

CLASSIFICATION IN GENERAL FINITE DIMENSIONAL SPACES WITH THE k -NEAREST NEIGHBOR RULE

BY SÉBASTIEN GADAT*, THIERRY KLEIN[†] AND CLÉMENT MARTEAU[‡]

*Toulouse School of Economics**, *Université Paul Sabatier[†]*
and Université Lyon 1[‡]

Given an n -sample of random vectors $(X_i, Y_i)_{1 \leq i \leq n}$ whose joint law is unknown, the long-standing problem of supervised classification aims to *optimally* predict the label Y of a given new observation X . In this context, the k -nearest neighbor rule is a popular flexible and intuitive method in non-parametric situations. Even if this algorithm is commonly used in the machine learning and statistics communities, less is known about its prediction ability in general finite dimensional spaces, especially when the support of the density of the observations is \mathbb{R}^d . This paper is devoted to the study of the statistical properties of the k -nearest neighbor rule in various situations. In particular, attention is paid to the marginal law of X , as well as the smoothness and margin properties of the *regression function* $\eta(X) = \mathbb{E}[Y|X]$. We identify two necessary and sufficient conditions to obtain uniform consistency rates of classification and derive sharp estimates in the case of the k -nearest neighbor rule. Some numerical experiments are proposed at the end of the paper to help illustrate the discussion.

1. Introduction. The supervised classification model has been at the core of numerous contributions to statistical literature in recent years. It continues to provide interesting theoretical and practical problems. Supervised classification aims to predict a feature $Y \in \mathcal{M}$ when a variable of interest $X \in \mathbb{R}^d$ is observed, the set \mathcal{M} being finite ($\mathcal{M} = \{0, 1\}$ for a binary classification). In order to provide a prediction of the label Y of X , it is assumed that a training set $\mathcal{S}_n = (X_i, Y_i)_{1 \leq i \leq n}$ is at our disposal and makes it possible to provide a prediction *via* an inference on the joint law (X, Y) . Many methods have been proposed over the years and we refer to [5] for an extended introduction. These methods can be divided in (at least) three families:

- *Pure entropy considerations and Empirical Risk Minimization (ERM)*. It selects a classifier that yields the ERM among a family of candidates (see, e.g., [2, 25] and [24] for a detailed description). In an almost similar context, aggregation schemes (see, e.g., [14] and [21]) have been shown to be adaptive to margin and complexity.

Received November 2014; revised September 2015.

MSC2010 subject classifications. Primary 62G05, 62F15; secondary 62G20.

Key words and phrases. Supervised classification rates, k -nearest neighbor, plug-in rules.

- *Geometric interpretation or information theory.* The SVM (see [29, 32] among others) aims to maximize the margin of the classification rule. CART is another intuitive standard method, improved by a bagging procedure in [1] and [6], referred to as Random Forest.
- *Plug-in rules.* The main idea is to mimic the Bayes classifier using a plug-in rule after a preliminary estimation of the *regression* function. We refer to [18] for a general overview (see also [2]).

In this general overview, the k -nearest neighbor rule (k -N.N. for short) belongs to the last two classes: it is a plug-in classifier with a simple geometrical interpretation. It has attracted a great deal of attention for the past few decades, from the seminal work of [13]. In particular, a famous positive result of [30] is its *universal* consistency (see also [11]), meaning that the k -N.N. can be carefully tuned to be consistent under mild assumptions on the model. Recently, this algorithm has received further attention and is still at the core of several studies: [8] identifies the importance of the Besicovitch assumption, [19] is concerned with two notions of the sample structure, [28] describes an improvement of the algorithm that allows to deal with smoother regression functions, while [9] studies the consistency and the rate of convergence of the algorithm in abstract metric spaces.

We investigate here the achievable consistency rate of the k -N.N. under various conditions. Most of the results obtained for penalized ERM, SVM or plug-in classifiers are based on complexity considerations (entropy or VC dimension). In this paper, we mainly use the asymptotic behavior of the small ball probabilities instead (see [23] and the references therein), which is a dual quantity of the entropy (see [22]) and we deal with the intricate situation of not bounded away from zero densities (and noncompactly supported measures). For this purpose, we handle *smoothness* and *minimal mass* assumptions that will provide a pertinent estimation of the function η . We also consider an additional *margin* parameter α : [25] proved that *fast rates* (faster than \sqrt{n}^{-1}) can be obtained by exploiting the law of (X, Y) near $\{\eta = 1/2\}$. Our contributions can be gathered in 3 different axes.

Rate for bounded from below densities. We state the optimality of the k -N.N. Φ_n and show that mild assumptions implies the minimax consistency rate

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \leq C n^{-(1+\alpha)/(2+d)},$$

where α denotes the margin parameter, d the dimension of the problem, $\mathcal{R}(\Phi)$ the miss-classification error of Φ and Φ^* the Bayes classifier.¹

¹This result has also been established in the recent work of [28].

Rate for general densities. We study the behavior of Φ_n when the marginal density μ of X is not bounded from below on its support. Such an improvement is of first importance since it corresponds to many practical situations. To do this, we add an assumption on the *tail* of μ and prove that generically

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \leq Cn^{-(1+\alpha)/(2+\alpha+d)},$$

as soon as the bandwidth k involved in the classifier is allowed to depend on the spatial position of X . The *tail* assumption on μ involved in this result describes the behavior of μ near the set $\{\mu = 0\}$.

Lower bounds. Finally, we derive some lower bounds for the supervised classification problem, which extends the results obtained in [2]. We prove that our tail assumption is unavoidable to ensure uniform consistency rates for classification in a noncompact case. We then see how these upper and lower bounds are linked and show that a very unfavorable situation of classification occurs when the regression function η oscillates in the tail of the distribution μ : it is even impossible in these situations to obtain uniform consistency rates, and thus elucidate two open questions in [7].

The paper is organized as follows. Section 2 reminds some basics of the k -N.N. rule. Section 3 is devoted to the bounded from below case. We then extend our study to the general (typically noncompact) case in Section 4. Proofs of the upper bound are included in Appendix. Technical results and the proofs of the lower bounds can be found in the supplementary material [15].

We use the following notation throughout the paper. The term $\mathbb{P}_{X,Y}$ denotes the distribution of the couple (X, Y) and \mathbb{P}_X is the marginal distribution of X , which possesses a density μ w.r.t. the Lebesgue measure. Similarly, we set $\mathbb{P}_{\otimes^n} = \prod_{i=1}^n \mathbb{P}_{(X_i, Y_i)}$ and $\mathbb{P} = \mathbb{P}_{(X,Y)} \times \mathbb{P}_{\otimes^n}$. Naturally, $\mathbb{E}[\cdot]$, $\mathbb{E}_X[\cdot]$ and $\mathbb{E}_{\otimes^n}[\cdot]$ correspond hereafter to the expectations w.r.t. the measures \mathbb{P} , \mathbb{P}_X and \mathbb{P}_{\otimes^n} respectively. Finally, given two real sequences $(a_n)_{n \in \mathbb{N}}$ and $(b_n)_{n \in \mathbb{N}}$, we write $a_n \lesssim b_n$ (resp., $a_n \sim b_n$) if a constant $C \geq 1$ exists such that $a_n \leq Cb_n$ (resp., $C^{-1}b_n \leq a_n \leq Cb_n$) for all $n \in \mathbb{N}$.

2. Statistical setting and k -nearest neighbor classifier.

2.1. Statistical classification problem. We observe an i.i.d. sample $\mathcal{S}_n := (X_i, Y_i)_{i=1, \dots, n} \in \Omega \times \{0, 1\}$, whose distribution is $\mathbb{P}_{X,Y}$ and where $\Omega = \text{Supp}(\mu)$ is an open set of \mathbb{R}^d . For a new incoming observation X , our goal is to predict its corresponding label Y . To do this, we use a classifier Φ that provides a decision rule for this problem: Φ is a measurable mapping from \mathbb{R}^d to $\{0, 1\}$, whose corresponding miss-classification error is then defined as

$$\mathcal{R}(\Phi) = \mathbb{P}(\Phi(X) \neq Y).$$

In practice, the most interesting classifiers are those associated with the smallest error. It is well known that the Bayes classifier Φ^* defined as

$$(2.1) \quad \Phi^*(X) = \mathbf{1}_{\{\eta(X) > 1/2\}} \quad \text{where } \eta(x) := \mathbb{E}[Y|X = x] \quad \forall x \in \Omega,$$

minimizes the misclassification error, that is,

$$\mathcal{R}(\Phi^*) \leq \mathcal{R}(\Phi) \quad \forall \Phi : \mathbb{R}^d \longrightarrow \{0, 1\}.$$

Unfortunately, Φ^* is not available since the regression function η depends on the underlying distribution of (X, Y) . The Bayes classifier can be considered as an *oracle* that provides a benchmark error and the main challenge is to construct a classifier Φ that possesses a small excess risk given by

$$\mathcal{R}(\Phi) - \mathcal{R}(\Phi^*).$$

We study the properties of the excess risk of a given classifier Φ_n through the *minimax* paradigm. Given a set \mathcal{F} of possible distributions F for (X, Y) , we define

$$\delta_n(\mathcal{F}) := \inf_{\Phi} \sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi) - \mathcal{R}(\Phi^*)],$$

where the infimum in the above formula is taken over all \mathcal{S}_n measurable classifiers. A classifier Φ_n is then said to be minimax over the set \mathcal{F} if

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \leq C \delta_n(\mathcal{F}),$$

for a positive constant C . The considered set \mathcal{F} will be detailed later on and will depend on some smoothness, margin and minimal mass assumptions.

2.2. The k -nearest neighbor rule. The k -N.N. rule is one of the simplest and widespread classification procedures. Suppose that the state space is \mathbb{R}^d with a reference norm $\|\cdot\|$. Given any sample \mathcal{S}_n and for any $x \in \mathbb{R}^d$, we build the reordered sample $(X_{(j)}(x), Y_{(j)}(x))_{1 \leq j \leq n}$ w.r.t. the distances $\|X_i - x\|$:

$$\|X_{(1)}(x) - x\| \leq \|X_{(2)}(x) - x\| \leq \dots \leq \|X_{(n)}(x) - x\|.$$

In this context, $X_{(m)}(x)$ is the m -nearest neighbor of x w.r.t. the distance $\|\cdot\|$ and $Y_{(m)}(x)$ its corresponding label. Given any integer k in \mathbb{N} , the principle of the k -nearest neighbor algorithm is to construct a decision rule based on the k -nearest neighbors of the input X : the \mathcal{S}_n -measurable classifier $\Phi_{n,k}$ is

$$(2.2) \quad \Phi_{n,k}(X) = \begin{cases} 1, & \text{if } \frac{1}{k} \sum_{j=1}^k Y_{(j)}(X) > \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases}$$

For all $x \in \Omega$, the term $\frac{1}{k} \sum_{j=1}^k Y_{(j)}(x)$ appears to be an estimator of the regression function $\eta(x)$. In particular, we can write the classifier $\Phi_{n,k}$ as

$$(2.3) \quad \Phi_{n,k}(X) = \mathbf{1}_{\{\hat{\eta}_n(X) > 1/2\}} \quad \text{where } \hat{\eta}_n(x) = \frac{1}{k} \sum_{j=1}^k Y_{(j)}(x) \quad \forall x \in \Omega.$$

Hence, the k -N.N. is a plug-in classifier, *that is*, a preliminary estimator of the function η is plugged in our decision rule. It is worth noting that the integer k is a *regularization parameter*. The k -N.N. is quite robust since universal consistency is obtained as soon as $k_n \rightarrow +\infty$ and $k_n/n \rightarrow 0$, but a careful tuning of the number of neighbors k_n is needed to obtain an acceptable rate of convergence. Indeed, if k is too small, the classifier $\Phi_{n,k}$ only uses a small amount of the neighbors of X , inducing a large variance of the classification process. On the other hand, large values of k generate some bias into the decision rule since we use observations that may be far away from the input X .

For this purpose, we introduce some baseline assumptions into the following section that will make it possible to characterize an optimal value for k_n .

2.3. Baseline assumptions. It is well known that no reliable prediction can be made in a distribution-free setting (see [18]). We restrict the class of possible distributions of (X, Y) below.

Since the k -nearest neighbor rule is a plug-in classification rule, we expect to take advantage of some smoothness properties of η in order to improve the classification process. In fact, when η is smooth, the respective values of $\eta(x_1)$ and $\eta(x_2)$ are comparable for close enough x_1, x_2 . In other words, we can infer the sign of $\eta(x) - \frac{1}{2}$ from those of the neighbors of x .

ASSUMPTION A1 (Smoothness). The regression function η belongs to the Hölder class of parameter 1 with a *radius* L , which is denoted $\mathcal{C}^{1,0}(\Omega, L)$ and corresponds to the set of functions such that

$$\forall (x_1, x_2) \in \Omega^2 \quad |\eta(x_1) - \eta(x_2)| \leq L|x_1 - x_2|.$$

REMARK 2.1. It would be tempting to consider some more general smoothness classes for the regression function η . Nevertheless, the standard k -nearest neighbor algorithm does not make it possible to use smoothness indexes greater than 1. An alternative procedure has been proposed in [28]: the idea is then to balance the $(Y_{(j)})_{j=1,\dots,k}$ with a suitable monotonous weighting sequence. However, this modification complicates the statistical analysis and may alter the ideas developed below. We therefore chose to fix the smoothness of η to 1 [i.e., restrict our study to $\mathcal{C}^{1,0}(\Omega, L)$].

Our second assumption was introduced by [31] in the binary supervised classification model (see [25] in a *smooth discriminant analysis* setting).

ASSUMPTION A2 (Margin assumption). Some constants $\alpha > 0$ and $C > 0$ exist such that

$$\mathbb{P}_X(0 < |\eta(X) - \frac{1}{2}| < \varepsilon) \leq C\varepsilon^\alpha \quad \forall \varepsilon > 0.$$

In such a case, we write $(\mu, \eta) \in \mathbf{M}_\alpha$.

The Bayes classifier depends on the sign of $\eta(X) - 1/2$. Intuitively, it would be easier to mimic the behavior of this classifier when the mass around the set $\{\eta = 1/2\}$ is small. On the other hand, the decision process may be more complicated when $\eta(X)$ is close to $1/2$ with a large probability. Quantifying this closeness is the purpose of this margin assumption.

For the sake of convenience, we use the set $\mathcal{F}_{L,\alpha}$ throughout the paper, which contains distributions that satisfy both Assumptions A1 and A2, namely

$$\mathcal{F}_{L,\alpha} := \left\{ \mathbb{P}_{(X,Y)} : \mathcal{L}(X) \sim \mu, \mathcal{L}(Y|X) \sim \mathcal{B}(\eta(X)) \text{ such that} \right. \\ \left. \eta \in \mathcal{C}^{1,0}(\Omega, L) \text{ and } (\mu, \eta) \in \mathbf{M}_\alpha \right\}.$$

We now turn to our last assumption that involves the marginal distribution of the variable X .

2.4. *Minimal mass assumption.* In the sequel, this type of hypothesis will play a very important role.

ASSUMPTION A3 (Strong Minimal Mass Assumption). There exists $\kappa > 0$ such that the marginal density μ of X satisfies $\mu \in \mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ where

$$\mathfrak{M}_{\text{mma}}(\Omega, \kappa) := \left\{ \mathbb{P}_X : \mathcal{L}(X) \sim \mu \mid \exists \delta_0 > 0, \forall \delta \leq \delta_0, \right. \\ \left. \forall x \in \Omega : \mathbb{P}_X(X \in B(x, \delta)) \geq \kappa \mu(x) \delta^d \right\}.$$

This assumption guarantees that \mathbb{P}_X possesses a minimal amount of mass on each ball $B(x, \delta)$; this lower bound being balanced by the level of the density on x . In some sense, distributions in $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ will make it possible to obtain reliable predictions of the regression function η according to its Lipschitz property.

The strong minimal mass Assumption A3 is much stronger than the so-called Besicovitch assumption that is quite popular in the statistical literature (see, e.g., [10] for a version of the Besicovitch assumption used for pointwise consistency or [8] for a general discussion on this hypothesis in finite or infinite dimension). It is worth pointing out that the Besicovitch assumption introduced in [8] states that η satisfies:

$$(2.4) \quad \forall \varepsilon > 0 \quad \lim_{\delta \rightarrow 0} \mu \left\{ x : \frac{1}{\mu(B(x, \delta))} \int_{B(x, \delta)} |\eta(z) - \eta(x)| d\mu(z) > \varepsilon \right\} = 0.$$

As shown in [8], (2.4) is always true in finite dimensional space.

We can also remark that if η is L -Lipschitz (Assumption A1), we have

$$\forall x \in \Omega \quad \int_{B(x, \delta)} |\eta(z) - \eta(x)| \mu(z) dz \leq L \int_{B(x, \delta)} |x - z| \mu(z) dz \\ \leq L \delta \mu(B(x, \delta)),$$

which implies that the set involved in (2.4) is empty as soon as $L\delta \leq \varepsilon$. Hence, (2.4) is true when $\eta \in \mathcal{C}^{1,0}(\Omega, L)$, whatever the dimension of Ω is. In fact, we can also show that the continuity of η implies (2.4) we refer to Section 3 in [8] for more details.

We will see that Assumption A3 is necessary to obtain quantitative estimates for any finite dimensional classification problem.

In a slightly different setting, our Assumption A3 is used in [2] when the density μ is lower bounded on its (compact) support, which is assumed to possess some geometrical properties [(c_0, r_0) regularity]. This setting is at the core of the study presented in Section 3 below. Assumption A3 also recalls the notion of standard sets used in [27] for the estimation of compact support sets. More generally, the following examples present some standard distributions that satisfy Assumption A3.

EXAMPLE 2.1.

- In \mathbb{R}^d , it is not difficult to check that Gaussian measures with non-degenerated covariance matrices satisfy $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$. As a simple example, consider a standard Gaussian law $\mu \sim \mathcal{N}(0, 1)$. For any $x \in \mathbb{R}$ and $\delta > 0$, if x belongs to a compact set K , then a constant C_K exists such that $(2\pi)^{-1/2} \int_{x-\delta}^{x+\delta} e^{-t^2/2} dt \geq C_K e^{-x^2/2} \delta$. Now, if $x \rightarrow +\infty$, we can check that

$$(2\pi)^{-1/2} \int_{x-\delta}^{x+\delta} e^{-t^2/2} dt \sim (2\pi)^{-1/2} e^{-x^2/2} \left[\frac{e^{x\delta}}{x-\delta} - \frac{e^{-x\delta}}{x+\delta} \right] e^{-\delta^2/2}.$$

The bracket above is always greater than δ when $(x\delta)^{-1} = O(1)$. Now, if $\delta = o(1/x)$, a simple Taylor expansion yields

$$(2\pi)^{-1/2} \int_{x-\delta}^{x+\delta} e^{-t^2/2} dt \sim \mu(x) \frac{1 + 2x\delta}{x} \gtrsim \mu(x)\delta.$$

- The same computations are still possible for symmetric Laplace distributions ($e^t \int_{t-\delta}^{t+\delta} e^{-x} dx = [e^\delta - e^{-\delta}] \sim 2\delta$ when δ is small. Thus, any Laplace distributions belongs to $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$. In a same way, when μ is a standard Cauchy distribution, we can check that

$$\begin{aligned} \int_{x-\delta}^{x+\delta} \frac{dt}{1+t^2} &= \frac{1}{1+x^2} \int_{-\delta}^{\delta} \frac{1}{1+h(2x+h)/(1+x^2)} dh \\ &\sim \frac{1}{1+x^2} \left[2\delta - \frac{2}{3} \frac{\delta^3}{1+x^2} + +8 \frac{\delta^3 x^2}{(1+x^2)^2} o(\delta^3) \right] \\ &\gtrsim \frac{\delta}{1+x^2}. \end{aligned}$$

In the case of compactly supported distribution, it is intuitive to see that $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ is related to the regularity of the boundary of the support. For example, consider the uniform law on the $\ell_{1/2}$ ball of \mathbb{R}^2 given by $\Omega := \{(x_1, x_2) \in \mathbb{R}^2 \mid \sqrt{|x_1|} + \sqrt{|x_2|} \leq 1\}$. In this case, we can check that $\mathbb{P}_X(X \in B((1, 0), \delta)) = 4\delta^3/3$ for $\delta \leq 1$, which is much smaller than δ^2 when $\delta \rightarrow 0$ and this distribution does not belong to $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$. This last distribution exemplifies that Assumption A3 is not only an assumption on the boundedness of the derivatives of μ , but is also strongly related to the geometrical shape of its support.

Other typical distributions that do not satisfy the strong minimal mass Assumption A3 possess some important oscillations in their tails (when the density μ is close to 0). In such a setting, the alternative set $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ defined as follows, may be considered:

$$\begin{aligned} \widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa) := & \{ \mathbb{P}_X : \mathcal{L}(X) \sim \mu \mid \exists(\rho, C) \in]0; +\infty[^2, \exists \delta_0 > 0, \forall \delta \leq \delta_0, \\ & \forall x \in \Omega : \mu(x) \geq e^{-C\delta^{-\rho}} \implies \mu(B(x, \delta)) \geq \kappa \mu(x) \delta^d \}. \end{aligned}$$

The interest of the weaker $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ compared to $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ is that the statistical abilities of the k -nearest neighbor rule are still the same with $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ or $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$. Moreover, an analytic criterion that ensures $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ can be found (see Proposition 4.1). This is not the case for the *uniform* assumption $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ (it is indeed more difficult to ensure the lower bound on the global set Ω).

Although all the subsequent results may be established for a weaker version of the minimal mass assumption [based on the set $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$], we will restrict ourselves to its strong formulation (Assumption A3). In Section 3, we prove that the k -nearest neighbor rule is optimal in the minimax sense provided that the margin and smoothness assumptions hold, with a marginal density of the variable X bounded away from 0 and a suitable choice of k . In Section 4, we will see that $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ is not yet sufficient to derive some uniform consistency rates for classifiers with non-compactly supported densities and a last additional hypothesis is needed.

3. Bounded away from zero densities.

3.1. *Minimax consistency of the k -nearest neighbor rule.* In this section, we are interested in the special case of a marginal density μ bounded from below by a strictly positive constant μ_- . In this context, we can state an upper bound on the consistency rate of the k -nearest neighbor rule.

THEOREM 3.1. *Assume that Assumptions A1–A3 hold. The k -nearest neighbor classifier Φ_{n, k_n} with $k_n = \lfloor n^{2/(2+d)} \rfloor$ satisfies*

$$\sup_{\mathbb{P}_{X, Y} \in \mathcal{F}_{L, \alpha} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}} [\mathcal{R}(\Phi_{n, k_n}) - \mathcal{R}(\Phi^*)] \lesssim n^{-(1+\alpha)/(2+d)},$$

where $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}$ denotes the subset of densities of $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ that are bounded from below by μ_- .

Theorem 3.1 establishes a consistency rate of the k -nearest neighbor rule over $\mathcal{F}_{L,\alpha} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}$. A detailed proof of is presented in Appendix. Implicitly, we restrict our analysis to compactly supported observations, these assumptions being at the core of several statistical analyses (see, e.g., [5, 18, 25] or [19] among others). It is worth pointing out that this setting falls into the framework considered in [2].

DEFINITION 3.1 (Strong Density Assumption (SDA), [2]). *The marginal distribution of the variable X satisfies the strong density assumption if:*

- it admits a density μ w.r.t. the Lebesgue measure of \mathbb{R}^d ;
- the density μ satisfies

$$\mu_- \leq \mu(x) \leq \mu_+ \quad \forall x \in \text{Supp}(\mu)$$

for some constants $(\mu_-, \mu_+) \in]0, +\infty[^2$;

- the support of μ is (c_0, r_0) -regular, namely

$$\lambda[\text{Supp}(\mu) \cap B(x, r)] \geq c_0 \lambda[B(x, r)] \quad \forall r \leq r_0,$$

for some positive constants c_0 and r_0 .

As soon as the marginal density is bounded from below by a strictly positive constant, then both SDA and strong minimal mass Assumption A3 are equivalent, as stated in the following proposition.

PROPOSITION 3.1. *For bounded away from zero density, the SDA is equivalent to the strong minimal mass assumption.*

PROOF OF PROPOSITION 3.1. As soon as the support of μ is (c_0, r_0) -regular and the density is lower bounded by $\mu_- > 0$, then SDA implies a minimal mass type assumption:

$$\begin{aligned} \forall \delta \leq r_0 \quad \mu(B(x, \delta)) &= \int_{B(x, \delta)} \mu(z) dz \geq \mu_- \times \lambda[B(x, \delta) \cap \text{Supp}(\mu)] \\ &\geq c_0 \gamma_d \mu_- \delta^d. \end{aligned}$$

Conversely, we can also check the fact that the strong minimal mass Assumption A3 implies the SDA [including the (c_0, r_0) -regularity of μ]. Indeed, since for any x and $\delta \leq \delta_0$

$$1 \geq \int_{B(x, \delta)} \mu(y) dy \geq \kappa \mu(x) \delta^d,$$

then the density μ is upper bounded and we obtain that

$$\int_{B(x,\delta)} \mu(y) dy \leq \|\mu\|_\infty \lambda[\text{Supp}(\mu) \cap B(x, r)].$$

We therefore obtain

$$\lambda[\text{Supp}(\mu) \cap B(x, r)] \geq \kappa \frac{\mu(x)}{\|\mu\|_\infty} \delta^d \geq \kappa \frac{\mu_-}{\|\mu\|_\infty} \delta^d.$$

This completes the proof of this proposition. \square

It is possible to link the constants (c_0, r_0) and κ involved in SDA with $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}$, we have omitted their relationships here for the sake of simplicity. Minimax rates of excess risk under the SDA are established in Theorem 3.5 of [2]. A consequence of Proposition 3.1 is that the same lower bound is still valid with $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}$. Consequently, under Assumptions A1–A3 and if a constant $\mu_- > 0$ exists such that $\mu(x) > \mu_-$ on Ω , then we obtain that

$$\inf_{\Phi} \sup_{\mathbb{P}_{X,Y} \in \mathcal{F}_{L,\alpha} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}} [\mathcal{R}(\Phi) - \mathcal{R}(\Phi^*)] \gtrsim n^{-(1+\alpha)/(2+d)}.$$

This inequality and Theorem 3.1 show that the k -N.N. rule achieves the minimax rate of convergence in the particular case where the density μ is lower bounded on its (compact) support. As already discussed in [25] or [2], the higher the margin index α is, the smaller the excess risk will be. On the other hand, the performance deteriorates as the dimension of the considered problem increases. The lower bound obtained by [2] is based on an adaptation of standard tools from non-parametric statistics (Assouad’s lemma). This proof is of primary importance for next lower bound results. It is recalled in Section 1 of the supplementary material [15] for the sake of convenience.

4. Non-compact finite dimensional case.

4.1. *The tail assumption.* Results of the previous section are designed for the problem of supervised binary classification with compactly supported inputs and lower bounded densities. Such an assumption is an important prior on the problem that may be improper in several practical settings. Various situations involve Gaussian, Laplace, Cauchy or Pareto distributions on the observations, and both the compactness and the boundedness away from zero assumptions may seem to be very unrealistic. This is even more problematic when dealing with functional classification with a Gaussian White Noise model (GWN). In such a case, observations are described through an infinite sequence of Gaussian random variables, and the SDA or $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)^{\mu_-}$ are far from being well tailored for this situation (see [23] for a discussion and further references).

This section is dedicated to a more general case of binary supervised classification problems where the marginal density μ of X is no longer assumed to be lower

bounded on its support. The main problem related to such a setting is that we have to predict labels in places where few (or even no) observations are available in the training set. In order to address this problem, we take the following assumption.

ASSUMPTION A4 (Tail Assumption). A function ψ that satisfies $\psi(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$ and that increases in a neighborhood of 0 exists such that

$$\mathbb{P}_{(X,Y)} \in \mathcal{P}_{\mathcal{T},\psi} := \{ \mathbb{P}_X : \exists \varepsilon_0 \in \mathbb{R}_+^* : \forall \varepsilon < \varepsilon_0, \mathbb{P}_X(\{\mu < \varepsilon\}) \leq \psi(\varepsilon) \},$$

where $\mathcal{P}_{\mathcal{T},\text{Id}}$ corresponds to the particular case where $\psi = \text{Id}$.

The aim of this tail assumption is to ensure that the set where μ is small has a small mass. We use the notation \mathcal{T} because of the interpretation on the *tail* of μ , but $\mathcal{P}_{\mathcal{T},\psi}$ is not just an assumption on the tail of the μ . It is, in fact, an assumption on the behavior of μ near the set $\{\mu = 0\}$. We provide below some examples of marginal distributions that satisfy this tail requirement. In Section 4.2 below, we prove that the tail Assumption A4 is unavoidable in this setting. In Section 4.3, we investigate the performances of the k -nearest neighbor rule in this context.

EXAMPLE 4.1. The following are several families of densities in $\mathcal{P}_{\mathcal{T},\psi}$.

- Laplace distributions obviously satisfy $\mathcal{P}_{\mathcal{T},\text{Id}}$, and a straightforward integration by parts shows that Gamma distributions $\Gamma(k, \theta)$ satisfy $\mathcal{P}_{\mathcal{T},\psi}$ with $\psi(\varepsilon) = \varepsilon \log(\varepsilon^{-1})^{k-1}$ (the term around $x = 0$ is on the order of $\varepsilon^{k/(k-1)}$, and thus negligible compared to the term around $+\infty$).
- An immediate computation shows that the family of Pareto distributions of parameters (x_0, k) satisfies $\mathcal{P}_{\mathcal{T},\psi}$ where $\psi(\varepsilon) = \varepsilon^{k/(k+1)}$, regardless of the value of x_0 .
- The family of Cauchy distributions satisfies $\mathcal{P}_{\mathcal{T},\psi}$ with $\psi(\varepsilon) = \sqrt{\varepsilon}$.
- Univariate Gaussian laws γ_{m,σ^2} with mean m and variance σ^2 satisfy

$$\gamma_{m,\sigma^2}(x) \leq \varepsilon \iff |x - m| \geq t_{\sigma,\varepsilon} := \sqrt{2}\sigma \sqrt{\log\left(\frac{1}{\varepsilon}\right) + \log\left(\frac{1}{\sigma\sqrt{2\pi}}\right)},$$

and a standard result on the size of Gaussian tails (see [3]) yields

$$\gamma_{m,\sigma^2}(\gamma_{m,\sigma^2} \leq \varepsilon) = \frac{\varepsilon}{t_{\sigma,\varepsilon}} \left[1 - \frac{1}{t_{\sigma,\varepsilon}^2} + \frac{1.3}{t_{\sigma,\varepsilon}^4} - \dots \right] \lesssim \frac{\varepsilon}{\sqrt{\log(1/\varepsilon)}}.$$

Hence, univariate Gaussian laws satisfy $\mathcal{P}_{\mathcal{T},\psi}$ with $\psi(\varepsilon) = \varepsilon \log(\varepsilon^{-1})^{-1/2}$.

- **Multivariate Gaussian laws.** If \mathbf{m} is any real vector of \mathbb{R}^d and Σ^2 a covariance matrix whose spectrum is $\lambda_1 \geq \dots \geq \lambda_d \geq 0$:

$$\gamma_{\mathbf{m},\Sigma^2}(\gamma_{\mathbf{m},\Sigma^2} \leq \varepsilon) = \gamma_{\mathbf{0},\Sigma^2}(\gamma_{\mathbf{0},\Sigma^2} \leq \varepsilon) \lesssim \gamma_{\mathbf{0},\Sigma^2} \left(\|X\| \geq \sqrt{2\lambda_1 \log\left(\frac{1}{\varepsilon}\right)} \right).$$

Careful inspection of Theorem 1 of [20] now yields

$$\gamma_{\mathbf{0}, \Sigma^2} \left(\|X\| \geq \sqrt{2\lambda_1 \log\left(\frac{1}{\varepsilon}\right)} \right) \sim C_{\Sigma^2} \log\left(\frac{1}{\varepsilon}\right)^{r/2-1} \varepsilon,$$

where C_{Σ^2} is a constant that only depends on the spectrum of Σ^2 and r is the multiplicity of the eigenvalue λ_1 . In particular, $\gamma_{\mathbf{m}, \Sigma^2}$ satisfy $\mathcal{P}_{\mathcal{T}, \psi}$ where $\psi(\varepsilon) = C_{\Sigma^2} \varepsilon \log(\varepsilon^{-1})^{r/2-1}$.

4.2. *Non-uniform consistency results.* We first justify the introduction of the sets $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ and $\mathcal{P}_{\mathcal{T}, \psi}$ and discuss their influences regarding feasibility to derive lower bounds and even uniform consistency of any estimator. To do this, we first state that the minimal mass Assumption A3 is necessary to obtain uniformly consistent classification rules. Second, we assert that the tail Assumption A4 is also unavoidable.

THEOREM 4.1. *Assume that the law $\mathbb{P}_{X,Y}$ belongs to $\mathcal{F}_{L,\alpha}$, then:*

(i) *No classification rule can be uniformly consistent if Assumptions A1–A3 hold and not A4. For any integer n , any discrimination rule Φ_n and for any $\varepsilon < 4^{-\alpha}$, a distribution $\mathbb{P}_{(X,Y)}^{(n)}$ in $\mathcal{F}_{L,\alpha} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ exists such that*

$$\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \geq \varepsilon.$$

(ii) *No classification rule can be uniformly consistent if Assumptions A1, A2, A4 hold and not A3. For any integer n , any discrimination rule Φ_n and for any $\varepsilon < 4^{-\alpha}$, a distribution $\mathbb{P}_{(X,Y)}^{(n)}$ in $\mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T}, \text{Id}}$ exists such that*

$$\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \geq \varepsilon.$$

We briefly comment on the two points raised by this last theorem. First and foremost, Theorem 4.1 does not contradict the seminal result of [30] that establishes the universal consistency of the k -N.N. as soon as $k_n \rightarrow +\infty$ with $k_n/n \rightarrow 0$. This positive result corresponds to the consistency of the k -N.N. (without any rate) as soon as the distribution $\mathbb{P}_{(X,Y)}$ is fixed while the number of observations n is growing to infinity. Theorem 4.1 states that both Assumptions A3 and A4 are necessary to derive uniform consistency rates for a family of distributions. In particular, Theorem 4.1 is obtained via the construction of a set of distributions $\mathbb{P}_{(X,Y)}^{(n)}$ on the entries that depend on n . To sum up, results obtained in [30] and Theorem 4.1 illustrate essentially the difference between universal convergence and uniform convergence over a class of distribution.

The first result (i) asserts that even if the minimal mass Assumption A3 holds for the underlying density on X , it is not possible to expect a uniform consistency result over the entire class of non-compactly considered densities. In some

sense, the support of the variable X seems to be too large to obtain reliable predictions with any classifiers without additional assumptions. As discussed above, the tail Assumption A4 may make it possible to counterbalance this curse of support effect (see next section). Such statistical damage has also been observed for the estimation of densities that are supported on the real line instead of being compactly supported, even though such dramatic consequences are not shown here. We refer to [26] and the references therein for a more detailed description.

The second result (ii) states that the strong minimal mass Assumption A3 cannot be skipped for uniform consistency rates and no compactly supported densities. This is in line with the former studies of [17] and [11]. In particular, Lemma 2.2 of [11] takes advantage of some of the positive consequences of this type of assumption. Our proof relies on the construction of a sample size dependent law on (X, Y) that violates Assumption A3 but that keeps the regression function η in our smoothness class $\mathcal{F}_{L,\alpha}$. This is a major difference with former counterexamples built in [12] where the nonuniform consistency is obtained with a family of non-smooth regression functions η . In our study, we also obtained a family of smooth regression functions for which such phenomena occur. Even in this case, it is still possible to keep the excess risk larger than a fixed positive constant (independent on n) for any classifier Φ_n .

Finally, these uniform inconsistency results always occur when building a network of regression functions η that oscillate around the value $1/2$ at the neighborhood of the set $\{\mu = 0\}$. In a sense, Theorem 4.1 contributes to the understanding of one of the open questions put forth in [7] on the behavior of the k -nearest neighbor rule when η is oscillating about $1/2$ in the tail.

4.3. *Minimax rates of convergence.* In the meantime, when both Assumptions A2, A3 and A4 hold, we are able to precisely describe the corresponding minimax rate of convergence.

4.3.1. *Minimax lower bound.*

THEOREM 4.2. *Assume that Assumptions A1–A4 hold. Then*

$$\inf_{\Phi_n} \sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa) \cap \mathcal{P}_{\mathcal{T}, \text{Id}}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \gtrsim n^{-(1+\alpha)/(2+\alpha+d)}.$$

For the sake of convenience, we briefly outline the proof of Theorem 3.5 borrowed from [2] in Section 1 of the supplementary material [15]. It is then adapted to our new set of assumptions.

Theorem 4.5 below provides some lower bounds for different tails of distributions (through the function ψ). It should be noted that we recover the known rate of compactly supported densities with the so-called *Mild Density Assumption* of [2] in the particular case $\psi = \text{Id}$. This implies that in the noncompact case, the rate cannot be improved compared to the compact setting, even with an additional tail assumption.

4.3.2. *An upper bound for the k -nearest neighbor rule.* When the density is no longer bounded away from 0, the integer k_n will be chosen in order to counter-balance the vanishing probability of the small balls in the tail of the distributions. For example, when $\psi = \text{Id}$, we show that a suitable choice of the integer k_n is

$$k_n := \lfloor n^{2/(3+\alpha+d)} \rfloor,$$

which appears to be quite different from the one in the previous section.

THEOREM 4.3. *Assume that Assumptions A1–A3 hold and if the tail Assumption A4 is driven by $\psi = \text{Id}$, the choice $k_n := \lfloor n^{2/(3+\alpha+d)} \rfloor$ yields*

$$\sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\text{Id}} \cap \mathcal{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_{n,k_n}) - \mathcal{R}(\Phi^*)] \lesssim n^{-(1+\alpha)/(3+\alpha+d)}.$$

The proof of Theorem 4.3 is provided in Section A.1. The above results indicate that the price to pay for the classification from entries in compact sets to arbitrary large sets of \mathbb{R}^d is translated by the degradation from $n^{-(1+\alpha)/(2+d)}$ to at least $n^{-(1+\alpha)/(2+\alpha+d)}$ [see, e.g., Theorem 4.2 when $\psi(\varepsilon) \sim \varepsilon$]. Our upper bound for the k -nearest neighbor rule does not exactly match this lower bound since we obtain $n^{-(1+\alpha)/(3+\alpha+d)}$ in a similar situation. At this step, obtaining the appropriate minimax rate requires slight changes inside the construction of the k -nearest neighbor rule. This is the purpose of the next paragraph.

4.3.3. *Minimax upper bound for an optimal k -nearest neighbor rule.* The upper bound proposed in the theorem can be improved if we change the way in which the regularization parameter k_n is constructed. We use a k -nearest neighbor algorithm with a number of neighbors that depends on the position of the observation x according to the value of the density $\mu(x)$. More formally, we define for all $j \in \mathbb{N}$

$$\Omega_{n,0} := \{x \in \mathbb{R}^d : \mu(x) \geq n^{-\alpha/(2+\alpha+d)}\},$$

and

$$\Omega_{n,j} = \left\{ x \in \mathbb{R}^d : \frac{n^{-\alpha/(2+\alpha+d)}}{2^{j+1}} \leq \mu(x) < \frac{n^{-\alpha/(2+\alpha+d)}}{2^j} \right\}.$$

Setting $k_{n,0} = \lfloor n^{2/(2+\alpha+d)} \log(n) \rfloor$, we then use for all $j \in \mathbb{N}$

$$(4.1) \quad k_n(x) = \lfloor k_{n,0} 2^{-2j/(2+d)} \rfloor \vee 1 \quad \text{when } x \in \Omega_{n,j}.$$

According to (4.1), the number of neighbors involved in the decision process depends on the spatial position of the input X . In some sense, this position is linked to the tail. The statistical performances of the corresponding k -nearest neighbor classifier is displayed below. Such a construction of this sequence of “slices” may be interpreted as a spatial adaptive bandwidth selection. This bandwidth is smaller at points $x \in \mathbb{R}^d$ such that $\mu(x)$ is small. In a sense, this idea is close to the one introduced in [16] that provides a similar slicing procedure to obtain an adaptive minimax density estimation on \mathbb{R}^d .

THEOREM 4.4. *Assume that Assumptions A1–A3 hold and that the tail Assumption A4 is driven by $\psi = \text{Id}$. Then, if Φ_{n,k_n}^* is the classifier associated with (4.1), we have*

$$\begin{aligned} & \sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\text{Id}} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_{n,k_n}^*) - \mathcal{R}(\Phi^*)] \\ & \lesssim n^{-(1+\alpha)/(2+\alpha+d)} (\log n)^{1/2+1/d}. \end{aligned}$$

We stress that the upper bound obtained in Theorem 4.4 nearly matches the lower bound proposed in Theorem 4.2, up to a log-term. This log-term can be removed by the use of additional technicalities that are omitted in our proof. Hence, Theorems 4.4 and 4.2 make it possible to identify the exact minimax rate of classification when the tail assumption is driven by $\psi = \text{Id}$, that is,

$$\inf_{\Phi} \sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\text{Id}} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_{n,k_n}^*) - \mathcal{R}(\Phi^*)] \sim n^{-(1+\alpha)/(2+\alpha+d)}.$$

REMARK 4.1. Let us briefly compare our results with those obtained in the recent contribution [9] on k -N.N. in general metric spaces. In the particular case of $\Omega = \mathbb{R}^d$ with compactly supported measure, Theorem 3.1 yields an excess risk of the order $n^{-(1/d)(1+\alpha)/(2/d+1)}$, which is also the result stated in Theorem 4 of [9]. Both results are obtained with an additional smoothness assumption (see Assumption A1 in our framework and the (α, L) smoothness assumption of [9] related to the average value of η on small balls of radius r).

Now, when the measure is not compactly supported, [9] describes a geometric set $\mathcal{E}_{n,k}$ that involves all points for which a ball of relative mass k/n (in the sense of μ) leads to an average value of η around $1/2 \pm k^{-1/2}$. We refer to their Section 2.4 for a complete definition of $\mathcal{E}_{n,k}$. Taking together Theorems 5 and 6 of [9] shows that the performance of the k -N.N. classifier is almost proportional to $\mu(\mathcal{E}_{n,k})$ but the excess risk is not explicit.

In our work, $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ and our tail Assumption A4 may be seen as one way to obtain a quantitative description of the set \mathcal{E}_{n,k_n} and then derive explicit in n consistency rates. Finally, our lower bound stated in Theorem 4.1 is obtained *via* a construction of a set of pairs (μ, η) that keep a sufficiently large mass on the associated geometric set $\mathcal{E}_{n,k}$. It is important to note that this lower bound applies not only for the k -N.N. but for *any* classifier.

4.3.4. Generalizations. We propose several extensions of our previous results (lower and upper bounds) for more general tails of distribution. We also propose to enlighten the *minimal mass assumption* $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$.

Effect of the tail: From $\mathcal{P}_{\mathcal{T}, \text{Id}}$ to $\mathcal{P}_{\mathcal{T}, \psi}$.

THEOREM 4.5. *Assume that Assumptions A1–A4 hold. For any tail \mathcal{T} parameterized by a function ψ , we obtain the following results:*

(i) Lower bound: the minimax classification rate satisfies

$$\inf_{\Phi_n} \sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\psi} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \gtrsim \varepsilon_{n,\alpha,d}^{1+\alpha},$$

where $\varepsilon_{n,\alpha,d}$ satisfies the balance

$$(4.2) \quad n^{-1} = \{\varepsilon_{n,\alpha,d}\}^{2+d} \times \psi^{-1}(\{\varepsilon_{n,\alpha,d}\}^\alpha).$$

(ii) Upper bound: the k -nearest neighbor rule satisfies

$$\sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\psi} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_{n,k_n}) - \mathcal{R}(\Phi^*)] \leq C v_{n,\alpha,d}^{1+\alpha}$$

with $k_n = v_{n,\alpha,d}^{-2}$ where $v_{n,\alpha,d}$ fulfills the balance

$$(4.3) \quad n^{-1} = \psi^{-1}(\{v_{n,\alpha,d}\}^{1+\alpha}) \{v_{n,\alpha,d}\}^{2+d}.$$

It would also be possible to propose some generalizations using the sliced k -nearest neighbor rule presented in Sections 4.3.2 and 4.3.3 for tails driven by a general function ψ , even if we do not include this additional result for the purpose of clarity.

Meeting the minimal mass assumption $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$. We now obtain similar rates when using the weaker assumption $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ instead of $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$: the lower bounds of $\mu(B(x, \delta))$ are only useful for some points x such that $\mu(x)$ is large enough. We can state the next corollary.

COROLLARY 4.1. *Assume that Assumptions A1, A2, A4 hold and $P_{(X,Y)} \in \widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$, then*

$$\sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\psi} \cap \widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_{n,k_n}) - \mathcal{R}(\Phi^*)] \lesssim v_{n,\alpha,d}^{1+\alpha},$$

with $k_n = v_{n,\alpha,d}^{-2}$ where $v_{n,\alpha,d}$ satisfies the balance

$$n^{-1} = \psi^{-1}(\{v_{n,\alpha,d}\}^{1+\alpha}) \{v_{n,\alpha,d}\}^{2+d}.$$

The condition $\mathfrak{M}_{\text{mma}}(\Omega, \kappa)$ cannot be easily described through an analytical condition because of its uniform nature over Ω . In contrast, $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ is more tractable in view of the criterion given by the next result (Proposition 4.1). Using a log-density model, we write the density μ as

$$\mu(x) = e^{-\varphi(x)} \quad \forall x \in \mathbb{R}^d.$$

PROPOSITION 4.1. *Let $\varphi \in \mathcal{C}^1(\Omega)$ and assume that a real number $a > 0$ exists such that*

$$\lim_{x:\mu(x)\rightarrow 0} \frac{\|\nabla\varphi(x)\|}{\varphi(x)^a} = 0,$$

then a suitable κ can be found such that $\mu = e^{-\varphi} \in \widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$.

PROOF. For any $\delta > 0$, we compute a lower bound of

$$\mathbb{P}_X(B(x, \delta)) = \int_{B(x, \delta)} e^{-\varphi(z)} dz.$$

The Jensen inequality applied to the normalized Lebesgue measure over $B(x, t)$, which is denoted $\bar{d}z$, yields

$$(4.4) \quad \int_{B(x, \delta)} e^{-\varphi(z)} dz \geq \frac{\pi^{d/2}\delta^d}{\Gamma(d/2 + 1)} \exp\left(-\varphi(x) + \int_{B(x, \delta)} [\varphi(z) - \varphi(x)] \bar{d}z\right).$$

A first-order Taylor expansion leads to

$$\begin{aligned} \int_{B(x, \delta)} [\varphi(z) - \varphi(x)] \bar{d}z &\leq \sup_{z \in B(x, \delta)} \|\nabla\varphi(z)\| \int_{B(x, \delta)} \|z - x\| \bar{d}z \\ &\leq \delta \sup_{z \in B(x, \delta)} \|\nabla\varphi(z)\|. \end{aligned}$$

Now, our assumption on φ implies that a large enough C_a exists such that

$$\|\nabla\varphi(z)\| \leq C_a(1 + \varphi(z)^a).$$

Thus, the lower bound (4.4) becomes

$$\int_{B(x, \delta)} e^{-\varphi(z)} dz \geq \frac{\pi^{d/2}\delta^d}{\Gamma(d/2 + 1)} e^{-\varphi(x)} e^{-C_a\delta(1 + \sup_{z \in B(x, \delta)} \varphi^a(z))}.$$

It is now sufficient to consider points x such that $\varphi \leq \delta^{-1/a}$ (equivalent to $\mu \geq e^{-\delta^{-1/a}}$) to obtain a meaningful lower bound. Hence, $\widetilde{\mathfrak{M}}_{\text{mma}}(\Omega, \kappa)$ is satisfied choosing

$$\rho = 1/a \quad \text{and} \quad \kappa = \frac{\pi^{d/2}}{2\Gamma(d/2 + 1)} e^{-C_a}. \quad \square$$

4.4. *Practical settings on typical examples.* The aim of this section is to illustrate the results obtained above. We first describe a location model for which we can derive explicit upper and lower bounds in several different cases. We then propose a small numerical study in order to enhance the discussion regarding the importance of the tail assumption and we conclude by drawing a comparison between the standard k -nearest neighbor and sliced k -nearest neighbor rules.

TABLE 1
Convergence rates for locations models with several tail sizes

Law	Tail ψ	Margin	$k_n \sim n^\beta$	Upper bound
Gauss	$\psi(\varepsilon) \propto \varepsilon \log(1/\varepsilon)^{r/2-1}$	$\alpha = 1$	$\beta = 2/(4 + d)$	$n^{-2/(4+d)} \log(n)^{\beta(r)}$
Laplace	$\psi(\varepsilon) \propto \varepsilon$	$\alpha = 1$	$\beta = 2/(4 + d)$	$n^{-2/(4+d)}$
Gamma	$\psi(\varepsilon) \propto \varepsilon \log(1/\varepsilon)^{k-1}$	$\alpha = 1$	$\beta = 2/(4 + d)$	$n^{-2/(4+d)} \log(n)^{\beta(k)}$
Cauchy	$\psi(\varepsilon) \propto \sqrt{\varepsilon}$	$\alpha = 1$	$\beta = 1/(3 + d)$	$n^{-2/(3+d)}$
Power laws/Pareto	$\psi(\varepsilon) \propto \varepsilon^{p/(p+1)}$	$\alpha = 1 \wedge p$	$\beta = \frac{2(p+1)}{p(3+\alpha+d)+2+d}$	$n^{\frac{-4(p+1)}{p(3+\alpha+d)+2+d}}$

Explicit rates for specific location models. We investigate here the influence of the function ψ in $\mathcal{P}_{\mathcal{T},\psi}$ as well as the one of the margin parameter on the convergence rates through several specific location models. These models are defined as follows: given any positive random variable Z (whose cumulative distribution function is denoted as F) and two real location values a and b , the random variable X is given by

$$(4.5) \quad X = \varepsilon Z + Yb + (1 - Y)a,$$

where ε is a Rademacher random variable (whose values is ± 1) independent of Z , and Y is the label of the observation, sampled independently of ε and Z with a Bernoulli law $\mathcal{B}(1/2)$. Using a translation invariance argument, it is enough in the next study to consider $a = 0$ and $b > 0$. Table 1 illustrates the rate reached by the k -nearest neighbor procedure in each situation.

A numerical study for “power laws.” In order to illustrate equations (4.2) and (4.3), we consider some specific cases of “power laws” such that

$$\mathbb{P}_X(\mu(X) < \varepsilon) = \psi(\varepsilon) \sim \varepsilon^g \quad \text{when } \varepsilon \rightarrow 0^+,$$

for some $g > 0$. In this case, the upper bound on the k -nearest neighbor classifier is given by

$$\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \lesssim n^{-(1+\alpha)/(1+\alpha+(2+d)/g)}$$

although the lower bound derived from (4.2) is

$$\inf_{\Phi_n} \sup_{\mathbb{P}_{(X,Y)} \in \mathcal{F}_{L,\alpha} \cap \mathcal{P}_{\mathcal{T},\psi} \cap \mathfrak{M}_{\text{mma}}(\Omega, \kappa)} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \gtrsim n^{-(1+\alpha)/(\alpha+(2+d)/g)}.$$

We immediately observe that the classification rates are seriously damaged when g is small. In contrast, for very thin tails, the rate can be arbitrarily close to n^{-1} . For this purpose, we illustrate this phenomenon with a family of distributions \mathcal{P}_g , where the parameter $g > 0$ influences the tail size. We define the cumulative distribution function of the positive random variable Z :

$$\forall t \geq 0 \quad F_g(t) = 1 - \frac{1}{(t + 1)^g}.$$




FIG. 1. Example of observed empirical rates and upper bound theoretical rates given by (4.3) for several power law distributions of parameter g .

Then, for two real values (a, b) , we sample n observations (X_i, Y_i) according to the previous model and the Bayes classifier is given by

$$\Phi^*(X) = \mathbf{1}_{\{X > (a+b)/2\}}.$$

In this example, the margin α is equal to 1 and η is L -Lipschitz. We then consider $k_n = \lfloor n^{2/5} \rfloor + 1$ to assess the statistical performance of the k -nearest neighbor classifier. Figure 1 represents the excess risk obtained by the k -nearest neighbor classifier and the successive degradation of the convergence rate when g decreases to 0 [on the left, the empirical performance of the k -nearest neighbor rule with the underlying distributions and on the right for the upper bound theoretically derived from Theorem (4.3)]. These numerical experiments are consistent with the theoretical result obtained in Theorem 4.5.

Comparison between the standard k -nearest neighbor and its sliced counterpart. We provide here a short numerical study that aims to compare the results reached by the standard k -nearest neighbor rule described in Theorem 4.3 and the ones obtained by its sliced counterpart described in Section 4.3.3 and in Theorem 4.4. To measure such an improvement, we have chosen to once again use some noncompactly supported distributions and several different location models.

On the one hand, as pointed out in Theorem 4.3, the standard k -nearest neighbor procedure will be tuned with a number of neighbor $k_n := \lfloor n^{2/(3+\alpha+d)} \rfloor + 1$. On the other hand, the sliced k -nearest neighbor rule described in Theorem 4.4 requires a preliminary estimation of the law of observation \mathbb{P}_X . To do this, we used the recent kernel density estimation package² provided by [4], which is an adaptive estimator based on linear diffusion processes. Given any training set $(X_i, Y_i)_{1 \leq i \leq n}$, we first built the preliminary estimator $\hat{\mu}_n$ of the unknown density μ . This estimator is

²kde.m is available on the author's Website of [4].

interesting because of its adaptive smoothing properties and because it includes a very fast automatic bandwidth selection algorithm.

The sliced k -nearest neighbor rule then uses a number of neighbors that depends on the design point X . If the density estimate is large enough, that is, if $\hat{\mu}_n(X) \geq n^{-\alpha/(2+\alpha+d)}$,

$$k_n(X) := \lfloor n^{2/(2+\alpha+d)} \rfloor + 1.$$

Otherwise, when $2^{-(j+1)} \leq \hat{\mu}_n(X)n^{\alpha/(2+\alpha+d)} \leq 2^{-j}$, the number $k_n(X)$ is

$$k_n(X) := \lfloor n^{2/(2+\alpha+d)} 2^{-2j/(2+d)} \rfloor + 1.$$

To draw some reliable comparisons, we also used some various univariate laws ($d = 1$) for the random variable Z involved in the definition of the location model (4.5) (Normal distributions, Cauchy distributions and Power laws) whose parameters are described in Table 1. The two location parameters are still denoted a and b and fixed such that $a = -b$.

In each situation, we used a Monte-Carlo strategy with 1000 replications to compute the mean excess risk of each k -nearest neighbor rule. We used a training set of cardinal n , as well as a test set of size 200. Results are given in Table 2.

We may observe in Table 2 that the sliced version of the k -nearest neighbor always outperforms the standard one. Such a numerical result is consistent with the theoretical ones of Theorems 4.3 and 4.4. Note also that the relative improvement of the sliced k -nearest neighbor rule seems to increase when the number of observations n growth, meaning that each excess risk of the two procedures varies with a different power of n .

Finally, it should be mentioned that we have not tried to modify the dimension of the observations X . Indeed, the difference of the upper bounds given by Theorems 4.3 and 4.4 becomes more and more negligible when the dimension is increasing. This should also be the case in the empirical study that will be in the subject of a future work. Likewise, the statistical study of the empirical sliced k -nearest neighbor rule should also be addressed in a future study, since a balance between the

TABLE 2
 Mean excess risk multiplied by 100 (left: standard k -nearest neighbor; middle: sliced k -nearest neighbor; right: percentage of improvement). Standard errors are lower than 0.2

Law of Z	$n = 100$			$n = 500$			$n = 1000$		
Gauss, $a = 1, \sigma = 2$	4.51	4.47	8.9%	3.34	2.48	25.7%	2.71	1.81	33%
Cauchy, $a = \frac{1}{2}, \gamma = \frac{1}{2}$	2.43	2.08	14.5%	1.26	1.07	14.8%	1	0.83	16%
Cauchy, $a = \frac{1}{2}, \gamma = 1$	4.6	3.78	17.8%	2.89	2.08	27%	2.3 _{0,2}	1.55	32.5%
Power, $a = \frac{1}{2}, \gamma = 1$	4.13	3.32	19.6%	2.48	1.9	23.7%	2.06	1.49	27.3%
Power, $a = \frac{1}{2}, \gamma = 2$	2.18	1.92	12.1%	1.09	0.95	12.4%	0.79	0.69	12.6%

estimation $\hat{\mu}_n$ of the density μ and the excess risk of classification with the sliced rule may exist. We have left this problem open for a future study.

APPENDIX: PROOF OF THE UPPER BOUNDS

Recall that \mathbb{E} (resp., $\mathbb{E}_X, \mathbb{E}_{\otimes^n}$) denote the expectation with respect to the measure \mathbb{P} (resp., $\mathbb{P}_X, \mathbb{P}_{\otimes^n}$). Let $\varepsilon > 0$ be a given real number (whose value will be specified later), and define

$$\mathcal{B}_\varepsilon := \{x \in \mathbb{R}^d \mid |\eta(x) - 1/2| \leq \varepsilon\}.$$

Applying Proposition 2.1 in Section 2 of the supplementary material [15], the excess risk can be decomposed as follows:

$$\begin{aligned} \mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) &= \mathbb{E}[|2\eta(X) - 1| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}}] \\ &= \underbrace{\mathbb{E}[|2\eta(X) - 1| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{X \in \mathcal{B}_\varepsilon}]}_{:=T_{1,\varepsilon}} \\ &\quad + \underbrace{\mathbb{E}[|2\eta(X) - 1| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{X \in \mathcal{B}_\varepsilon^c}]}_{:=T_{2,\varepsilon}}. \end{aligned}$$

Now, the margin Assumption A2 yields

$$(A.1) \quad T_{1,\varepsilon} \leq 2\mathbb{E}[|\eta(X) - 1/2| \mathbf{1}_{X \in \mathcal{B}_\varepsilon}] \leq 2\varepsilon \mathbb{P}_X(X \in \mathcal{B}_\varepsilon) \leq 2C\varepsilon^{1+\alpha}.$$

In order to control $T_{2,\varepsilon}$, define

$$\forall j \geq 1 \quad \mathcal{B}_{\varepsilon,j} := \{x \in \mathbb{R}^d \mid 2^{j-1}\varepsilon \leq |\eta(x) - 1/2| \leq 2^j\varepsilon\}.$$

Now,

$$\begin{aligned} T_{2,\varepsilon} &= 2 \sum_{j \geq 1} \mathbb{E}[|\eta(X) - 1/2| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{\{X \in \mathcal{B}_{\varepsilon,j}\}}] \\ &\leq 2\varepsilon \sum_{j \geq 1} 2^j \mathbb{E}_X[\mathbf{1}_{\{X \in \mathcal{B}_{\varepsilon,j}\}} \mathbb{E}_{\otimes^n}(\mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}})]. \end{aligned}$$

We can apply Proposition 2.2 of the supplementary material [15] to obtain

$$(A.2) \quad T_{2,\varepsilon} \leq 4\varepsilon \sum_{j \geq 1} 2^j \mathbb{E}_X[\mathbf{1}_{\{X \in \mathcal{B}_{\varepsilon,j}\}} \exp(-2k_n [2^{j-1}\varepsilon - \Delta_n(X)]_+^2)].$$

Since μ is lower bounded by $a > 0$ on Ω , we can apply Proposition 2.3 of the supplementary material [15] with $a = \mu_-$ to obtain

$$\Delta_n(X) \leq C \left(\left(\frac{k_n}{n} \mu_-^{-1} \right)^{1/d} + \exp(-3k_n/14) \right).$$

Now, we consider $\varepsilon = \varepsilon_n \geq 2\Delta_n(X)$, for example, by choosing

$$(A.3) \quad \varepsilon_n := 2C \left(\left(\frac{k_n}{n} a^{-1} \right)^{1/d} + \exp(-3k_n/14) \right).$$

With ε_n defined as in (A.3), we deduce that $2^{j-1}\varepsilon_n - \Delta_n(X) \geq 2^{j-1}\varepsilon_n - \frac{\varepsilon_n}{2} \geq \varepsilon_n(2^{j-1} - \frac{1}{2}) > 0$. Thus, (A.2) becomes

$$T_{2,\varepsilon_n} \leq 4\varepsilon_n \sum_{j \geq 1} 2^j \mathbb{E}_X [\mathbf{1}_{\{0 < |\eta(X) - 1/2| < 2^j \varepsilon_n\}} \exp(-2k_n \varepsilon_n^2 (2^{j-1} - 1/2)^2)].$$

Now, in order to control the previous bound, we get the constraint on k_n :

$$(A.4) \quad k_n = \varepsilon_n^{-2}.$$

Thanks to (A.3), the constraint (A.4) then yields

$$(A.5) \quad \varepsilon_n \sim n^{-1/(2+d)} \quad \text{and} \quad k_n \sim n^{2/(2+d)}.$$

We then obtain that

$$\begin{aligned} T_{2,\varepsilon_n} &\leq 4\varepsilon_n \sum_{j \geq 1} 2^j \mathbb{E}_X \left[\mathbf{1}_{\{0 < |\eta(X) - 1/2| < 2^j \varepsilon_n\}} \exp\left(-\frac{2^{2j}}{8}\right) \right] \\ &\leq \varepsilon_n \sum_{j \geq 1} 2^{j+2} \exp\left(-\frac{2^{2j}}{8}\right) \mathbb{P}_X(|\eta(X) - 1/2| < 2^j \varepsilon_n). \end{aligned}$$

The margin assumption applied to $\mathbb{P}_X(|\eta(X) - 1/2| < 2^j \varepsilon_n)$ leads to

$$T_{2,\varepsilon_n} \leq \varepsilon_n^{1+\alpha} \sum_{j \geq 1} 2^{j(1+\alpha)+2} \exp\left(-\frac{2^{2j}}{8}\right).$$

The series on the right-hand side converges. This last bound associated with (A.1) leads to

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \leq Cn^{-(1+\alpha)/(2+d)}.$$

A.1. Proof of the upper bounds: Theorem 4.3 and Theorem 4.5(ii).

PROOF OF THEOREM 4.3. We consider a constant γ and use the following decomposition of \mathbb{R}^d for a suitable $\gamma > 0$ (that will be chosen later on):

$$\mathbb{R}^d = \underbrace{\{x : 0 \leq \mu(x) \leq n^{-\gamma}\}}_{R_n} \cup \underbrace{\{x : \mu > n^{-\gamma}\}}_{Q_n}.$$

We follow the road map of the proof of Theorem 3.1 and keep the notation \mathcal{B}_ε , which refers to $\mathcal{B}_\varepsilon := \{x \in \mathbb{R}^d : |\eta(x) - 1/2| \leq \varepsilon\}$. Thanks to Proposition 2.1 of the supplementary material [15], we obtain

$$\begin{aligned} \mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) &= \mathbb{E}[2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}}] \\ &= \underbrace{\mathbb{E}[2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{X \in R_n}]}_{:=T_{R_n}} \\ &\quad + \underbrace{\mathbb{E}[2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{X \in Q_n}]}_{:=T_{Q_n}}. \end{aligned}$$

Study of R_n . The tail Assumption A4 in the particular case where $\psi = \text{Id}$ leads to

$$T_{R_n} \leq \mathbb{P}_X(X \in R_n) = \mathbb{P}_X(\mu(X) \leq n^{-\gamma}) \lesssim n^{-\gamma}.$$

Study of Q_n . Following the proof of Theorem 3.1 with $a = n^{-\gamma}$, equations (A.2)–(A.4) yield

$$(A.6) \quad T_{Q_n} \leq C \varepsilon_n^{1+\alpha},$$

where ε_n and k_n satisfy the balance equations

$$\varepsilon_n \sim 2C \left(\frac{k_n}{n} a^{-1}\right)^{1/d} = 2C \left(\frac{k_n}{n^{1-\gamma}}\right)^{1/d} \quad \text{and} \quad k_n = \varepsilon_n^{-2}.$$

The equilibria are met in the two terms above with

$$(A.7) \quad k_n \sim Cn^{2(1-\gamma)/(2+d)} \quad \text{and} \quad \varepsilon_n \lesssim n^{-(1-\gamma)/(2+d)}.$$

Final control of the risk. From the previous bounds, we obtain that

$$(A.8) \quad \mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \lesssim n^{-(1-\gamma)(1+\alpha)/(2+d)} + n^{-\gamma}.$$

We optimize the last expression with respect to γ by setting

$$(1 - \gamma)(1 + \alpha) = \gamma(2 + d) \quad \Leftrightarrow \quad \gamma = \frac{1 + \alpha}{3 + \alpha + d}.$$

The above choices allow us to conclude that

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \leq Cn^{-(1+\alpha)/(3+\alpha+d)}. \quad \square$$

PROOF OF THEOREM 4.5. (ii) We follow the road map of the previous proof and replace the threshold $n^{-\gamma}$ with a_n , which should be carefully chosen. The key balance is still $k_n = v_n^{-2}$ on the set $\{\mu \geq a_n\}$ with the optimal setting

$$\frac{k_n}{na_n} \lesssim v_n^d.$$

Since we want to obtain a minimal value for v_n , this last equation leads to the choice

$$(A.9) \quad a_n = \frac{1}{nv_n^{2+d}},$$

and the upper bound of the excess risk we obtained is then

$$\sup_{F \in \mathcal{F}} [\mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*)] \lesssim v_n^{1+\alpha} + \psi(a_n).$$

The natural equilibrium is found when plug-in (A.9) in this last upper bound and v_n are be fixed so that $\psi^{-1}(v_n^{1+\alpha}) = n^{-1}v_n^{-(2+d)}$. We then obtain the rate $v_n^{1+\alpha}$ with the balance equation

$$v_n^{2+d} \psi^{-1}(v_n^{1+\alpha}) \sim n^{-1}. \quad \square$$

A.2. Proof of Theorem 4.4 (sliced k -nearest neighbor).

PROOF. We use the partition of Ω naturally derived from the slices $\Omega_{n,0}$ and $(\Omega_{n,j})_{j \geq 1}$:

$$\Omega_{n,0} := \{x | \mu(x) \geq n^{-\gamma}\}$$

and

$$\Omega_{n,j} := \{x | n^{-\gamma} 2^{-(j+1)} \leq \mu(x) \leq n^{-\gamma} 2^{-j}\}.$$

For this purpose, let $\gamma \in (0, 1)$ (that will be specified later) and

$$k_{n,j} = k_{n,0} 2^{-2j/(2+d)} \quad \text{with } k_{n,0} = n^{2(1-\gamma)/(2+d)} \log(n).$$

We then use the following decomposition of the excess risk:

$$\begin{aligned} & \mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \\ &= \mathbb{E} \left[2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \left(\mathbf{1}_{\Omega_{n,0}} + \sum_{j=1}^{+\infty} \mathbf{1}_{\Omega_{n,j}} \right) \right] \\ &= \mathbb{E} [2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{\Omega_{n,0}}] \\ &\quad + \sum_{j=1}^{+\infty} \mathbb{E} [2\eta(X) - 1 | \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} \mathbf{1}_{\Omega_{n,j}}] \\ &:= T_n + \sum_{j=1}^{+\infty} R_{n,j}. \end{aligned}$$

Study of T_n . The density is lower bounded by $n^{-\gamma}$ on $\Omega_{n,0}$. The proof of Theorem 4.3 yields [see (A.6) and (A.7)]

$$T_n \leq n^{-(1+\alpha)(1-\gamma)/(2+d)}.$$

Study of $R_{n,j}$. For any $j > J_0(n) := (1 - \gamma) \frac{\log(n)}{\log(2)}$ and for any $x \in \Omega_{n,j}$ with $j > J_0$, we have

$$\mu(x) < n^{-\gamma} 2^{-(1-\gamma) \log(n) / \log(2)} = 1/n.$$

The tail assumption with $\psi = \text{Id}$ leads to

$$\sum_{j > J_0} R_{n,j} \leq \mathbb{P}_X[\mu(X) \leq n^{-1}] \lesssim n^{-1}.$$

Study of $R_{n,j}$, $j \leq J_0(n)$. We consider the intermediary slices, and for $1 \leq j \leq J_0(n)$,

$$R_{n,j} = \underbrace{\mathbb{E} \left[|2\eta(X) - 1| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} [\mathbf{1}_{|\eta(X) - 1/2| \leq \varepsilon_{n,j}}] \mathbf{1}_{X \in \Omega_{n,j}} \right]}_{:=R_{n,j,1}} + \underbrace{\mathbb{E} \left[|2\eta(X) - 1| \mathbf{1}_{\{\Phi_n(X) \neq \Phi^*(X)\}} [\mathbf{1}_{|\eta(X) - 1/2| > \varepsilon_{n,j}}] \mathbf{1}_{X \in \Omega_{n,j}} \right]}_{:=R_{n,j,2}},$$

where $\varepsilon_{n,j}$ will be chosen later. To bound $R_{n,j,1}$, we use the fact that $|\eta - 1/2| \leq \varepsilon_{n,j}$ as well as the tail assumption on the set $\Omega_{n,j} \subset \{\mu(X) \leq n^{-\gamma} 2^{-j}\}$ to obtain

(A.10)
$$R_{n,j,1} \leq 2\varepsilon_{n,j} n^{-\gamma} 2^{-j}.$$

Thanks to Proposition 2.2 of the supplementary material [15], we can bound the term $R_{n,j,2}$ as follows:

$$R_{n,j,2} \leq 4\mathbb{E}[\mathbf{1}_{X \in \Omega_{n,j}} \exp(-2k_{n,j}[\varepsilon_{n,j} - \Delta_n(X)]_+^2)].$$

The term $\varepsilon_{n,j}$ is then chosen such that $\varepsilon_{n,j} - \Delta_n(X) \leq \varepsilon_{n,j}/2$. According to Proposition 2.3 of the supplementary material [15], we obtain

$$\varepsilon_{n,j} = c \left(k_{n,j} \frac{n^\gamma 2^{j+1}}{n} \right)^{1/d},$$

where c is chosen large enough. With this value, we obtain the following simplifications:

(A.11)
$$\begin{aligned} \varepsilon_{n,j}^2 k_{n,j} &= c^2 \frac{k_{n,j}^{1+2/d} 2^{(j+1)/d}}{n^{(1-\gamma)/d}} = c^2 \frac{k_{n,0}^{1+2/d}}{n^{(1-\gamma)/d}} 2^{-(2j/(2+d))(1+2/d)} 2^{(2j+2)/d} \\ &= c^2 2^{2/d} \log(n)^{1+2/d}. \end{aligned}$$

Taken together, (A.10) and (A.11) lead to

$$R_{n,j} \leq 2\varepsilon_{n,j} n^{-\gamma} 2^{-j} + 4 \exp(-c^2 2^{2/d} \log(n)^{1+2/d}) \mathbb{P}(X \in \Omega_{n,j}).$$

We then sum up all these terms for $j \geq 1$:

$$\begin{aligned} R_n &\lesssim \sum_{j=1}^{J_0} \frac{\log(n)^{1/2+1/d}}{\sqrt{k_{n,j}}} n^{-\gamma} 2^{-j} + \frac{1}{n} \sum_{j=1}^{J_0} \mathbb{P}(X \in \Omega_{n,j}) \\ &\lesssim n^{-\gamma} \log(n)^{1/2+1/d} \sum_{j=1}^{J_0} 2^{-j} k_{n,j}^{-1/2} + \frac{1}{n} \\ &\lesssim n^{-\gamma} n^{-(1-\gamma)/(2+d)} \log(n)^{1/2+1/d} \sum_{j=1}^{J_0} 2^{-j+(j+1)/(2+d)} + \frac{1}{n} \\ &\lesssim n^{-\gamma} n^{-(1-\gamma)/(2+d)} \log(n)^{1/2+1/d} + \frac{1}{n}. \end{aligned}$$

We can see in this last upper bound that we obtain an improvement between the standard rule and the one fixed here since the term $n^{-\gamma}$ that appears in the tail of μ on the right-hand side of (A.8) is transformed into $n^{-\gamma} \times n^{-(1-\gamma)/(2+d)}$ up to a log term.

Final equilibrium. We now fix the optimal value of γ with the conjunction of the upper bounds for R_n and T_n :

$$\begin{aligned} \mathcal{R}(\Phi_n) - \mathcal{R}(\Phi^*) \\ \lesssim n^{-(1+\alpha)((1-\gamma)/(2+d))} + n^{-\gamma-(1-\gamma)/(2+d)} \log(n)^{1/2+1/d} + O(1/n). \end{aligned}$$

The balance equilibrium is reached with $(1 + \alpha)\frac{1-\gamma}{2+d} = \gamma + \frac{1-\gamma}{2+d}$, meaning that $\gamma = \frac{1}{2+\alpha+d}$. This completes the proof. \square

Acknowledgements. The authors are grateful to an Associate Editor and the referees for helpful and valuable comments. We also thank C. Bordenave for fruitful discussions on an earlier version of the paper.

SUPPLEMENTARY MATERIAL

Supplement to “Classification in general finite dimensional spaces with the k -nearest neighbor rule”. (DOI: [10.1214/15-AOS1395SUPP](https://doi.org/10.1214/15-AOS1395SUPP); .pdf). Supplement contains some technical results and the proofs of Theorem 4.1, 4.2 and 4.5(i).

REFERENCES

- [1] AMIT, Y. and GEMAN, D. (1997). Shape quantization and recognition with randomized trees. *Neural Comput.* **9** 1545–1588.
- [2] AUDIBERT, J.-Y. and TSYBAKOV, A. B. (2007). Fast learning rates for plug-in classifiers. *Ann. Statist.* **35** 608–633. [MR2336861](#)
- [3] BARNDORFF-NIELSEN, O. E. and COX, D. R. (1989). *Asymptotic Techniques for Use in Statistics*. Chapman & Hall, London. [MR1010226](#)
- [4] BOTEV, Z. I., GROTHOWSKI, J. F. and KROESE, D. P. (2010). Kernel density estimation via diffusion. *Ann. Statist.* **38** 2916–2957. [MR2722460](#)
- [5] BOUCHERON, S., BOUSQUET, O. and LUGOSI, G. (2005). Theory of classification: A survey of some recent advances. *ESAIM Probab. Stat.* **9** 323–375. [MR2182250](#)
- [6] BREIMAN, L. (2001). Random forests. *Mach. Learn.* **45** 5–32.
- [7] CANNINGS, T. (2013). Nearest neighbour classification in the tails of a distribution. Preprint.
- [8] CÉROU, F. and GUYADER, A. (2006). Nearest neighbor classification in infinite dimension. *ESAIM Probab. Stat.* **10** 340–355 (electronic). [MR2247925](#)
- [9] CHAUDHURI, K. and DASGUPTA, S. (2014). Rates of convergence for nearest neighbor classification. In *Advances in Neural Information Processing Systems* (Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence and K. Q. Weinberger, eds.) **27** 3437–3445. Curran Associates, Red Hook, NY.

- [10] DEVROYE, L. (1981). On the almost everywhere convergence of nonparametric regression function estimates. *Ann. Statist.* **9** 1310–1319. [MR0630113](#)
- [11] DEVROYE, L., GYÖRFI, L., KRZYŻAK, A. and LUGOSI, G. (1994). On the strong universal consistency of nearest neighbor regression function estimates. *Ann. Statist.* **22** 1371–1385. [MR1311980](#)
- [12] DEVROYE, L., GYÖRFI, L. and LUGOSI, G. (1996). *A Probabilistic Theory of Pattern Recognition. Applications of Mathematics (New York)* **31**. Springer, New York. [MR1383093](#)
- [13] FIX, E. and HODGES, J. L. (1951). Discriminatory analysis, nonparametric discrimination, consistency properties. Randolph Field, Texas, Project 21-49-004, Report 4.
- [14] FREUND, Y. and SCHAPIRE, R. E. (1997). A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* **55** 119–139. [MR1473055](#)
- [15] GADAT, S., KLEIN, T. and MARTEAU, C. (2015). Supplement to “Classification in general finite dimensional spaces with the k -nearest neighbor rule.” DOI:10.1214/15-AOS1395SUPP.
- [16] GOLDENSHLUGER, A. and LEPSKI, O. (2014). On Adaptive Minimax Density Estimation on R^d . *Probab. Theory Related Fields* **159** 479–543. [MR3230001](#)
- [17] GYÖRFI, L. (1978). On the Rate of Convergence of Nearest Neighbor Rules. *IEEE Trans. Inform. Theory* **24** 509–512. [MR0501595](#)
- [18] GYÖRFI, L., KOHLER, M., KRZYŻAK, A. and WALK, H. (2002). *A Distribution-Free Theory of Nonparametric Regression*. Springer, New York. [MR1920390](#)
- [19] HALL, P., PARK, B. U. and SAMWORTH, R. J. (2008). Choice of neighbor order in nearest-neighbor classification. *Ann. Statist.* **36** 2135–2152. [MR2458182](#)
- [20] HÜSLER, J., LIU, R. Y. and SINGH, K. (2002). A formula for the tail probability of a multivariate normal distribution and its applications. *J. Multivariate Anal.* **82** 422–430. [MR1921635](#)
- [21] LECUÉ, G. (2007). Simultaneous adaptation to the margin and to complexity in classification. *Ann. Statist.* **35** 1698–1721. [MR2351102](#)
- [22] LI, W. V. and SHAO, Q.-M. (2001). Gaussian processes: Inequalities, small ball probabilities and applications. In *Stochastic Processes: Theory and Methods. Handbook of Statist.* **19** 533–597. North-Holland, Amsterdam. [MR1861734](#)
- [23] LIAN, H. (2011). Convergence of functional k -nearest neighbor regression estimate with functional responses. *Electron. J. Stat.* **5** 31–40. [MR2773606](#)
- [24] LOUSTAU, S. and MARTEAU, C. (2015). Minimax fast rates for discriminant analysis with errors in variables. *Bernoulli* **21** 176–208. [MR3322316](#)
- [25] MAMMEN, E. and TSYBAKOV, A. B. (1999). Smooth discrimination analysis. *Ann. Statist.* **27** 1808–1829. [MR1765618](#)
- [26] REYNAUD-BOURET, P., RIVOIRARD, V. and TULEAU-MALOT, C. (2011). Adaptive density estimation: A curse of support? *J. Statist. Plann. Inference* **141** 115–139. [MR2719482](#)
- [27] RODRÍGUEZ CASAL, A. (2007). Set estimation under convexity type assumptions. *Ann. Inst. Henri Poincaré Probab. Stat.* **43** 763–774. [MR3252430](#)
- [28] SAMWORTH, R. J. (2012). Optimal weighted nearest neighbour classifiers. *Ann. Statist.* **40** 2733–2763. [MR3097618](#)
- [29] STEINWART, I. (2005). Consistency of support vector machines and other regularized kernel classifiers. *IEEE Trans. Inform. Theory* **51** 128–142. [MR2234577](#)
- [30] STONE, C. J. (1977). Consistent nonparametric regression. *Ann. Statist.* **5** 595–645. [MR0443204](#)
- [31] TSYBAKOV, A. B. (2004). Optimal aggregation of classifiers in statistical learning. *Ann. Statist.* **32** 135–166. [MR2051002](#)

[32] VAPNIK, V. N. (1998). *Statistical Learning Theory*. Wiley, New York. [MR1641250](#)

S. GADAT
TOULOUSE SCHOOL OF ECONOMICS
UNIVERSITÉ TOULOUSE 1—CAPITOLE
21 ALLÉE DE BRIENNE
31000 TOULOUSE
FRANCE
E-MAIL: sebastien.gadat@math.univ-toulouse.fr

T. KLEIN
INSTITUT MATHÉMATIQUES DE TOULOUSE
UNIVERSITÉ TOULOUSE 3—PAUL SABATIER
118 ROUTE DE NARBONNE
31400 TOULOUSE
FRANCE
E-MAIL: thierry.klein@math.univ-toulouse.fr

C. MARTEAU
INSTITUT CAMILLE JORDAN
UNIVERSITÉ LYON 1—CLAUDE BERNARD
43 BOULEVARD DU 11 NOVEMBRE 1918
69622 VILLEURBANNE
FRANCE
E-MAIL: marteau@math.univ-lyon1.fr