbooks such as Tsybakov [62]. There is one practical reason: we require that all the kernels satisfy Assumption 3.6; another reason that appears to be more intrinsically related to our model: since the Euclidean distance is only a second order approximation of the Riemannian distance on M , defining a kernel through orthogonality relations with respect to a family of polynomials is not sufficient in our framework.

Proposition 3.8. Let ℓ be an integer greater than α and β , and let $K^{(\ell)}$ be the kernel defined in (7) starting from a kernel K satisfying Assumption 3.6. Then, for any $P \in \Sigma_{\alpha,\beta}^d$ and any bandwidth $h < \tau/2$, the estimator $\widehat{f}_h(x_0)$ defined as in (4) using $K^{(\ell)}$ is such that

$$|\mathcal{B}_h(P, x_0)| \lesssim h^{\alpha \wedge \beta}. \quad (8)$$

Patching together Proposition 3.5 and Proposition 3.8 yields

$$\sup_{P \in \Sigma_{\alpha,\beta}^d} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}_h(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \Omega(h) + h^{\alpha \wedge \beta},$$

for any $h < \tau/2$. Therefore the estimator $\widehat{f}(x_0) = \widehat{f}_h(x_0)$ specified with $h = n^{-1/(2\alpha \wedge \beta + d)}$ indeed satisfies the conclusion of Theorem 3.1, for n large enough.




FIG 3. Plots of the kernel $K^{(d, \ell)}$ for $d = 1$ and $\ell = 1, 2, 3$.

3.3. A minimax estimator in the case $d = 1$, $\alpha < \beta$

The gap we observe between the two rates in Theorem 3.1 and Theorem 3.2 leads to the following question: does the regularity $\alpha \geq 1$ of M have a genuine limiting effect in the estimation of $f_P(x_0)$, or does it rather reveal a weakness of the estimator described in Section 3.2? We do not have a definitive answer to this question except for $d = 1$ *i.e.* when M is a closed curve in an Euclidean space. We can then show that the parameter α does not interfere at all with the density estimation. The proofs of this section can be found in Appendix A.5.

If $d = 1$, any submanifold M in $\mathcal{C}_1(\tau)$ is a closed smooth injective curve that can be parametrized by a unit-speed path $\gamma_M : [0, L_M] \rightarrow \mathbb{R}^D$ with $\gamma_M(0) = \gamma_M(L_M)$ and with $L_M = \text{vol } M$ being the length of the curve. In that case, the volume density function is trivial.

Lemma 3.9. *For $M \in \mathcal{C}_1(\tau)$, for any $x \in M$ and any $v \in T_x M$, we have $\det g^x(v) = 1$.*

Thanks to Lemma 3.9, the estimator proposed by Pelletier [54] takes a simpler form, which we will try to take advantage of. Indeed, in the representation (2) of Pelletier [54], only d_M remains unknown. We now show how to efficiently estimate d_M thanks to the Isomap method as coined by Tenenbaum et al. [60]. The analysis of this algorithm essentially comes from Bernstein et al. [8] and is pursued in Arias-Castro and Le Gouic [5], but the bounds obtained there are manifold dependent. We thus propose a slight modification of their proofs in order to obtain uniform controls over $\mathcal{C}_1(\tau)$, and make use of the simplifications coming from the dimension 1. Indeed, for $d = 1$, we have the following simple and explicit formula for the intrinsic distance on M :

$$d_M(\gamma_M(s), \gamma_M(t)) = |t - s| \wedge (L_M - |t - s|) \quad \forall s, t \in [0, L_M].$$

The Isomap method can be described as follows: let $\varepsilon > 0$, and let \mathcal{G}_ε be

the ε -neighbourhood graph built upon the data (X_1, \dots, X_n) and x_0 — namely, $\mathcal{G}_\varepsilon = (V, E)$ where $V = (x_0, X_1, \dots, X_n)$, and where $E = \{(y, z) \in V \mid \|y - z\| \leq \varepsilon\}$. For a path in \mathcal{G}_ε (meaning: a sequence of adjacent vertices) $s = (p_0, \dots, p_m)$, we define its length as $L_s = \|p_1 - p_0\| + \dots + \|p_m - p_{m-1}\|$. The distance between x and a vertice y in the graph \mathcal{G}_ε is then defined as

$$\widehat{d}_\varepsilon(x, y) = \min \{L_s \mid s \text{ path in } \mathcal{G}_\varepsilon \text{ connecting } x \text{ to } y\}, \quad (9)$$

and we set this distance to ∞ if x and y are not connected. We are now ready to describe our estimators $\widehat{f}^{1D}(x_0)$. For any $h, \varepsilon > 0$, we set

$$\widehat{f}_h^{1D}(x_0) = \frac{1}{nh} \sum_{i=1}^n K^{1D}(\widehat{d}_\varepsilon(x_0, X_i)/h), \quad (10)$$

for some kernel $K^{1D} : \mathbb{R} \rightarrow \mathbb{R}$. Notice that the kernel $K^{(1,\ell)}(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}$ defined in Section 3.2 starting from kernel $K = \lambda_1^{-1}\Lambda$ can be put in the form $K^{(1,\ell)}(x) = K^{(1,\ell)}(\|x\|)$ with $K^{(1,\ell)}(\cdot)$ denoting thus (with a slight abuse of notation) both functions starting from either \mathbb{R} or \mathbb{R}^D . We choose this kernel in the next statement.

Proposition 3.10. *Assume that $f_{\min} > 0$. The estimator defined in (10) above and specified with $K^{1D} = K^{(1,\ell)}(\cdot)$ satisfies the following property: for any $\beta \in [0, \ell]$ and any $\alpha \geq 1$, we have*

$$\sup_{P \in \Sigma_{\alpha,\beta}^1} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}_h^{1D}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \frac{\varepsilon^2}{h^2} + \Omega(h) + h^\beta + \frac{1}{nh},$$

with

$$\varepsilon = \frac{32(p+1)}{f_{\min}} \times \frac{\log n}{n},$$

for $h < \tau/4$ and n large enough.

The proof of Theorem 3.3 readily follows from Proposition 3.10 using the estimator $\widehat{f}^{1D} = \widehat{f}_h^{1D}$ with $h = n^{-1/(2\beta+1)}$.

3.4. Smoothness adaptation

We implement Lepski's algorithm, following closely Lepski et al. [46] in order to automatically select the bandwidth from the data (X_1, \dots, X_n) . We know from Section 3.2 that the optimal bandwidth on $\Sigma_{\alpha,\beta}^d$ is of the form $n^{-1/(2\alpha\wedge\beta)+d}$. Hence, without prior knowledge of the value of α and β , we can restrict our search for a bandwidth in a bounded interval of the form $[h^-, 1]$ discretized as follow

$$\mathbb{H} = \{2^{-j}, \text{ for } 0 \leq j \leq \log_2(1/h^-)\}$$

We pick

$$h^- = \left(\frac{\|K\|_\infty}{2\omega} \right)^{1/d} \frac{1}{n^{1/d}};$$

this bandwidth is always smaller than the optimal bandwidth $n^{-1/(2\alpha\wedge\beta+d)}$ on $\Sigma_{\alpha,\beta}^d$ for n large, and is such that $\Omega(h) \leq 2\sqrt{2\omega/(nh^d)}$ for all $h \geq h^-$. For $h, \eta \in \mathbb{H}$, we introduce the following quantities:

$$\begin{aligned} \lambda(h) &= 1 \vee \sqrt{\Theta d \log(1/h)}, \\ \psi(h, \eta) &= \Omega(h)\lambda(h) + \Omega(\eta)\lambda(\eta) \end{aligned} \tag{11}$$

where Θ is a positive constant (to be specified). For $h \in \mathbb{H}$ we define the subset of bandwidths $\mathbb{H}(h) = \{\eta \in \mathbb{H}, \eta \leq h\}$. The selection rule for h is the following:

$$\widehat{h}(x_0) = \max \left\{ h \in \mathbb{H} \mid \forall \eta \in \mathbb{H}(h), |\widehat{f}_h(x_0) - \widehat{f}_\eta(x_0)| \leq \psi(h, \eta) \right\},$$

and we finally consider the estimator

$$\widehat{f}(x_0) = \widehat{f}_{\widehat{h}(x_0)}(x_0), \tag{12}$$

where $\widehat{f}_h(x_0)$ is defined at (4).

Proposition 3.11. *Assume $\Theta > p$. Let $\ell \in \mathbb{N}$, and let $\widehat{f}(x_0)$ be the estimator defined in (12) using $K^{(\ell)}$ originated from a kernel K satisfying Assumption 3.6. Then, for any $\alpha \geq 1, \beta \in [0, \ell]$, we have, for n large enough*

$$\sup_{P \in \Sigma_{\alpha,\beta}^d} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \left(\frac{\log n}{n} \right)^{\alpha \wedge \beta / (2\alpha \wedge \beta + d)}.$$

The proof of Proposition 3.11 can be found in Appendix A.6. Some remarks:
1) Proposition 3.11 provides us with a classical smoothness adaptation result in the spirit of Lepski et al. [46]: the estimator \widehat{f} has the same performance as the estimator \widehat{f}_h selected with the optimal bandwidth $n^{-1/(2\alpha\wedge\beta+d)}$, up to a logarithmic factor on each model $\Sigma_{\alpha,\beta}^d$ without the prior knowledge of $\alpha \wedge \beta$ over the range $[0, \ell]$. **2)** The extra logarithmic term is the unavoidable payment for the Lepski-Low phenomenon (Lepskii [48], Low [49]) when recovering a function in pointwise or in a uniform loss.

3.5. Simultaneous adaptation to smoothness and dimension

The estimators considered in Theorem 3.1 or Proposition 3.11 heavily rely on the intrinsic dimension d through the choice of kernel satisfying Assumption 3.6, through the normalisation h^d and either through the choice of an optimal bandwidth h , or the selection procedure (11)-(12). We now show how to adapt to d considered as an unknown and nuisance parameter. The proofs of this section can be found in Appendix A.7.

We redefine all the quantities introduced before as now depending on d . Namely, for $h, \eta > 0$, and a given family of kernel $K(d; \cdot)$, we set

$$\widehat{f}_h(d; x_0) = \frac{1}{nh^d} \sum_{i=1}^n K(d; (X_i - x_0)/h).$$

$$\begin{aligned} \Omega(d; h) &= \sqrt{\frac{2\omega_d}{nh^d}} + \frac{\|K(d; \cdot)\|_\infty}{nh^d} \quad \text{with} \quad \omega_d = 4^d \zeta_d \|K(d; \cdot)\|_\infty^2 f_{\max}, \\ \lambda(d; h) &= 1 \vee \sqrt{\Theta d \log(1/h)}, \\ \psi(d; h, \eta) &= v(d; h)\lambda(d; h) + v(d; \eta)\lambda(d; \eta), \\ h_d^- &= (\|K(d; \cdot)\|_\infty / 2\omega_d)^{1/d} n^{1/d}, \\ \mathbb{H}_d &= \{2^{-j}, \text{ for } 0 \leq j \leq \log_2(1/h_d^-)\}, \end{aligned}$$

where Θ is a constant. We also define

$$\widehat{h}(d; x_0) = \max\{h \in \mathbb{H}_d \mid \forall \eta \in \mathbb{H}_d(h), \quad |\widehat{f}_h(d; x_0) - \widehat{f}_\eta(d; x_0)| \leq \psi(d; h, \eta)\} \quad (13)$$

with $\mathbb{H}_d(h) = \{\eta \in \mathbb{H}_d, \eta \leq h\}$. We are now left with the choice of kernel family $K(d; \cdot)$. For any $1 \leq d \leq D - 1$ and $h > 0$, we define

$$K^{(1)}(d; x) = \lambda_d^{-1} \Lambda(x)$$

where Λ and λ_d have been introduced in Section 3.2. We then pick an integer $\ell \in \mathbb{N}$ and choose

$$K(d; \cdot) = K^{(d, \ell)}(d; \cdot) \quad (14)$$

where $K^{(d, \ell)}(d; \cdot)$ is defined by recursion in (7) starting from the kernel $K^{(1)}(d; \cdot)$.

We assume that we have an estimator \widehat{d} of the dimension d of M with values in $\{1, \dots, D\}$. More precisely, we need the following property:

Assumption 3.12. For any $1 \leq d < D$ and $\alpha \geq 1, \beta \in (0, \ell]$, we have

$$\sup_{P \in \Sigma_{\alpha, \beta}^d} P^{\otimes n} (\widehat{d} \neq d) \lesssim n^{-3p/2},$$

where p is the exponent of the loss function.

If we are given such a estimator of the dimension d , then we can build a estimator that adapts to this parameter.

Proposition 3.13. Let $\widehat{f}(x_0) = \widehat{f}_{\widehat{h}}(\widehat{d}, x_0)$ built with the kernel family (14), where \widehat{d} is a estimator satisfying Assumption 3.12 and where $\widehat{h} = \widehat{h}(\widehat{d}, x_0)$ is defined at (13). Then, for any $1 \leq d \leq D - 1$ and $\alpha \geq 1, \beta \in (0, \ell]$, we have, for n large enough

$$\sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \left(\frac{\log n}{n}\right)^{\alpha \wedge \beta / (2\alpha \wedge \beta + d)}.$$

It only remains to show that there exists an estimator \widehat{d} satisfying Assumption 3.12 to obtain Theorem 3.4.

There are various way to define such an estimator, see Farahmand et al. [20] or even Kim et al. [40] where an estimator with super-exponential minimax rate on a wide class of probability measures is constructed. For sake of completeness and simplicity, we will mildly adapt the work of Farahmand et al. [20] to our setting.

The resulting estimator will behave well as soon as we add the assumption that $f_{\min} > 0$.

Definition 3.14. For a probability measure P , we write $P_\eta = P(B(x_0, \eta))$ for any $\eta > 0$, and $\hat{P}_\eta = \hat{P}_n(B(x_0, \eta))$ where $\hat{P}_n = n^{-1} \sum_{i=1}^n \delta_{X_i}$ denotes the empirical measure of the sample (X_1, \dots, X_n) . Define

$$\hat{\delta}_\eta = \log_2 \hat{P}_{2\eta} - \log_2 \hat{P}_\eta,$$

and set $\hat{\delta}_\eta = D$ when $\hat{P}_\eta = 0$. We define \hat{d}_η to be the closest integer of $\{1, \dots, D\}$ to $\hat{\delta}_\eta$, namely $\hat{d}_\eta = \lfloor \hat{\delta}_\eta + 1/2 \rfloor$.

Proposition 3.15. Assume that $f_{\min} > 0$. Then, for any $1 \leq d \leq D - 1$, and any $\alpha \geq 1, \beta > 0$, the estimator $\hat{d} = \hat{d}_\eta$ for $\eta = n^{-1/(2D+2)}$ verifies for n large enough

$$\sup_{P \in \Sigma_{\alpha, \beta}^d} P^{\otimes n} (\hat{d} \neq d) \leq 4 \exp \left(-2n^{1-(d+1)/(D+1)} \right).$$

4. Numerical illustration

In this section we propose a few simulations to illustrate the results presented above. The goal is two-fold

- To highlight the rate obtained in Theorem 3.1 using estimator $\hat{f}_h(x_0)$, in the case where $\beta \leq \alpha$, on arbitrary submanifold and for a carefully chosen bandwidth h ;
- To show the computational feasibility and performance of estimator $\hat{f}^{\text{adapt}}(x_0)$ described in Section 3.4.

For the sake of visualisation and simplicity, we focus on two typical examples of submanifold of \mathbb{R}^D , namely non-isometric embeddings of the flat circle $\mathbb{T}^1 = \mathbb{R}/\mathbb{Z}$ and of the flat torus $\mathbb{T}^2 = \mathbb{T}^1 \times \mathbb{T}^1$. In particular, these embeddings will be chosen in such way that their images, as submanifolds of \mathbb{R}^D , are not homogeneous compact Riemannian manifolds, so that the work of Kerkycharian et al. [38] for instance cannot be of use here.

For a given embedding $\Phi : N \rightarrow M \subset \mathbb{R}^D$ where N is either \mathbb{T}^1 or \mathbb{T}^2 , we construct absolutely continuous probabilities on M by pushing forward probability densities of N w.r.t. their volume measure. Indeed, if $Q = g \cdot \mu_N$, the push-forward measure $P = \Phi_* Q$ has density f with respect to μ_M given by

$$\forall x \in M, \quad f(x) = \frac{g(\Phi^{-1}(x))}{|\det d\Phi(\Phi^{-1}(x))|} \tag{15}$$

where the determinant is taken in an orthonormal basis of $T_{\Phi^{-1}(x)}N$ and $T_x M$, so that, if Φ is chosen smooth enough, f has the same regularity as g . If Φ is an embedding of \mathbb{T}^1 , we simply have $|\det d\Phi(y)| = \|\Phi'(y)\|$ for all $y \in \mathbb{T}^1$. If now Φ

maps \mathbb{T}^2 to M , we have

$$\begin{aligned}
 |\det d\Phi(y)| &= \sqrt{\det \langle d\Phi(y)[e_i], d\Phi(y)[e_j] \rangle_{1 \leq i, j \leq 2}} \\
 &= \sqrt{\|d\Phi(y)[e_1]\|^2 \|d\Phi(y)[e_2]\|^2 - \langle d\Phi(y)[e_1], d\Phi(y)[e_2] \rangle^2} \quad (16)
 \end{aligned}$$

where (e_1, e_2) is an orthonormal basis of $\mathbb{R}^2 \simeq T_y\mathbb{T}^2$.

Strictly speaking, the probability measures P exhibited below are not elements of the models $\Sigma_{\alpha, \beta}^d$, but we know that they locally coincide with some $\tilde{P} \in \Sigma_{\alpha, \beta}^d$ around our candidate point x , meaning that

$$P_{|B(x,r)} = \tilde{P}_{|B(x,r)} \quad \text{for some } r > 0.$$

This ensures that all the results displayed in Section 3 hold for P — see Remark 2.3 for a discussion on the local character of our setting.

4.1. An example of a density supported by a one-dimensional submanifold

Let $\beta \in \mathbb{N}^*$ and define the following function for $v \in [-1/2, 1/2]$

$$g_\beta(v) = C_\beta \times (1 - (-2v)^\beta) \mathbb{1}_{[-1/2, 0)}(v) + C_\beta (1 - (2v)^{\beta+1}) \mathbb{1}_{[0, 1/2]}(v). \quad (17)$$

where C_β is an explicit normalisation constant. The function g_β is positive and $\int_0^1 g_\beta(v) dt = 1$; it defines a probability density over $[-1/2, 1/2]$. Also, because the $(\beta - 1)$ -th derivative of g_β is 1-Lipschitz, but its β -derivative is discontinuous at $v = 0$, the function g_β is β -Hölder but *not* $(\beta + \varepsilon)$ -Hölder for any $\varepsilon > 0$. See Figure 4 for a few plots of the functions g_β . We next consider the parametric




FIG 4. Plots of the densities g_β for $\beta = 2, 4, 8$.

curve

$$\Phi : \begin{cases} \mathbb{T}^1 \rightarrow \mathbb{R}^2 \\ t \mapsto (\cos(2\pi v) + a \cos(2\pi\omega v), \sin(2\pi v) + a \sin(2\pi\omega v)). \end{cases}$$

Short computations show that Φ is indeed an embedding as soon as $a\omega < 1$, in which case $M = \Phi(\mathbb{T}^1)$ is indeed a smooth compact submanifold of \mathbb{R}^2 . For the rest of this section, we set $a = 1/8$ and $\omega = 6$. See Figure 5 for a plot of M with these parameters. We are interested in estimating the density f_β with




FIG 5. Plot of the submanifold M (Left) for parameters $a = 1/8$ and $\omega = 6$, and 500 points sampled independently from $\Phi_*g_\beta \cdot \mu_{\mathbb{T}^1}$ for $\beta = 3$ (Right). The black cross denotes the point $x = \Phi(0)$.

respect to $d\mu_M$ of the push-forward measure $P_\beta = \Phi_*g_\beta \cdot \mu_{\mathbb{T}^1}$, at point $x_0 = (1 + a, 0) \in M$. We use formula (15) to compute $f_\beta(x_0)$: We have $\Phi^{-1}(x_0) = 0$ and $\|\Phi'(0)\| = 2\pi(1 + a\omega)$ hence

$$f_\beta(x_0) = \frac{C_\beta}{2\pi(1 + a\omega)} \text{ at } x_0 = (1 + a, 0).$$

Our aim here is to provide an empirical measure for the convergence of the risk $n \mapsto \mathbb{E}_{P_\beta^{\otimes n}}[|\hat{f}_h(x_0) - f_\beta(x_0)|^p]^{1/p}$ when h is tuned optimally (in an oracle way). We pick $p = 2$. Our numerical procedure is detailed in Algorithm 1 below, and the numerical results are presented in Figure 8.

4.2. An example of a density supported by a two-dimensional submanifold

We consider a non-isometric embedding of the flat torus \mathbb{T}^2 . We first construct a density function. For an integer $\beta \geq 1$, define

$$G_\beta : (v, u) \in [-1/2, 1/2]^2 \mapsto g_\beta(v)g_\beta(u) \tag{18}$$

where g_β is defined as in (17). Obviously, G_β defines a density function on \mathbb{T}^2 that is β -Hölder (but not $(\beta + \varepsilon)$ -Hölder for any $\varepsilon > 0$).

Algorithm 1 MSE rate of convergence estimation

-
- 1: Provide integers $\beta \geq 1$ and $\ell \geq \beta$.
 - 2: Set a grid of increasing number of points $\mathbf{n} = (n_1, \dots, n_k) \in \mathbb{N}^k$ and a number of repetition N .
 - 3: **for** $n_i \in \mathbf{n}$ **do**
 - 4: Sample n_i points independently from P_β ,
 - 5: Compute $\hat{f}_{h_i}(x_0)$ with kernel $K^{(\ell)}$ and bandwidth $h_i = n_i^{-1/(2\beta+1)}$,
 - 6: Compute the square error $(\hat{f}_{h_i}(x) - f_\beta(x))^2$,
 - 7: Repeat the three previous steps N times,
 - 8: Average the errors to get a Monte-Carlo approximation $\hat{R}(n_i)$ of $\mathbb{E}_{P_\beta \otimes n_i} [|\hat{f}_h(x_0) - f_\beta(x_0)|^2]$.
 - 9: **end for**
 - 10: Perform an Ordinary Least Square Linear Regression on the curve $\log n_i \mapsto \log \hat{R}(n_i)$.
 - 11: **return** The coefficient of the linear regression.
-

We next consider the parametric surface

$$\Psi : \begin{cases} \mathbb{T}^2 \rightarrow \mathbb{R}^3 \\ (v, u) \mapsto \begin{pmatrix} (b + \cos(2\pi v)) \cos(2\pi u) + a \sin(2\pi \omega v) \\ (b + \cos(2\pi v)) \sin(2\pi u) + a \cos(2\pi \omega v) \\ \sin(2\pi v) + a \sin(2\pi \omega u) \end{pmatrix} \end{cases}$$

for some $a, b, \omega \in \mathbb{R}$. In the remaining of the section, we set $a = 1/8$, $b = 3$ and $\omega = 5$. We show that Ψ indeed defines an embedding. See Figure 7 for a plot of the submanifold $M = \Psi(\mathbb{T}^2)$. For an integer $\beta \geq 1$, we denote by F_β the density of the push forward measure $Q_\beta = \Psi_* G_\beta \cdot \mu_{\mathbb{T}^2}$ with respect to the volume measure μ_M . Let $x = (b + 1, a, 0)$ be the image of 0 by Ψ . Simple calculations show that the differential of Ψ at 0 evaluated at $e_1 = (1, 0) \in T_0 \mathbb{T}^2$




FIG 6. Plot of the probability density function G_β for $\beta = 3$.




FIG 7. Plot of the submanifold M (Top) for parameters $a = 1/8$, $b = 3$ and $\omega = 6$, and 500 points sampled independently from $\Psi_*G_\beta \cdot \mu_{\mathbb{T}^2}$ with $\beta = 3$ (Bottom). The black cross marks the point x .

and $e_2 = (0, 1) \in T_0\mathbb{T}^2$ is equal to respectively $d\Psi(0)[e_1] = 2\pi(a\omega, 0, 1)$ and $d\Psi(0)[e_2] = 2\pi(0, b + 1, a\omega)$. Hence formula (16) yields

$$\det d\Psi(0) = (2\pi)^2 \left((1 + a^2\omega^2) \left((b + 1)^2 + a^2\omega^2 \right) - a^2\omega^2 \right)^{1/2} \quad (19)$$

and we obtain

$$F_\beta(x_0) = \frac{C_\beta^2}{4\pi^2} \left((1 + a^2\omega^2) \left((b + 1)^2 + a^2\omega^2 \right) - a^2\omega^2 \right)^{-\frac{1}{2}}.$$

In the same way as in the previous section, we aim at providing an empirical measure for the rate of convergence of the risk $\mathbb{E}_{P_\beta^{\otimes n}} [|\hat{f}_h(x_0) - f_\beta(x_0)|^2]$ when h is suitably tuned with respect to n and β . This is done using again Algorithm 1. The results are presented in Figure 8.

4.3. Adaptation

In this section we estimate a density when its regularity is unknown, contrary to the previous simulation where the regularity parameter β is pugged in the bandwidth choice $n^{-1/(2\beta+d)}$. This is performed using Lepski's method presented in Section 3.4. The rate is computed using Algorithm 1, for both the one-dimensional and the two-dimensional synthetic datasets.




FIG 8. Plot of the empirical mean square error (blue) for a density supported by a one-dimensional submanifold (Top) and two-dimensional submanifold (Bottom) with parameter $\beta = 2$. We use a log-regular grid \mathbf{n} of 21 points ranging from 100 to 10^4 . Each experiment is repeated $N = 500$ times.

For the adaptive estimation on the two-dimensional manifold, we observe that the corrective term $\det d\Psi(0)$ computed in (19) results in a density $F_\beta(x_0)$ that is quite small, while the function ψ defined at (11) and used to tune the bandwidth soars dramatically because of the retained value of $\omega_d = 4^d \zeta_d \|K^{(d,\ell)}\|_\infty^2 f_{\max}$, so that the values of $\hat{f}_h(x_0)$ and $\psi(h, \cdot)$ (defined at (11)) are not of the same order anymore at this scale (using maximum 10^6 observations). To circumvent this effect, we introduce a scaling parameter λ as follows

$$G_{\beta,\lambda} = v, u \mapsto \lambda^2 G_\beta(\lambda v, \lambda u).$$

Like before, we consider the push-forward probability measure $\Psi_* G_{\beta,\lambda} \cdot \mu_{\mathbb{T}^2}$ which has density $F_{\beta,\lambda}$ with respect to μ_M . For $\lambda = 4$, we find that $F_{\beta,\lambda}(x_0)$ is of order 1 for most values of β , and we use the function $\psi^{\text{num}}(h, \eta) = \Omega^{\text{num}}(h)\lambda(h) + \Omega^{\text{num}}(\eta)\lambda(\eta)$ using simply $\Omega^{\text{num}}(h) = \sqrt{1/nh^d}$. We have no




FIG 9. (Top) Plot of the empirical median square error for the one-dimensional submanifold with $\beta = 2$. The bandwidth h is chosen adaptively using Lepski's method of order $\ell = 3$ as in Section 3.4. We used a log-regular grid \mathbf{n} of 11 points ranging from 100 to 10^4 and each experience was repeated $N = 500$ times. (Bottom) Same experiment but for the two dimensional manifold, with a grid ranging from 10^4 to 10^6 and $N = 100$ repetitions.

theoretical guarantee that such a method will work but we recover nonetheless the right rate in the estimation of the value of the density, see Figure 9 for a plot of the estimated rate.

We find an empirical error with a relatively high dispersion, hence our choice to represent the median of the squared error instead of the more traditional mean squared error.

Acknowledgments

We are grateful to Krishnan (Ravi) Shankar and Hippolyte Verdier for insightful discussions and comments. The valuable input of two referees is greatly acknowledged

References

- [1] Aamari, E. (2017). *Vitesse de convergence en inférence géométrique*. Ph. D. thesis, Paris Saclay.
- [2] Aamari, E., J. Kim, F. Chazal, B. Michel, A. Rinaldo, and L. Wasserman (2019). Estimating the reach of a manifold. *Electron. J. Stat.* *13*(1), 1359–1399. [MR3938326](#)
- [3] Aamari, E. and C. Levrard (2019). Nonasymptotic rates for manifold, tangent space and curvature estimation. *Ann. Statist.* *47*(1), 177–204. [MR3909931](#)
- [4] Alexander, S. B. and R. L. Bishop (2006). Gauss equation and injectivity radii for subspaces in spaces of curvature bounded above. *Geometriae Dedicata* *117*(1), 65–84. [MR2231159](#)
- [5] Arias-Castro, E. and T. Le Gouic (2019). Unconstrained and curvature-constrained shortest-path distances and their approximation. *Discrete & Computational Geometry* *62*(1), 1–28. [MR3959920](#)
- [6] Balakrishnan, S., A. Rinaldo, D. Sheehy, A. Singh, and L. Wasserman (2012). Minimax rates for homology inference. In *Artificial Intelligence and Statistics*, pp. 64–72.
- [7] Baldi, P., G. Kerkycharian, D. Marinucci, D. Picard, et al. (2009). Adaptive density estimation for directional data using needlets. *The Annals of Statistics* *37*(6A), 3362–3395. [MR2549563](#)
- [8] Bernstein, M., V. De Silva, J. C. Langford, and J. B. Tenenbaum (2000). Graph approximations to geodesics on embedded manifolds. Technical report, Technical report, Department of Psychology, Stanford University.
- [9] Berry, T. and T. Sauer (2017). Density estimation on manifolds with boundary. *Computational Statistics & Data Analysis* *107*, 1–17. [MR3575055](#)
- [10] Besse, A. (1978). *Manifolds all of whose Geodesics are Closed*, Volume 93. Springer Science & Business Media. [MR0496885](#)
- [11] Boucheron, S., G. Lugosi, and P. Massart (2013). *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press. [MR3185193](#)
- [12] Bowman, A. W. (1984). An alternative method of cross-validation for the smoothing of density estimates. *Biometrika* *71*(2), 353–360. [MR0767163](#)
- [13] Castillo, I., G. Kerkycharian, and D. Picard (2014). Thomas Bayes’ walk on manifolds. *Probability Theory and Related Fields* *158*(3-4), 665–710. [MR3176362](#)
- [14] Chiu, S.-T. (1991). Bandwidth selection for kernel density estimation. *Ann. Statist.* *19*(4), 1883–1905. [MR1135154](#)
- [15] Cleanthous, G., A. Georgiadis, G. Kerkycharian, P. Petrushev, and D. Picard (2018). Kernel and wavelet density estimators on manifolds and more general metric spaces. *arXiv preprint 1805.04682*. [MR4091093](#)
- [16] Devroye, L. and L. Györfi (1985). *Nonparametric density estimation*. Wiley Series in Probability and Mathematical Statistics: Tracts on Probability and Statistics. John Wiley & Sons, Inc., New York. The L_1 view. [MR0780746](#)

- [17] Dubins, L. E. (1957). On curves of minimal length with a constraint on average curvature, and with prescribed initial and terminal positions and tangents. *American Journal of mathematics* 79(3), 497–516. [MR0089457](#)
- [18] Eugeciouglu, Ö. and A. Srinivasan (2000). Efficient nonparametric density estimation on the sphere with applications in fluid mechanics. *SIAM Journal on Scientific Computing* 22(1), 152–176. [MR1769490](#)
- [19] Evans, L. C. and R. F. Gariepy (1992). *Measure theory and fine properties of functions*. Chapman and Hall/CRC. [MR1158660](#)
- [20] Farahmand, A. M., C. Szepesvári, and J.-Y. Audibert (2007). Manifold-adaptive dimension estimation. In *Proceedings of the 24th international conference on Machine learning*, pp. 265–272. ACM.
- [21] Federer, H. (1959). Curvature measures. *Transactions of the American Mathematical Society* 93(3), 418–491. [MR0110078](#)
- [22] Federer, H. (1969). *Geometric measure theory*. Springer. [MR0257325](#)
- [23] Fefferman, C., S. Mitter, and H. Narayanan (2016). Testing the manifold hypothesis. *Journal of the American Mathematical Society* 29(4), 983–1049. [MR3522608](#)
- [24] Gaïffas, S. and G. Lecué (2007). Optimal rates and adaptation in the single-index model using aggregation. *Electron. J. Stat.* 1, 538–573. [MR2369025](#)
- [25] Gallot, S., D. Hulin, and J. Lafontaine (2004). *Riemannian geometry*, Volume 3. Springer. [MR2088027](#)
- [26] Genovese, C., M. Perone-Pacifico, I. Verdinelli, and L. Wasserman (2012). Minimax manifold estimation. *Journal of machine learning research* 13(May), 1263–1291. [MR2930639](#)
- [27] Giné, E. and R. Nickl (2016). *Mathematical foundations of infinite-dimensional statistical models*, Volume 40. Cambridge University Press. [MR3588285](#)
- [28] Goldenshluger, A. and O. Lepski (2008). Universal pointwise selection rule in multivariate function estimation. *Bernoulli* 14(4), 1150–1190. [MR2543590](#)
- [29] Goldenshluger, A. and O. Lepski (2011). Bandwidth selection in kernel density estimation: oracle inequalities and adaptive minimax optimality. *Ann. Statist.* 39(3), 1608–1632. [MR2850214](#)
- [30] Goldenshluger, A. and O. Lepski (2014). On adaptive minimax density estimation on R^d . *Probab. Theory Related Fields* 159(3-4), 479–543. [MR3230001](#)
- [31] Gromov, M. L. and V. A. Rokhlin (1970). Embeddings and immersions in riemannian geometry. *Russian Mathematical Surveys* 25(5), 1. [MR0290390](#)
- [32] Hall, P., G. Watson, and J. Cabrera (1987). Kernel density estimation with spherical data. *Biometrika* 74(4), 751–762. [MR0919843](#)
- [33] Hartman, P. (1951). On geodesic coordinates. *American Journal of Mathematics* 73(4), 949–954. [MR0046087](#)
- [34] Hendriks, H. (1990). Nonparametric estimation of a probability density on a Riemannian manifold using Fourier expansions. *The Annals of Statistics*, 832–849. [MR1056339](#)
- [35] Henry, G. S., A. L. Muñoz, and D. A. Rodriguez (2013). Locally adaptative

- density estimation on Riemannian manifolds. [MR3183696](#)
- [36] Hoffmann, M. and O. Lepski (2002). Random rates in anisotropic regression. *Ann. Statist.* *30*(2), 325–396. With discussions and a rejoinder by the authors. [MR1902892](#)
- [37] Izeddin, I., J. Boulanger, V. Racine, C. Specht, A. Kechkar, D. Nair, A. Triller, D. Choquet, M. Dahan, and J. Sibarita (2012). Wavelet analysis for single molecule localization microscopy. *Optics express* *20*(3), 2081–2095.
- [38] Kerkycharian, G., R. Nickl, and D. Picard (2012). Concentration inequalities and confidence bands for needlet density estimators on compact homogeneous manifolds. *Probability Theory and Related Fields* *153*(1-2), 363–404. [MR2925578](#)
- [39] Kerkycharian, G., T. M. Pham Ngoc, and D. Picard (2011). Localized spherical deconvolution. *Ann. Statist.* *39*(2), 1042–1068. [MR2816347](#)
- [40] Kim, J., A. Rinaldo, and L. Wasserman (2016). Minimax rates for estimating the dimension of a manifold. *arXiv preprint* [1605.01011](#). [MR3918925](#)
- [41] Kim, P. T. and J.-Y. Koo (2002). Optimal spherical deconvolution. *J. Multivariate Anal.* *80*(1), 21–42. [MR1889831](#)
- [42] Kim, P. T., J.-Y. Koo, and Z.-M. Luo (2009). Weyl eigenvalue asymptotics and sharp adaptation on vector bundles. *J. Multivariate Anal.* *100*(9), 1962–1978. [MR2543079](#)
- [43] Klein, T., S. Proppert, and M. Sauer (2014). Eight years of single-molecule localization microscopy. *Histochemistry and cell biology* *141*(6), 561–575.
- [44] Lee, J. M. (2006). *Riemannian manifolds: an introduction to curvature*, Volume 176. Springer Science & Business Media. [MR1468735](#)
- [45] Lepski, O. and N. Serdyukova (2014). Adaptive estimation under single-index constraint in a regression model. *Ann. Statist.* *42*(1), 1–28. [MR3161459](#)
- [46] Lepski, O. V., E. Mammen, V. G. Spokoiny, et al. (1997). Optimal spatial adaptation to inhomogeneous smoothness: an approach based on kernel estimates with variable bandwidth selectors. *The Annals of Statistics* *25*(3), 929–947. [MR1447734](#)
- [47] Lepskii, O. (1992). Asymptotically minimax adaptive estimation. i: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications* *36*(4), 682–697. [MR1147167](#)
- [48] Lepskiĭ, O. V. (1990). A problem of adaptive estimation in Gaussian white noise. *Teor. Veroyatnost. i Primenen.* *35*(3), 459–470. [MR1091202](#)
- [49] Low, M. G. (1992). Nonexistence of an adaptive estimator for the value of an unknown probability density. *Ann. Statist.* *20*(1), 598–602. [MR1150366](#)
- [50] Nikol’skii, S. M. (2012). *Approximation of functions of several variables and imbedding theorems*, Volume 205. Springer Science & Business Media.
- [51] Niyogi, P., S. Smale, and S. Weinberger (2008). Finding the homology of submanifolds with high confidence from random samples. *Discrete & Computational Geometry* *39*(1-3), 419–441. [MR2383768](#)
- [52] Ozakin, A. and A. G. Gray (2009). Submanifold density estimation. In *Advances in Neural Information Processing Systems*, pp. 1375–1382.

- [53] Parzen, E. (1962). On estimation of a probability density function and mode. *Ann. Math. Statist.* 33, 1065–1076. [MR0143282](#)
- [54] Pelletier, B. (2005). Kernel density estimation on Riemannian manifolds. *Statistics & probability letters* 73(3), 297–304. [MR2179289](#)
- [55] Richardson, L. F. (1911). IX. the approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character* 210(459-470), 307–357.
- [56] Rudemo, M. (1982). Empirical choice of histograms and kernel density estimators. *Scand. J. Statist.* 9(2), 65–78. [MR0668683](#)
- [57] Silverman, B. W. (1986). *Density estimation for statistics and data analysis*. Monographs on Statistics and Applied Probability. Chapman & Hall, London. [MR0848134](#)
- [58] Silverman, B. W. (1998). *Density estimation for statistics and data analysis*. Routledge. [MR0848134](#)
- [59] Stone, C. J. (1982). Optimal global rates of convergence for nonparametric regression. *Ann. Statist.* 10(4), 1040–1053. [MR0673642](#)
- [60] Tenenbaum, J. B., V. De Silva, and J. C. Langford (2000). A global geometric framework for nonlinear dimensionality reduction. *science* 290(5500), 2319–2323.
- [61] Triebel, H. (1987). Characterizations of function spaces on a complete Riemannian manifold with bounded geometry. *Mathematische Nachrichten* 130(1), 321–346. [MR0885637](#)
- [62] Tsybakov, A. (2008). *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York. [MR2724359](#)
- [63] Van der Vaart, A. W. (2000). *Asymptotic statistics*, Volume 3. Cambridge university press. [MR1652247](#)
- [64] Watson, G. S. and E. J. Williams (1956). On the construction of significance tests on the circle and the sphere. *Biometrika* 43, 344–352. [MR0082243](#)
- [65] Wu, H.-T. and N. Wu (2020). Strong uniform consistency with rates for kernel density estimators with general kernels on manifolds. *arXiv preprint 2007.06408*.
- [66] Yu, B. (1997). Assouad, Fano, and Le Cam. In *Festschrift for Lucien Le Cam*, pp. 423–435. Springer. [MR1462963](#)

A. Appendix

A.1. Additional results of geometry

We first state a few classical results that we will need in the upcoming proofs. We start with a quantitative bound that links the reach to the curvature of a submanifold. We denote by \mathbb{I} the second fundamental form.

Proposition A.1. Niyogi et al. [51, Prp. 6.1] *Let M be a compact smooth submanifold of \mathbb{R}^D . Then, for any $x \in M$, we have $\|\mathbb{I}_x\|_{\text{op}} \leq 1/\tau_M$.*

Since \mathbb{I}_x is the differential of order two of the mapping \exp_x at the $0 \in T_x M$, Proposition A.1 has several convenient implications. First, it gives a uniform lower bound for the injectivity radii of M as stated in Proposition 2.3. Second, it also yields nice bounds on how well the Euclidean distance on \mathbb{R}^D approximates the Riemannian distance d_M on $M \times M$.

Proposition A.2. Niyogi et al. [51, Prp. 6.3] *For any compact submanifold M of \mathbb{R}^D and any $x, y \in M$ such that $\|x - y\| \leq \tau_M/2$, we have*

$$\|x - y\| \leq d_M(x, y) \leq \tau_M \left(1 - \sqrt{1 - \frac{2\|x - y\|}{\tau_M}} \right).$$

Proposition A.2 allows in turn to compare the volume measure μ_M to the Lebesgue measure on its tangent spaces.

Lemma A.3. *For any d -dimensional compact smooth submanifold M of \mathbb{R}^D , for any $x \in M$ and any $\eta \leq \tau_M/2$, we have*

$$(1 - \eta^2/6\tau_M^2)^d \zeta_d \eta^d \leq \mu_M(B(x, \eta)) \leq \{(1 + (\xi(\eta/\tau_M)\eta)^2/\tau_M^2)\xi(\eta/\tau_M)\}^d \zeta_d \eta^d$$

where $\xi(s) = (1 - \sqrt{1 - 2s})/s$ and ζ_d is the volume of the unit Euclidean ball in \mathbb{R}^d .

Proof. This result already appears in Aamari [1, Lem III.23] but we prove it here to make constants explicit. Let us denote by Leb the Lebesgue measure on $T_x M$. Using Aamari [1, Prp III.22.v], we know that, as long as $\xi(\eta/\tau_M)\eta \leq \tau_M$ (which holds if $\eta \leq \tau_M/2$),

$$\begin{aligned} & (1 - \eta^2/6\tau^2)^d \text{Leb}(B_{T_x M}(0, \eta)) \\ & \leq \mu_M(\exp_x(B_{T_x M}(0, \eta))) \\ & \leq \mu_M(\exp_x(B_{T_x M}(0, \xi(\eta/\tau_M)\eta))) \\ & \leq (1 + (\xi(\eta/\tau_M)\eta)^2/\tau_M^2)^d \text{Leb}(B_{T_x M}(0, \xi(\eta/\tau_M)\eta)). \end{aligned}$$

Thanks to Proposition A.2, if $\eta \leq \tau_M/2$, then $\exp_x(B_{T_x M}(0, \eta)) \subset M \cap B(x, \eta) \subset \exp_x(B_{T_x M}(0, \xi(\eta/\tau_M)\eta))$. These inclusions combined with the last inequalities yield the result. \square

A.2. Proof of Theorem 2.6

We go along a classical line of arguments, thanks to a Bayesian two-point inequality by means of Le Cam's lemma (Yu [66, Lem. 1]), restated here in our context. For two probability measures P_1, P_2 , we write $\text{TV}(P_1, P_2) = \sup_A |P_1(A) - P_2(A)|$ for their variational distance and $H^2(P_1, P_2) = \int (\sqrt{dP_1} - \sqrt{dP_2})^2$ for their (squared) Hellinger distance.

Lemma A.4. (Le Cam) For any $P_1, P_2 \in \Sigma_{\alpha, \beta}^d$, we have,

$$\begin{aligned} & \inf_{\hat{f}} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \\ & \geq \frac{1}{2} |f_{P_1}(x_0) - f_{P_2}(x_0)| (1 - \text{TV}(P_1^{\otimes n}, P_2^{\otimes n})) \\ & \geq \frac{1}{2} |f_{P_1}(x_0) - f_{P_2}(x_0)| \left(1 - \sqrt{2 - 2(1 - H^2(P_1, P_2)/2)^n}\right). \end{aligned} \tag{20}$$

Proof. The proof of (20) can be found in Yu [66, Lem. 1]. It only remains to see that $\text{TV}(P_1^{\otimes n}, P_2^{\otimes n}) \leq \sqrt{2 - 2(1 - H^2(P_1, P_2)/2)^n}$. This comes from classical inequalities on the Hellinger distance, see Tsybakov [62, Lem. 2.3 p.86] and Tsybakov [62, Prp.(i)-(iv) p.83]. \square

Proof of Theorem 2.6. With no loss of generality, we pick $x_0 = 0$. We work in $\mathbb{R}^{d+1} \subset \mathbb{R}^D$, and denote (e_1, \dots, e_{d+1}) the canonical basis of \mathbb{R}^{d+1} . We consider a family of submanifolds $M_\delta \subset \mathbb{R}^{d+1}$ such that

$$M_\delta \cap \{(z, t) \in \mathbb{R}^d \times \mathbb{R} \mid \|z\| \leq 1\} = O(\delta) \cup O(-\delta)$$

where $O(t) = \{(z, t) \mid z \in \mathbb{R}^d, \|z\| \leq 1\}$. We do not give the construction explicitly, but refer instead to Figure 10 for a diagram of such a manifold.




FIG 10. Diagram of a candidate for M_δ .

We endow each M_δ with a density f_δ such that

$$\forall z \in O(\delta), f_\delta(z) = f_{\max} \quad \text{and} \quad \forall z \in O(-\delta), f_\delta(z) = f_{\min}$$

and we denote $Q_\delta = f_\delta d\mu_{M_\delta}$. If f_{\min} is small enough (due to the constraint $\text{vol supp } P \leq 1/f_{\min}$ for any $P \in \Sigma_{\alpha, \beta}^d$) we can always choose M_δ and f_δ so that $|\iota_{M_\delta}|_{\alpha+1} \leq R/2$ and $|f_\delta|_\beta \leq L/2$.

Let now $\Phi : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ be a smooth, positive, radial function with support in $B(0, 1)$ with $\Phi(0) = 1$. Because the exponential map smoothly depends on the

metric, for any $h < 1$, there exists $\delta_h \in (0, h)$ sufficiently small such that the push-forward measures of Q_{δ_h} through the mappings

$$\Psi_h^+(z) = \text{Id} - \delta_h \Phi \left(\frac{z - \delta_h e_{d+1}}{h} \right) e_{d+1} \text{ and } \Psi_h^-(z) = \text{Id} + \delta_h \Phi \left(\frac{z + \delta_h e_{d+1}}{h} \right) e_{d+1}$$

are both in $\Sigma_{\alpha, \beta}^d$. We write $N_h^\pm = \Psi_h^\pm(M_{\delta_h})$, $P_h^\pm = (\Psi_h^\pm)_*(Q_{\delta_h})$ and g_h^\pm for the continuous version of the density $dP_h^\pm/d\mu_{N_h^\pm}$. See Figure 11 for a diagram of N_h^+ and N_h^- .




FIG 11. Diagram of manifolds N_h^+ (Left) and N_h^- (right).

Using Lemma A.4, we obtain

$$\inf_{\hat{f}} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(0) - f_P(0)|^p]^{1/p} \geq \frac{1}{2} |g_h^+(0) - g_h^-(0)| (1 - n \text{TV}(P_h^+, P_h^-)).$$

But now $g_h^+(0) = f_{\delta_h}(\delta_h e_{d+1}) \times |\det d\Psi_h^+(\delta_h e_{d+1})|^{-1} = f_{\max}$ and, likewise, $g_h^-(0) = f_{\min}$. As for the total variation distance, we have that $\text{TV}(P_h^+, P_h^-)$ is equal to

$$\begin{aligned} & P_h^+(\Psi_h^+(hO(\delta_h))) + P_h^+(hO(-\delta_h)) + P_h^-(\Psi_h^-(hO(-\delta_h))) + P_h^-(hO(\delta_h)) \\ & = 2 \text{vol}(hO(\delta_h)) \times (f_{\min} + f_{\max}) = 2\zeta_d h^d (f_{\min} + f_{\max}) \end{aligned}$$

where we recall that ζ_d is the volume of the d -dimensional unit-ball. Putting all the estimates together, we conclude

$$\inf_{\hat{f}} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(0) - f_P(0)|^p]^{1/p} \geq \frac{1}{2} (f_{\max} - f_{\min}) (1 - 2n\zeta_d h^d (f_{\min} + f_{\max})).$$

Letting h goes to 0 yields the result. \square

A.3. Proof of Theorem 3.2

Proof of Theorem 3.2. Suppose without loss of generality that $x = 0$ and consider a smooth submanifold M of $\mathbb{R}^{d+1} \subset \mathbb{R}^D$ that contains the disk $B_{\mathbb{R}^d}(0, 1) \subset \mathbb{R}^d \times \{0_{\mathbb{R}^{D-d}}\}$ with reach greater than τ , see Figure 12 for a diagram of such an M . By smoothness and compactity of M , there exists L_* (depending on τ) such that $M \in \mathcal{C}_{d, \alpha}(\tau, L_*)$. Let P be the uniform probability measure over M , with density $f : x \mapsto 1/\text{vol } M$. We have $P \in \Sigma_{\alpha, \beta}^d$ as long as $L_* \leq L$ and $f_{\min} \leq 1/\text{vol } M \leq f_{\max}$ an assumption we make from now on. For $0 < \delta \leq 1$, let $P_\delta = f_\delta \cdot d\mu_M$ with

$$f_\delta(x) = \begin{cases} f(x) + \delta^\beta G(x/\delta) & \text{if } x \in B_{\mathbb{R}^d}(0, \delta) \\ f(x) & \text{otherwise} \end{cases}$$

with $G : \mathbb{R}^d \rightarrow \mathbb{R}$ a smooth function with support in $B_{\mathbb{R}^d}(0, 1)$ and such that $\int_{\mathbb{R}^d} G(y) dy = 0$. We pick G such that $f_\delta \in \mathcal{F}_\beta$ for small enough δ , depending on τ . Such a G can be chosen to depend on R only.

For δ small enough (depending on τ), we thus have $P_\delta \in \Sigma_{\alpha, \beta}^d$ as well. By Lemma A.4, we infer

$$\inf_{\hat{f}} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x) - f_P(x)|^p]^{1/p} \geq \frac{1}{2} \delta^\beta |G(0)| \left(1 - \sqrt{2 - 2(1 - H^2(P, P_\delta))^n}\right)$$

so that it remains to compute $H^2(P, P_\delta)$. We have the following bound

$$\begin{aligned} H^2(P, P_\delta) &= \int_{B_{\mathbb{R}^d}(0, \delta)} (1 - \sqrt{1 + \text{vol } M \delta^\beta G(x/\delta)})^2 dx \\ &\leq \int_{B_{\mathbb{R}^d}(0, \delta)} (\text{vol } M)^2 \delta^{2\beta} G^2(x/\delta) dx \\ &\leq (C \vee 1) \delta^{2\beta+d} \end{aligned}$$

with $C = \text{vol } M \times \int_{B_{\mathbb{R}^d}(0, 1)} G(z)^2 dz$ depending on τ and R only. Taking $\delta = (1/(C \vee 1)n)^{1/(2\beta+d)}$ we obtain, for large enough n (depending on τ)

$$\begin{aligned} &\inf_{\hat{f}} \sup_{P \in \Sigma_{\alpha, \beta}^d} \mathbb{E}_{P^{\otimes n}} [|\hat{f}(x) - f_P(x)|^p]^{1/p} \\ &\geq \frac{1}{2} ((C \vee 1)n)^{-\beta/(2\beta+d)} \sqrt{2 - 2(1 - 1/n)^n} \geq C_* n^{-\beta/(2\beta+d)}, \end{aligned}$$

with $C_* = (C \vee 1)^{-1/2}$ depending on τ and R . \square

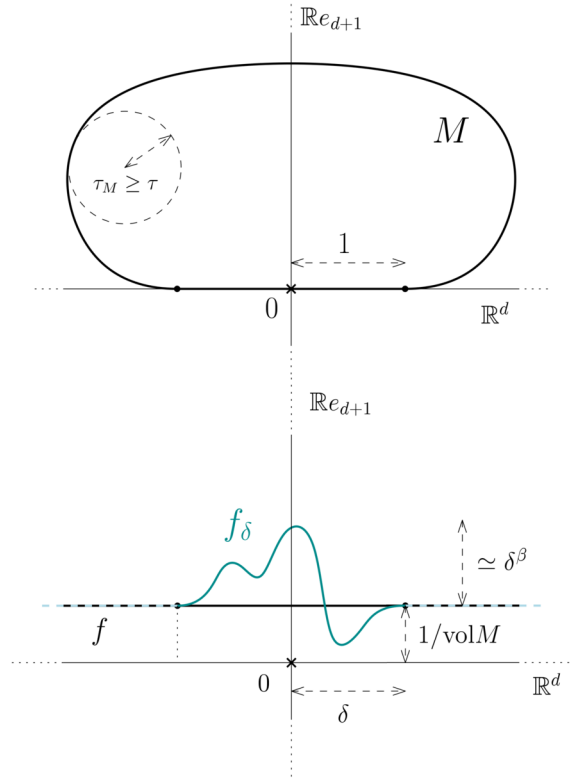


FIG 12. Diagram of a candidate for M (Left) and of the densities f and f_δ around 0 (Right).

A.4. Proofs of Section 3.2

We set $K_h(x) = h^{-d}K(x/h)$ and start with bounding the variance of $K_h(X - x_0)$ when X is distributed according to $P \in \Sigma_{\alpha,\beta}^d$. Let first observe that

$$|K_h(X - x_0)| \leq \frac{\|K\|_\infty}{h^d} \mathbb{1}_{B_D(x_0,h)}(X) \leq \frac{\|K\|_\infty}{h^d} \tag{21}$$

Lemma A.5. For any $P \in \Sigma_{\alpha,\beta}^d$ and for any $h \leq \tau/2$,

$$\text{Var}_P(K_h(X - x_0)) \leq \frac{\omega}{h^d} \quad \text{where } \omega = 4^d \zeta_d \|K\|_\infty^2 b.$$

with ζ_d being the volume of the unit ball in \mathbb{R}^d .

Proof. We have

$$\text{Var}_P(K_h(X - x_0)) \leq \mathbb{E}_P[K_h(X - x_0)^2] \leq \frac{\|K\|_\infty^2}{h^{2d}} P(B(x_0, h)) \leq \frac{4^d \zeta_d f_{\max} \|K\|_\infty^2}{h^d}$$

where we used (21) and Lemma A.3 with $\eta = \tau/2$. □

Using Bernstein inequality Boucheron et al. [11, Thm. 2.10 p.37], for any $P \in \Sigma_{\alpha,\beta}^d$ and any $t > 0$, we infer

$$\mathbb{P} \left(\left| \widehat{\xi}_h(P, x) \right| \geq \sqrt{\frac{2\omega t}{nh^d}} + \frac{\|K\|_\infty t}{nh^d} \right) \leq 2e^{-t}, \tag{22}$$

where \mathbb{P} is a short-hand notation for the distribution $P^{\otimes n}$ of the n -sample X_1, \dots, X_n taken under P . The bound (22) is the main ingredient needed to bound the L_p -norm of the stochastic deviation of \widehat{f}_h .

Proof of Proposition 3.5. We denote by $u_+ = \max\{u, 0\}$ the positive part of a real number u . We start with

$$\mathbb{E}_{P^{\otimes n}} [|\widehat{\xi}_h(P, x_0)|^p] \leq 2^{p-1} \left(\Omega(h)^p + \mathbb{E}_{P^{\otimes n}} [(|\widehat{\xi}_h(P, x_0)| - \Omega(h))_+^p] \right).$$

The first term has the right order. For the second one, we make use of (22) to infer

$$\begin{aligned} & \mathbb{E}_{P^{\otimes n}} [(|\widehat{\xi}_h(P, x_0)| - \Omega(h))_+^p] \\ &= \int_0^\infty \mathbb{P} \left(|\widehat{\xi}_h(P, x_0)| > \Omega(h) + u \right) pu^{p-1} du \\ &= p\Omega(h)^p \int_0^\infty \mathbb{P} \left(|\widehat{\xi}_h(P, x_0)| > \Omega(h)(1 + u) \right) u^{p-1} du \\ &\leq p\Omega(h)^p \left(1 + \int_1^\infty \mathbb{P} \left(|\widehat{\xi}_h(P, x_0)| > \sqrt{\frac{2\omega(1+u)}{nh^d}} + \frac{\|K\|_\infty(1+u)}{nh^d} \right) u^{p-1} du \right) \\ &\leq p\Omega(h)^p \left(1 + \int_1^\infty 2e^{-1-u} u^{p-1} du \right) \\ &\leq p\Omega(h)^p (1 + \Gamma(p)) \end{aligned}$$

which ends the proof. □

The proof of Lemma 3.7 partly relies on the following elementary lemma.

Lemma A.6. *Let $\gamma \geq 0$ be a real number and let $g : \mathbb{R}^m \rightarrow \mathbb{R}$ for $m \in \mathbb{N}^*$ satisfying that $\|g\|_\infty \leq b$ and that the restriction of g to $B(0, r)$ (denoting here the open ball in \mathbb{R}^m) is β -Hölder, meaning that*

$$\forall v, w \in B(0, r), \quad \|d^k g(v) - d^k g(w)\|_{\text{op}} \leq A \|v - w\|^\delta$$

for some $A > 0$ with $k = \lceil \gamma - 1 \rceil$ and $\delta = \gamma - k$. Then there exists a constant C (depending on m, γ, r, b and A , and depending on m and γ when $r = \infty$) such that, for all $1 \leq j \leq k$,

$$\sup_{v \in B(0, r/2)} \|d^j g(v)\|_{\text{op}} \leq C b^{1-j/\gamma} A^{j/\gamma}.$$

Proof. Let $v \in B(0, r/2)$. Since g is γ -Hölder on $B(0, r)$, we know that there exists a function R_v such that, for any z such that $v + z \in B(0, r)$, we have

$$g(v + z) - \sum_{j=0}^k \frac{1}{j!} d^j g(v)[z^{\otimes j}] = R_v(z)$$

with $|R_v(z)| \leq A\|z\|^\beta/k!$. Let $h = (2bk!/A)^{1/\gamma}$, and $z_0 \in \mathbb{R}^m$ be unit-norm. Pick $a_1, \dots, a_k \in (0, 1)$ all distincts and small enough such that $ha_k z_0 \in B(0, r/2)$ for all k (if $r = \infty$, then we can pick the a_i independently from A , b and γ). Introducing the vectors of \mathbb{R}^k

$$X = \left(hdg(v)[z_0], \dots, \frac{h^k}{k!} d^k g(v)[z_0^{\otimes k}] \right) \quad \text{and}$$

$$Y = (g(v + ha_1 z_0) - g(v) - R_v(ha_1 z_0), \dots, g(v + ha_k z_0) - g(v) - R_v(ha_k z_0))$$

we have $Y = VX$ with V being the Vandermonde matrix associated with the real numbers (a_1, \dots, a_k) . The former being invertible, we have $\|X\| \leq \|V^{-1}\|_{\text{op}}\|Y\|$ and thus, for any $1 \leq j \leq k$

$$\left| \frac{h^j}{j!} d^j g(v)[z_0^{\otimes j}] \right| \leq \|V^{-1}\|_{\text{op}} \left(2b + \frac{A}{k!} h^\gamma \right).$$

Substituting the value of h and noticing that the former inequality holds for every unit-norm vector z_0 , we can conclude. \square

Proof of Lemma 3.7. We set $B_h = B(x_0, h)$. Since $\tau/2$ is smaller than the injectivity radius of \exp_{x_0} (see Proposition 2.3) we can write

$$\begin{aligned} f_h(P, x_0) &= \int_{B_h} K_h(p - x_0) f(p) d\mu_M(p) \\ &= \int_{\exp_{x_0}^{-1} B_h} K_h(\exp_{x_0} v - x_0) f(\exp_{x_0} v) \zeta(v) dv \end{aligned} \tag{23}$$

with $\zeta(v) = \sqrt{\det g^{x_0}(v)}$. We set $\gamma = \alpha \wedge \beta$ and $k = \lceil \gamma - 1 \rceil$. Let F denote the map $f \circ \exp_{x_0}$. For h smaller than $\tau/2$, we have $\exp_{x_0}^{-1} B_h \subset B_{T_{x_0}M}(0, 2h) \subset B_{T_{x_0}M}(0, \tau)$ (see Proposition A.2). We can thus write the following expansion, valid for all $v \in \exp_{x_0}^{-1} B_h$ and all $w \in T_{x_0}M$,

$$\exp_{x_0}(v) = x + v + \sum_{j=2}^{k+1} \frac{1}{j!} d^j \exp_{x_0}(0)[v^{\otimes j}] + R_1(v) \quad \text{with} \quad \|R_1(v)\| \leq C_1 \|v\|^{\gamma+1}, \tag{24}$$

$$F(v) = f(x_0) + \sum_{j=1}^k \frac{1}{j!} d^j F(0)[v^{\otimes j}] + R_2(v) \quad \text{with} \quad |R_2(v)| \leq C_2 \|v\|^\gamma, \tag{25}$$

$$K(v+w) = K(v) + \sum_{j=1}^k \frac{1}{j!} d^j K(v)[w^{\otimes j}] + R_3(v,w) \quad \text{with } |R_3(v,w)| \leq C_3 \|w\|^\gamma, \tag{26}$$

with C_1 depending on α, τ and L , C_2 depending on β, τ, f_{\max} and R (see Lemma A.6), and C_3 depending on K . Since now we know that $g_{ij}^{x_0}(v) = \langle d \exp_{x_0}(v)[e_i], d \exp_{x_0}(v)[e_j] \rangle$, we have a similar expansion for the mapping $\zeta(v) = \sqrt{\det g^{x_0}(v)}$

$$\zeta(v) = 1 + \sum_{j=1}^k \frac{1}{j!} d^j \zeta(0)[v^{\otimes j}] + R_4(v) \quad \text{with } |R_4(v)| \leq C_4 \|v\|^\gamma \tag{27}$$

with C_4 depending on α, τ and L . Making the change of variable $v = hw$ in (23), we get

$$f_h(P, x_0) = \sum_{k=0}^k G_j(h, P, x_0) + R_h(P, x_0)$$

with G_j corresponding to the integration of the j -th order terms in the expansion around 0 of the function $w \mapsto K\left(\frac{\exp_{x_0}(hw) - x_0}{h}\right) F(hw)\zeta(hw)$. In particular G_j can be written as a sum of terms of the type

$$I = h^j \int_{\frac{1}{h} \exp_{x_0}^{-1} B_h} d^m K(w)[\phi(w)^{\otimes m}] \psi(w) dw$$

where ψ and ϕ are monomials in w satisfying $m \deg \phi + \deg \psi = j$, with coefficients bounded by constants depending on $\alpha, \tau, L, \beta, f_{\max}$ and R (again, use Lemma A.6 to bound the derivatives). Since now $B_{T_{x_0} M}(0, 1) \subset \frac{1}{h} \exp_{x_0}^{-1} B_h$, and since $d^j K$ is zero outside of $B(0, 1)$, we have that G_j can actually be written $G_j(h, P, x) = h^j G_j(P, x_0)$ with $|G_j(P, x_0)| \leq C$ for some C depending on $K, \alpha, \tau, L, \beta, f_{\max}$ and R . Similar reasoning leads to $R_h(P, x_0) \leq Ch^\gamma$ with C depending again on $K, \alpha, \tau, L, \beta, f_{\max}$ and R . To conclude, it remains to compute $G_0(P, x_0)$. Looking at the zero-th order terms in the expansions (24) to (27), we find that

$$G_0(P, x_0) = \int_{B_{T_{x_0} M}(0,1)} K(w) f(x_0) dw = f(x_0)$$

where we used Assumption 3.6. The proof of Lemma 3.7 is complete. □

Proof of Proposition 3.8. For a positive integer $\ell \geq 1$, let $f_h^{(\ell)}(P, x_0)$ be the mean of the estimator $\widehat{f}_h(x_0)$ computed using $K^{(\ell)}$. Let $\gamma = \alpha \wedge \beta$ and $k = \lceil \gamma - 1 \rceil$. We recursively prove on $1 \leq \ell < \infty$ the following identity

$$\forall h \leq \tau/2, \quad f_h^{(\ell)}(P, x_0) = f(x_0) + \sum_{j=\ell}^k h^j G_j^{(\ell)}(P, x_0) + R_h^{(\ell)}(P, x_0) \tag{28}$$

where $|R_h^{(\ell)}(P, x_0)| \leq C^{(\ell)} h^\gamma$ for some constant $C^{(\ell)}$ depending on $\tau, \ell, L, R, f_{\max}$ and β . The initialisation step $\ell = 1$ has been proven in Lemma 3.7. Let now $1 \leq \ell \leq k$. By linearity of $f_h(P, x_0)$ with respect to K , we have

$$f_h^{(\ell+1)}(P, x_0) = 2f_{2^{-1/\ell}h}^{(\ell)}(P, x_0) - f_h^{(\ell)}(P, x_0).$$

Since $2^{-1/\ell}h \leq h$, we can use our induction hypothesis (28) and find

$$f_h^{(\ell+1)}(P, x_0) = f(x_0) + \sum_{j=\ell}^k (2^{1-j/\ell} - 1) h^j G_j^{(\ell)}(P, x_0) + 2R_{2^{-1/\ell}h}^{(\ell)}(P, x_0) - R_h^{(\ell)}(P, x_0).$$

We conclude noticing that $2^{1-j/\ell} - 1 = 0$ for $j = \ell$, and setting $G_j^{(\ell+1)}(P, x_0) = (2^{1-j/\ell} - 1)G_j^{(\ell)}(P, x_0)$ and $R_h^{(\ell+1)}(P, x_0) = 2R_{2^{-1/\ell}h}^{(\ell)}(P, x_0) - R_h^{(\ell)}(P, x_0)$. The new remainder term verifies

$$|R_h^{(\ell+1)}(P, x_0)| \leq (2^{1-\gamma/\ell} + 1)C^{(\ell)}h^\gamma \leq 3C^{(\ell)}h^\gamma \quad (29)$$

ending the induction by setting $C^{(\ell+1)} = 3C^{(\ell)}$. When $\ell \geq k + 1$, the induction step is trivial. \square

A.5. Proofs of Section 3.3

Proof of Lemma 3.9. Let $\gamma : [0, L_M] \rightarrow M$ be a unit speed parametrisation of M and extend γ to a smooth function on \mathbb{R} by L_M -periodicity. Suppose without loss of generality that $\gamma(0) = x$. For any $t \in \mathbb{R}$, there is a canonical identification between $T_{\gamma(t)}M$ and \mathbb{R} through the map $v \mapsto \langle \dot{\gamma}(t), v \rangle$. With such an identification, we can write that for $s \in \mathbb{R} \simeq T_x M$, $\exp_x(s) = \gamma(s)$ because γ is unit-speed. We thus have $d\exp_x(s)[h] = h\dot{\gamma}(s)$ for any $h \in T_{\gamma(s)}M \simeq \mathbb{R}$. It follows that $\det g^x(s) = \|d\exp_x(s)[1]\|^2 = \|\dot{\gamma}(s)\|^2 = 1$ and this completes the proof. \square

We write $V = (x_0, X_1, \dots, X_n)$ for the vertices of \mathcal{G}_ε and $\widehat{\eta} = \sup_{x \in M} d(x, V)$. For small enough $\widehat{\eta}$ we have that \mathcal{G}_ε is connected, therefore the distance \widehat{d}_ε is well-defined on V . We have in that case a good reverse control of d_M by \widehat{d}_ε , as shown in the next two lemmata.

Lemma A.7. *If $\varepsilon \leq 8\tau$ and $16\widehat{\eta} \leq \varepsilon$, then $\widehat{d}_\varepsilon(x, y) \leq d_M(x, y)$ for any $x, y \in V$.*

Lemma A.8. *If $\varepsilon \leq \tau/2$, then $d_M(x, y) \leq \left(1 + \frac{\pi^2}{48\tau^2}\varepsilon^2\right)\widehat{d}_\varepsilon(x, y)$ for any $x, y \in V$.*

Proof of Lemma A.7. We can take the shortest path in M between x and y as a unit-speed path of the form $\gamma : [0, \ell] \rightarrow \mathbb{R}^D$ with $\ell = d_M(x, y) \leq L_M/2$. We let $\delta = \ell/(4\lfloor \ell/\varepsilon \rfloor)$ and $N = 4\lfloor \ell/\varepsilon \rfloor$. Notice that $\varepsilon/4 \leq \delta \leq \varepsilon/2$. Let us define $p_j = \gamma(j\delta)$, so that $p_0 = x$ and $p_N = y$. Since $\widehat{\eta} \leq \varepsilon/16$, for every $1 \leq j \leq N - 1$, there exists among our vertices V a point denoted by \widehat{p}_j such that $\|p_j - \widehat{p}_j\| \leq \varepsilon/16$. We set $\widehat{t}_j \in [0, L_M]$ for its coordinate, namely $\widehat{p}_j = \gamma(\widehat{t}_j)$.

Let us show first that for $1 \leq j < N$, we have $\widehat{t}_j \in [0, \ell]$. Indeed, thanks to Proposition A.2, since $\varepsilon/16 \leq \tau/2$, we have $|t_j - \widehat{t}_j| \leq 2\|p_j - \widehat{p}_j\| \leq \varepsilon/8$. Since $\delta \geq \varepsilon/4$, we thus have $0 \leq \widehat{t}_1 \leq \dots \leq \widehat{t}_{N-1} \leq \ell$. Furthermore, writing $\widehat{p}_0 = x$ and $\widehat{p}_N = y$, we have

$$\|\widehat{p}_j - \widehat{p}_{j+1}\| \leq \|\widehat{p}_j - p_j\| + \|p_j - p_{j+1}\| + \|p_{j+1} - \widehat{p}_{j+1}\| \leq \varepsilon$$

for any $0 \leq j < \ell$. The sequence $s = (\widehat{p}_0, \dots, \widehat{p}_N)$ is thus a path in \mathcal{G}_ε and so

$$\widehat{d}_\varepsilon(p, q) \leq L_s = \|\widehat{p}_1 - \widehat{p}_0\| + \dots + \|\widehat{p}_N - \widehat{p}_{N-1}\| \leq |\widehat{t}_1 - \widehat{t}_0| + \dots + |\widehat{t}_N - \widehat{t}_{N-1}| = \widehat{t}_N - \widehat{t}_0$$

where we set $\widehat{t}_0 = 0$ and $\widehat{t}_N = \ell = d_M(x, y)$, ending the proof. \square

Proof of Lemma A.8. Following the proof of Arias-Castro and Le Gouic [5, Lem. 5] if there exists $\delta > 0$ such that $\|x - y\| \leq \delta$ implies $d_M(x, y) \leq \pi\tau$ for all $x, y \in M$, then we must have that for any $x, y \in M$ satisfying $\|x - y\| \leq \delta$,

$$d_M(x, y) \leq \left(1 + \frac{\pi^2}{48\tau^2}\|x - y\|^2\right)\|x - y\|.$$

Thanks to Proposition A.2, this must hold for $\delta = \tau/2$. Now let p_0, \dots, p_m be one shortest path in \mathcal{G}_ε between x and y . Since $\varepsilon \leq \tau/2$, we have

$$\begin{aligned} d_M(x, y) &\leq \sum_{j=1}^m d_M(p_j, p_{j-1}) \leq \sum_{j=1}^m \left(1 + \frac{\pi^2}{48\tau^2}\|p_j - p_{j-1}\|^2\right)\|p_j - p_{j-1}\| \\ &\leq \left(1 + \frac{\pi^2}{48\tau^2}\varepsilon^2\right)\widehat{d}_\varepsilon(x, y) \end{aligned}$$

which ends the proof. \square

In view of Lemma A.7 and Lemma A.8, we want to tune ε so that it is the smallest possible and so that $16\widehat{\eta} \leq \varepsilon$ holds with high probability. This is achieved for ε of order $\log n/n$.

Lemma A.9. *Setting $\varepsilon = \frac{32(p+1)\log n}{f_{\min}n}$, for every $n \geq 3$, we have $\mathbb{P}(16\widehat{\eta} \leq \varepsilon) \geq 1 - 1/n^p$.*

Proof. Let $\delta > 0$, and let $N = \lfloor L_M/\delta \rfloor$. We split $[0, L_M]$ into N intervals I_1, \dots, I_N of length L_M/N . We denote \mathcal{A} the event for which each I_j contains at least one coordinate among those of the sample of observations (X_1, \dots, X_n) . On \mathcal{A} , we have $\widehat{\eta} \leq L_M/N \leq 2\delta$. Moreover,

$$\begin{aligned} \mathbb{P}(\mathcal{A}) &= 1 - \mathbb{P}(\exists j, \gamma(I_j) \text{ contains no observation}) \\ &\geq 1 - N \left(1 - \min_{1 \leq j \leq N} P(\gamma(I_j))\right)^n \geq 1 - N \left(1 - \frac{aL_M}{N}\right)^n. \end{aligned}$$

Using that $N \leq L_M/\delta$ and that $L_M \leq 1/a$ we infer

$$\mathbb{P}(\widehat{\eta} \leq 2\delta) \geq 1 - \frac{1}{a\delta}(1 - a\delta)^n \geq 1 - \frac{e^{-a\delta n}}{a\delta}.$$

Setting $\delta = \frac{(p+1)\log n}{an}$ and $\varepsilon = 32\delta$ yields

$$\mathbb{P}(16\widehat{\eta} \leq \varepsilon) \geq 1 - \frac{n}{(p+1)n^{p+1}\log n} \geq 1 - \frac{1}{n^p}$$

as soon as $\log n \geq 1$, *i.e.* for $n \geq 3$. \square

Proof of Proposition 3.10. Recall that we set $K^{1D} = K^{(1,\ell)}$ where $K^{(1,\ell)}$ is defined starting from kernel $\lambda_1^{-1}\Lambda$. Let \mathcal{A} be the event $\{16\widehat{\eta} \leq \varepsilon\}$. By triangle inequality, $\mathbb{E}_{P^{\otimes n}}[|\widehat{f}_h^{1D}(x_0) - f_P(x_0)|^p]^{1/p} \leq \mathcal{R}_{\mathcal{A}} + \mathcal{R}_{\mathcal{A}^c}$, with

$$\mathcal{R}_{\mathcal{A}} = \left(\mathbb{E}_{P^{\otimes n}} [|\widehat{f}_h^{1D}(x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}}] \right)^{1/p}$$

and

$$\mathcal{R}_{\mathcal{A}^c} = \left(\mathbb{E}_{P^{\otimes n}} [|\widehat{f}_h^{1D}(x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}^c}] \right)^{1/p}.$$

On \mathcal{A} , we have, for n large enough (depending on p, f_{\min} and τ) such that $\varepsilon \leq \tau/2$ holds, $|\widehat{d}_\varepsilon(X_i, x_0) - d_M(X_i, x_0)| \leq C_1\varepsilon^2$ with C_1 depending on τ only. This is inferred from Lemmas A.7 and A.8. We deduce that, on this event,

$$\left| \widehat{f}_h^{1D}(x_0) - \widehat{g}_h^{1D}(x_0) \right| \leq \frac{C_1 \|K^{1D'}\|_\infty \varepsilon^2}{h^2} \quad \text{with} \quad \widehat{g}_h^{1D}(x_0) = \frac{1}{n} \sum_{i=1}^n K_h^{1D}(d_M(X_i, x_0)).$$

It follows that

$$\begin{aligned} \mathcal{R}_{\mathcal{A}} &\leq \frac{C_1 \|K^{1D'}\|_\infty \varepsilon^2}{h^2} + \left(\mathbb{E}_{P^{\otimes n}} [|\widehat{g}_h^{1D}(x_0) - f_P(x_0)|^p] \right)^{1/p} \\ &\leq \frac{C_1 \|K^{1D'}\|_\infty \varepsilon^2}{h^2} + \left(\mathbb{E}_{P^{\otimes n}} [|\widehat{\xi}_h^*(P, x_0)|^p] \right)^{1/p} + |\mathcal{B}_h^*(P, x_0)| \end{aligned}$$

with \mathcal{B}_h^* and $\widehat{\xi}_h^*$ denoting the bias and stochastic deviation of estimator $\widehat{g}_h^{1D}(x_0)$. Following the same arguments as in proof of Proposition 3.5, we have

$$\mathbb{E}_{P^{\otimes n}} [|\widehat{\xi}_h^*(P, x_0)|^p]^{1/p} \leq c_p \Omega(h)^p,$$

with c_p depending only on p . For the bias term, as soon as $h \leq \pi\tau$, we have

$$\begin{aligned} \mathcal{B}_h^*(P, x) &= \mathbb{E}_{P^{\otimes n}} [\widehat{g}_h^{1D}(x_0)] - f(x_0) = \int_M K_h^{1D}(d(p, x_0)) f(p) d\mu_M(p) - f(x_0) \\ &= \int_{B_{T_{x_0}M}(0,1)} K^{1D}(\|v\|) (f \circ \exp_{x_0}(hv) - f(x_0)) dv. \end{aligned}$$

Since now $f \circ \exp_x$ is β -Hölder on $B_{T_{x_0}M}(0, \pi\tau)$, we know that all the terms in the development of $\mathcal{B}_h^*(P, x_0)$ up to order $\lceil \beta - 1 \rceil$ cancels. We deduce $|\mathcal{B}_h^*(P, x_0)| \leq C_2 h^\beta$ with C_2 depending on ℓ and R only. For the other term $\mathcal{R}_{\mathcal{A}^c}$, we write $|\widehat{f}_h^{1D}(x_0) - f(x_0)| \leq \frac{\|K^{1D}\|_\infty}{h} + f_{\max}$, so that, according to Lemma A.9,

$$\mathcal{R}_{\mathcal{A}^c} \leq \left(\frac{\|K^{1D}\|_\infty}{h} + f_{\max} \right) \mathbb{P}(\mathcal{A}^c)^{1/p} \leq C_3 \frac{1}{nh}$$

with C_3 depending on ℓ and f_{\max} . Putting all these estimates together yields the result. \square

A.6. Proofs of Section 3.4

Lemma A.10. For any $P \in \Sigma_{\alpha,\beta}^d$, and $\Theta > p$, we have

$$\mathbb{E}_{P^{\otimes n}}[|\widehat{f}(x_0) - f_P(x_0)|^p]^{1/p} \lesssim \Omega(h^*(P, x_0))\lambda(h^*(P, x_0))$$

up to a constant depending on p and Θ , with

$$h^*(P, x_0) = \max \{h \in \mathbb{H} \mid \forall \eta \in \mathbb{H}(h), |f_\eta(P, x_0) - f(x_0)| \leq \frac{1}{2}\Omega(h)\lambda(h)\}.$$

Proof. We fix $P \in \Sigma_{\alpha,\beta}^d$ and write \widehat{h} and h^* for $\widehat{h}(x_0)$ and $h^*(P, x_0)$ respectively. Let $\mathcal{A} = \{\widehat{h} \geq h^*\}$. We can write $\mathbb{E}_{P^{\otimes n}}[|\widehat{f}(x_0) - f_P(x_0)|^p] = \mathcal{R}_{\mathcal{A}} + \mathcal{R}_{\mathcal{A}^c}$, where

$$\mathcal{R}_{\mathcal{A}} = \mathbb{E}_{P^{\otimes n}}[|\widehat{f}(x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}}] \quad \text{and} \quad \mathcal{R}_{\mathcal{A}^c} = \mathbb{E}_{P^{\otimes n}}[|\widehat{f}(x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}^c}].$$

We start with bounding $\mathcal{R}_{\mathcal{A}}$. Firstly,

$$\begin{aligned} \mathcal{R}_{\mathcal{A}} &\leq 3^{p-1} (\mathbb{E}_{P^{\otimes n}}[|\widehat{f}_{\widehat{h}}(x_0) - \widehat{f}_{h^*}(x_0)|^p \mathbb{1}_{\mathcal{A}}] + \mathbb{E}_{P^{\otimes n}}[|\widehat{f}_{h^*}(x_0) - f_{h^*}(P, x_0)|^p \mathbb{1}_{\mathcal{A}}] \\ &\quad + \mathbb{E}_{P^{\otimes n}}[|f_{h^*}(P, x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}}]). \end{aligned}$$

Next, by definition of \widehat{h} and \mathcal{A} , we have

$$|\widehat{f}_{\widehat{h}}(x_0) - \widehat{f}_{h^*}(x_0)| \mathbb{1}_{\mathcal{A}} \leq \psi(\widehat{h}, h^*) \mathbb{1}_{\mathcal{A}} \leq 2\Omega(h^*)\lambda(h^*).$$

By definition of h^* , we also have $|f_{h^*}(P, x_0) - f(x_0)| \leq \frac{1}{2}\Omega(h^*)\lambda(h^*)$. Finally, using Proposition 3.5

$$\mathbb{E}_{P^{\otimes n}}[|f_{h^*}(P, x_0) - f_P(x_0)|^p \mathbb{1}_{\mathcal{A}}] \leq c_p \Omega(h^*)^p \leq c_p (\Omega(h^*)\lambda(h^*))^p$$

holds as well. Putting all three inequalities together yields

$$\mathcal{R}_{\mathcal{A}} \leq C_{\mathcal{A}} (\Omega(h^*)\lambda(h^*))^p \quad \text{with} \quad C_{\mathcal{A}} = 3^{p-1} (2^p + c_p + 2^{-p}).$$

We now turn to $\mathcal{R}_{\mathcal{A}^c}$. Notice that for any $h \in \mathbb{H}(h^*)$, we have

$$|f_h(P, x_0) - f(x_0)| \leq \frac{1}{2}\Omega(h^*)\lambda(h^*) \leq \frac{1}{2}\Omega(h)\lambda(h),$$

hence

$$|\widehat{f}_h(x_0) - f(x_0)| \leq \frac{1}{2}\Omega(h^*)\lambda(h^*) + |\widehat{\xi}_h(P, x_0)|.$$

We can thus write

$$\begin{aligned} \mathcal{R}_{\mathcal{A}^c} &= \sum_{h \in \mathbb{H}(h^*/2)} \mathbb{E}_{P^{\otimes n}}[|\widehat{f}_h(x_0) - f(x_0)|^p \mathbb{1}_{\{\widehat{h}=h\}}] \\ &\leq \sum_{h \in \mathbb{H}(h^*/2)} \mathbb{E}_{P^{\otimes n}} \left[\left(\frac{1}{2}\Omega(h^*)\lambda(h^*) + |\widehat{\xi}_h(P, x_0)| \right)^p \mathbb{1}_{\{\widehat{h}=h\}} \right]. \end{aligned}$$

Now, for any $h \in \mathbb{H}(h^*/2)$, we have

$$\begin{aligned} \{\widehat{h} = h\} &\subset \{\exists \eta \in \mathbb{H}(h), |\widehat{f}_{2h}(x_0) - \widehat{f}_\eta(x_0)| > \psi(2h, \eta)\} \\ &\subset \bigcup_{\eta \in \mathbb{H}(h)} \left\{ \Omega(h^*)\lambda(h^*) + |\widehat{\xi}_{2h,\eta}(P, x_0)| > \psi(2h, \eta) \right\}, \end{aligned}$$

where $\widehat{\xi}_{2h,\eta}(P, x_0) = \widehat{\xi}_{2h}(P, x_0) - \widehat{\xi}_\eta(P, x_0)$, and where we used the triangle inequality and the definition of h^* . Now, we have $\Omega(h^*)\lambda(h^*) \leq \Omega(2h)\lambda(2h)$ since $2h \leq h^*$ and by definition of $\psi(2h, \eta)$, we infer

$$\{\widehat{h} = h\} \subset \bigcup_{\eta \in \mathcal{H}(h)} \left\{ |\widehat{\xi}_{2h,\eta}(P, x_0)| > \Omega(\eta)\lambda(\eta) \right\}$$

so that

$$\mathbb{P}(\widehat{h} = h) \leq \sum_{\eta \in \mathbb{H}(h)} \mathbb{P} \left(|\widehat{\xi}_{2h,\eta}(P, x_0)| > \Omega(\eta)\lambda(\eta) \right) \tag{30}$$

$$\begin{aligned} &\leq \sum_{\eta \in \mathbb{H}(h)} \mathbb{P} \left(|\widehat{\xi}_{2h,\eta}(P, x_0)| > \sqrt{\frac{8\omega\lambda(\eta)}{n\eta^d}} + \frac{2\|K\|_\infty\lambda(\eta)}{n\eta^d} \right) \\ &\leq \sum_{\eta \in \mathbb{H}(h)} 2 \exp(-\lambda(\eta)^2). \end{aligned} \tag{31}$$

For (30) we use the fact that $\lambda(\eta) \geq 1$ and Bernstein's inequality on the random variable $\widehat{\xi}_{2h,\eta}(P, x_0)$ for (31). Noticing now that $\lambda(\eta)^2 \geq d\Theta \log(1/\eta)$, we further obtain

$$\mathbb{P}(\widehat{h} = h) \leq 2h^{\Theta d} \times \sum_{j=0}^{\lfloor \log_2(1/h^-) \rfloor} 2^{-j\Theta d} \leq \frac{2}{1 - 2^{-\Theta d}} h^{\Theta d}.$$

For any $h \in \mathbb{H}(h^*/2)$, we thus get the following bound, using Cauchy-Schwarz inequality

$$\begin{aligned} &\mathbb{E}_{P^{\otimes n}} \left[\left(\frac{1}{2}\Omega(h^*)\lambda(h^*) + |\widehat{\xi}_h(P, x_0)| \right)^p \mathbb{1}_{\{\widehat{h}=h\}} \right] \\ &\leq \mathbb{P}(\widehat{h} = h)^{1/2} E_{P^{\otimes n}} \left[\left(\frac{1}{2}\Omega(h^*)\lambda(h^*) + |\widehat{\xi}_h(P, x_0)| \right)^{2p} \right]^{1/2} \\ &\leq 2^{(2p-1)/2} \sqrt{\frac{2}{1 - 2^{-\Theta d}}} h^{\Theta d/2} \left(2^{-p}\Omega(h^*)^p \lambda(h^*)^p + c_{2p}^{1/2} \Omega(h)^p \right). \end{aligned} \tag{32}$$

We plan to sum over $h \in \mathbb{H}(h^*/2)$ the RHS of (32). Notice first that

$$\sum_{h < h^*} h^{\Theta d/2} \leq (h^*)^{\Theta d/2} (1 - 2^{-\Theta d/2})^{-1}.$$

Moreover, for any $h \geq h^-$, we have $\Omega(h) \leq 2\sqrt{2\omega/(nh^d)}$ by definition of h^- . It follows that

$$\Omega(h^*) \leq \Omega(h) \leq 2\Omega(h^*) \left(\frac{h^*}{h} \right)^{d/2}.$$

for any $h \leq h^*$. This enables us to bound the following sum

$$\begin{aligned} \sum_{h \in \mathbb{H}(h^*/2)} h^{\Theta d/2} \Omega(h)^p &\leq 2^p \Omega(h^*)^p h^{*pd/2} \sum_{h \in \mathbb{H}(h^*/2)} h^{\Theta d/2 - pd/2} \\ &\leq \frac{2^p}{1 - 2^{(p-\Theta)d/2}} \Omega(h^*) h^{*\Theta d/2} \end{aligned}$$

where we used that $\Theta > p$. Putting all these estimates together, using that $h^* \leq 1$ and $\lambda(h^*) \geq 1$, we eventually obtain $\mathcal{R}_{\mathcal{A}^c} \leq C_{\mathcal{A}^c} \Omega(h^*)^p \lambda(h^*)^p$, with

$$C_{\mathcal{A}^c} = \frac{2^p}{\sqrt{1 - 2^{-\Theta d}}} \left(\frac{2^{-p}}{1 - 2^{-\Theta d/2}} + \frac{\sqrt{C_{2p}} 2^p}{1 - 2^{(p-\Theta)d/2}} \right).$$

In conclusion $\mathbb{E}_{P \otimes n} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \leq (C_{\mathcal{A}} + C_{\mathcal{A}^c}) \Omega(h^*)^p \lambda(h^*)^p$ which completes the proof. \square

Proof of Theorem 3.4. Let $P \in \Sigma_{\alpha, \beta}^d$ and let $\bar{h} = (\rho \log n/n)^{1/(2\gamma+d)}$ with $\gamma = \alpha \wedge \beta$ and for some constant ρ to be specified later. By Proposition 3.8 we know that for n large enough (depending on ρ, α, β, d) such that $\bar{h} \leq \tau/2$, we have $|f_\eta(P, x_0) - f(x_0)| \leq C_1 \eta^\gamma$ for all $\eta \leq \bar{h}$ with C_1 depending on $K, \ell, \alpha, \tau, L, \beta, f_{\max}$ and R . Moreover, we also have

$$\frac{2^{-2} \Omega(\bar{h})^2 \lambda(\bar{h})^2}{C_1^2 \bar{h}^{2\gamma}} \geq \frac{d\Theta 2\omega \log(1/\bar{h})}{4C_1^2 n \bar{h}^{2\gamma+d}} = \frac{d\Theta \omega (2\gamma + d)^{-1} \log n - \log \log n - \log \rho}{2C_1^2 \rho \log n}.$$

Thus, picking $\rho = d\Theta \omega (2\gamma + d)^{-1} / (2C_1^2)$ yields $C_1 \bar{h}^\gamma \leq \frac{1}{2} \Omega(\bar{h}) \lambda(\bar{h})$ for n large enough (depending on ρ), and therefore $\bar{h} \leq h^*(P, x)$. By Lemma A.10 this implies

$$\mathbb{E}_{P \otimes n} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \leq C_2 \Omega(\bar{h}) \lambda(\bar{h})$$

where C_2 depends on p and Θ . But using that both $\bar{h} \geq h^-$ and $\lambda(\bar{h})^2 = d\Theta \log(1/\bar{h})$ for n large enough (depending on ρ, d, K and Θ), we also obtain

$$\begin{aligned} \Omega(\bar{h})^2 \lambda(\bar{h})^2 &\leq \frac{8\omega d \Theta \log(1/\bar{h})}{n \bar{h}^d} \\ &= \frac{8\omega d \Theta (2\gamma + d)^{-1} \log n - \log \log n - \log \omega}{\rho \log n} \bar{h}^{2\gamma} \leq 16C_1^2 \bar{h}^{2\gamma}. \end{aligned}$$

This last estimate yields

$$\mathbb{E}_{P \otimes n} [|\hat{f}(x_0) - f_P(x_0)|^p]^{1/p} \leq \left(4C_1 C_2 \rho^{\gamma/(2\gamma+d)} \right) \left(\frac{\log n}{n} \right)^{\gamma/(2\gamma+d)}$$

for n large enough depending on $\rho, \alpha, \beta, d, K$ and Θ , which completes the proof. \square

A.7. Proofs of Section 3.5

Proof of Proposition 3.13. By the triangle inequality, for any $P \in \Sigma_{\alpha,\beta}^d$, we write

$$\begin{aligned} \mathbb{E}_{P^{\otimes n}} [|\widehat{f}_{\widehat{d}}(\widehat{d}; x_0) - f_P(x_0)|^p]^{1/p} &\leq \left(\mathbb{E}_{P^{\otimes n}} \left[\left| \widehat{f}_{\widehat{d}}(\widehat{d}; x_0) - f_P(x_0) \right|^p \mathbb{1}_{\{\widehat{d}=d\}} \right] \right)^{1/p} \\ &\quad + \left(\mathbb{E}_{P^{\otimes n}} \left[\left| \widehat{f}_{\widehat{d}}(\widehat{d}; x_0) - f_P(x_0) \right|^p \mathbb{1}_{\{\widehat{d} \neq d\}} \right] \right)^{1/p}. \end{aligned}$$

The first term in the right-hand side has the right order thanks to Theorem 3.4. For the second one, using that $|f(x_0)| \leq f_{\max}$ and

$$|\widehat{f}_{\widehat{d}}(\widehat{d}, x_0)| \leq \sup_{1 \leq d < D} \frac{\|K^{(\ell)}(d, \cdot)\|_{\infty}}{(h_{\widehat{d}})^d} \lesssim n$$

up to a constant that depend on D, K and ℓ , we infer

$$\left(\mathbb{E}_P \left[\left| \widehat{f}_{\widehat{d}}(\widehat{d}; x_0) - f(x_0) \right|^p \mathbb{1}_{\{\widehat{d} \neq d\}} \right] \right)^{1/p} \lesssim \mathbb{P}(\widehat{d} \neq d)^{1/p} \times n.$$

Finally, since \widehat{d} satisfies Assumption 3.12, we have

$$\mathbb{E}_{P^{\otimes n}} [|\widehat{f}_{\widehat{d}}(\widehat{d}; x_0) - f_P(x_0)|^p]^{1/p} \lesssim \left(\frac{\log n}{n} \right)^{\alpha \wedge \beta / (2\alpha \wedge \beta + d)} + n^{-1/2}$$

for n large enough depending on $p, \Theta, K, \ell, \alpha, \tau, L, \beta, f_{\max}, f_{\min}$ and R , so that the result indeed holds up to a constant depending on the same parameters and D . \square

Proof of Proposition 3.15. Let $P \in \Sigma_{\alpha,\beta}^d$ and $\eta > 0$. Assume that $\widehat{P}_{\eta} > 0$. We have

$$\begin{aligned} |\widehat{\delta}_{\eta} - d| &\leq |\log_2 \widehat{P}_{2\eta} - \log_2 P_{2\eta}| + |\log_2 \widehat{P}_{\eta} - \log_2 P_{\eta}| + |\log_2 P_{2\eta} - \log_2 P_{\eta} - d| \\ &\leq \frac{1}{\log 2} \left(\frac{|\widehat{P}_{2\eta} - P_{2\eta}|}{\widehat{P}_{2\eta} \wedge P_{2\eta}} + \frac{|\widehat{P}_{\eta} - P_{\eta}|}{\widehat{P}_{\eta} \wedge P_{\eta}} \right) + |\log_2 (P_{2\eta}/(2^d P_{\eta}))| \end{aligned}$$

We first consider the determinist term. For $\eta \leq \tau/2$, we have, writing $r_{\eta} = \xi(\eta/\tau)\eta$ and using Lemma A.3,

$$L_{2\eta}(1 - \eta^2/6\tau^2)(2\eta)^2 \zeta_d \leq P_{2\eta} \leq U_{2\eta}(1 + r_{2\eta}^2/\tau^2)r_{2\eta}^d \zeta_d$$

and

$$L_{\eta}(1 - \eta^2/6\tau^2)\eta^d \zeta_d \leq P_{\eta} \leq U_{\eta}(1 + r_{\eta}^2/\tau^2)r_{\eta}^d \zeta_d,$$

where $L_{\eta} = \inf_{M \cap B(x_0, \eta)} f$ and $U_{\eta} = \sup_{M \cap B(x_0, \eta)} f$. Using again Lemma A.3, we have that for $\eta \leq \tau/2$, $M \cap B(x_0, \eta) \subset \exp_{x_0} B_{T_{x_0} M}(0, 2\eta)$, and, since $2\eta \leq \pi\tau/2$ and that $f \in \mathcal{F}_{\beta}$, we know, using Lemma A.6, that there exists $R_1 > 0$

(depending on β, τ, f_{\max} and R) such that $f(x) - R_1\eta \leq L_\eta \leq U_\eta \leq f(x) + R_1\eta$. If $\eta < \tau/4$, the same bounds apply for 2η and we thus obtain

$$(f(x) - R_1 2\eta)(1 - \eta^2/6\tau^2)(2\eta)^d \zeta_d \leq P_{2\eta} \leq (f(x) + R_1 2\eta)(1 + r_{2\eta}^2/\tau^2)r_{2\eta}^d \zeta_d$$

and

$$(f(x) - R_1\eta)(1 - \eta^2/6\tau^2)\eta^d \zeta_d \leq P_\eta \leq (f(x) + R_1\eta)(1 + r_\eta^2/\tau^2)r_\eta^d \zeta_d. \tag{33}$$

Using these two inequalities, and the fact that $r_\eta/\eta \rightarrow 1$ as $\eta \rightarrow 0$, we find that $|P_{2\eta}/(2^d P_\eta) - 1| \lesssim \eta$ up to a constant that depends on R_1, τ and f_{\min} , for η small enough (depending on R_1, τ and f_{\min} as well). For the other terms, a simple use of Hoeffding's inequality yields for any $\eta, \varepsilon > 0$,

$$\mathbb{P}\left(|\widehat{P}_\eta - P_\eta| > \varepsilon\right) \leq 2 \exp(-2n\varepsilon^2).$$

On the event $\mathcal{A}_\eta = \{|\widehat{P}_\eta - P_\eta| \leq \varepsilon\}$, we have moreover $\widehat{P}_\eta \wedge P_\eta \geq P_\eta - \varepsilon$. Setting $\varepsilon = \eta^{d+1}$, and using (33), we see that $P_\eta - \varepsilon \gtrsim \eta^d$ for η small enough (depending on R_1, τ and f_{\min}). Thus, on the event $\mathcal{A}_\eta \cap \mathcal{A}_{2\eta}$, with probability at least $1 - 4 \exp(-2n\eta^{2d+2})$, we derive

$$|\widehat{\delta}_\eta - d| \lesssim \eta + \frac{\varepsilon}{P_\eta - \varepsilon} + \frac{\varepsilon}{P_{2\eta} - \varepsilon} \lesssim \eta, \tag{34}$$

for η small enough (depending on R_1, τ and f_{\min}), up to a constant that depends on R_1, τ and f_{\min} . Now setting $\eta = n^{-1/(2D+2)}$, we have $\widehat{d} = \widehat{\delta}_\eta = d$ on the event $\mathcal{A}_\eta \cap \mathcal{A}_{2\eta}$ as soon as n is large enough so that the LHS of (34) is strictly smaller than $1/2$, ending the proof. \square