
Regression with reject option and application to k NN

Ahmed Zaoui*

LAMA, Université Gustave Eiffel
ahmed.zaoui@univ-eiffel.fr

Christophe Denis

LAMA, Université Gustave Eiffel
MIA-Paris, AgroParisTech, INRAE, Université Paris-Saclay
christophe.denis@univ-eiffel.fr

Mohamed Hebiri

LAMA, Université Gustave Eiffel
CREST, ENSAE, Institut Polytechnique de Paris
mohamed.hebiri@univ-eiffel.fr

Abstract

We investigate the problem of regression where one is allowed to abstain from predicting. We refer to this framework as *regression with reject option* as an extension of classification with reject option. In this context, we focus on the case where the rejection rate is fixed and derive the optimal rule which relies on thresholding the conditional variance function. We provide a semi-supervised estimation procedure of the optimal rule involving two datasets: a first *labeled* dataset is used to estimate both regression function and conditional variance function while a second *unlabeled* dataset is exploited to calibrate the desired rejection rate. The resulting predictor with reject option is shown to be almost as good as the optimal predictor with reject option both in terms of risk and rejection rate. We additionally apply our methodology with k NN algorithm and establish rates of convergence for the resulting k NN predictor under mild conditions. Finally, a numerical study is performed to illustrate the benefit of using the proposed procedure.

Keywords: Regression; Regression with reject option; k NN; Predictor with reject option.

1 Introduction

Confident prediction is a fundamental problem in statistical learning for which numerous efficient algorithms have been designed, *e.g.*, neural-networks, kernel methods, or k -Nearest-Neighbors (k NN) to name a few. However, even state-of-art methods may fail in some situations, leading to bad decision-making. Obvious damageable incidences of an erroneous decision may occur in several fields such as medical diagnosis, where a wrong estimation can be fatal. In this work, we provide a novel statistical procedure designed to handle these cases. In the specific context of regression, we build a prediction algorithm that allows to abstain from predicting when the doubt is too important. As a generalization of the classification with reject option setting [5, 6, 7, 13, 15, 18, 25], this framework is naturally referred to as *regression with reject option*. In the spirit of [7], we opt here for a strategy where the predictor can abstain up to a fraction $\varepsilon \in (0, 1)$ of the data. The merit of this approach is that it allows human action on the proportion of the data where the prediction is too difficult while

*ahmed.zaoui@univ-eiffel.fr

standard machine learning algorithms can be exploited to perform the predictions on the other fraction of the data. The difficulty to address a prediction is then automatically evaluated by the procedure. From this perspective, this strategy may improve the efficiency of the human intervention.

In this paper, we investigate the regression problem with reject option when the rejection (or abstention) rate is controlled. Specifically, we provide a statistically principled and computationally efficient algorithm tailored to this problem. We first formally define the regression with reject option framework, and explicitly exhibit the optimal predictor with bounded rejection rate in Section 2. This optimal rule relies on a thresholding of the conditional variance function. This result is the bedrock of our work and suggests the use of a plug-in approach. We propose in Section 3 a two-step procedure which first estimates both the regression function and the conditional variance function on a first *labeled* dataset and then calibrates the threshold responsible for abstention using a second *unlabeled* dataset. Under mild assumptions, we show that our procedure performs as well as the optimal predictor both in terms of risk and rejection rate. We emphasize that our procedure can be exploited with any off-the-shell estimator. As an example we apply in Section 4 our methodology with the k NN algorithm for which we derive rates of convergence. Finally, we perform numerical experiments in Section 5 which illustrate the benefits of our approach. In particular, it highlights the flexibility of the proposed procedure.

Rejection in regression is extremely rarely considered in the literature, an exception being [26] that views the reject option from a different perspective. There, the authors used the reject option from the side of ε -optimality, and therefore ensures that the prediction is inside a ball with radius ε around the regression function with high probability. Their methodology is intrinsically associated with empirical risk minimization procedures. In contrast, our method is applicable to any estimation procedure. Closer related works to ours appears in classification with reject option literature [2, 5, 6, 7, 13, 15, 18, 25]. In particular, the present work can be viewed as an extension of the classification with reject option setting. Indeed, from a general perspective, the present contribution brings a deeper understanding of the reject option. Importantly, the conditional variance function appears to capture the main feature behind the abstention decision. In [7], the authors also provide rates of convergence for plug-in type approaches in the case of bounded rejection rate. However, their rates of convergence holds only under some margin type assumption [1, 19] and a smoothness assumption on the considered estimator. On the contrary, we do not require these assumptions to get valid rates of convergence.

2 Regression with reject option

In this section we introduce the regression with reject option setup and derive a general form of the optimal rule in this context. We additionally highlight the case of fixed rejection rate as our main framework. First of all, before we proceed, let us introduce some preliminary notation. Let (X, Y) be a random couple taking its values in $\mathbb{R}^d \times \mathbb{R}$: here X denotes a feature vector and Y is the corresponding output. We denote by \mathbb{P} the joint distribution of (X, Y) and by \mathbb{P}_X the marginal distribution of the feature X . Let $x \in \mathbb{R}^d$, we introduce the regression function $f^*(x) = \mathbb{E}[Y|X = x]$ as well as the conditional variance function $\sigma^2(x) = \mathbb{E}[(Y - f^*(X))^2|X = x]$. We will give due attention to these two functions in our analysis. In addition, we denote by $\|\cdot\|$ the Euclidean on \mathbb{R}^d . Finally, $|\cdot|$ stands for the cardinality when dealing with a finite set.

2.1 Predictor with reject option

Let f be some measurable real-valued function which must be viewed as a prediction function. A *predictor with reject option* Γ_f associated to f is defined as being any function that maps \mathbb{R}^d onto $\mathcal{P}(\mathbb{R})$ such for all $x \in \mathbb{R}^d$, the output $\Gamma_f(x) \in \{\emptyset, \{f(x)\}\}$. We denote by Υ_f the set of all predictors with reject option that relies on f . Hence, in this framework, there are only two options for a particular $x \in \mathbb{R}^d$: whether the predictor with reject option outputs the empty set, meaning that no prediction is produced for x ; or the output $\Gamma_f(x)$ is of size 1 and the prediction coincides with the value $f(x)$. The framework of regression with reject option naturally brings into play two important characteristics of a given predictor Γ_f . The first one is the rejection rate that we denote by $r(\Gamma_f) = \mathbb{P}(|\Gamma_f(X)| = 0)$ and the second one is the L_2 error when prediction is performed

$$\text{Err}(\Gamma_f) = \mathbb{E}[(Y - f(X))^2 | |\Gamma_f(X)| = 1] \quad .$$

The ultimate goal in regression with reject option is to build a predictor Γ_f with a small rejection rate that achieves a small conditional L_2 error as well. A natural way to make this happen is to embed these quantities into a measure of performance. To this end, let consider the following risk

$$\mathcal{R}_\lambda(\Gamma_f) = \mathbb{E}[(Y - f(X))^2 \mathbb{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda r(\Gamma_f) \quad ,$$

where $\lambda \geq 0$ is a tuning parameter which is responsible for compromising error and rejection rate: larger λ 's result in predictors Γ_f with smaller rejection rates, but with larger errors. Hence, λ can be interpreted as the price to pay for using the reject option. Note that the above risk \mathcal{R}_λ has already been considered by [13] in the classification framework.

Minimizing the risk \mathcal{R}_λ , we derive an explicit expression of an optimal predictor with reject option.

Proposition 2.1. *Let $\lambda \geq 0$, and consider*

$$\Gamma_\lambda^* \in \arg \min \mathcal{R}_\lambda(\Gamma_f) \quad ,$$

where the minimum is taken over all predictors with reject option $\Gamma_f \in \Upsilon_f$ and all measurable functions f . Then we have that

1. The optimal predictor with rejected option Γ_λ^* can be written as

$$\Gamma_\lambda^*(X) = \begin{cases} \{f^*(X)\} & \text{if } \sigma^2(X) \leq \lambda \\ \emptyset & \text{otherwise} \end{cases} \quad . \quad (1)$$

2. For any $\lambda < \lambda'$, the following holds

$$\text{Err}(\Gamma_\lambda^*) \leq \text{Err}(\Gamma_{\lambda'}^*) \quad \text{and} \quad r(\Gamma_\lambda^*) \geq r(\Gamma_{\lambda'}^*) \quad .$$

Interestingly, this result shows that the oracle predictor relies on thresholding the conditional variance function σ^2 . We believe that this is an important remark that provides an essential characteristic of the reject option in regression but also in classification. Indeed, it has been shown that the optimal classifier with reject option for classification is obtained by thresholding the function f^* (see for instance [13]). However, in the binary case where $Y \in \{0, 1\}$, one has $\sigma^2(x) = f^*(x)(1 - f^*(x))$, and then thresholding σ^2 and f^* are equivalent.

The second point of the proposition shows that the error and the rejection rate of the optimal predictor are working in two opposite directions *w.r.t.* λ and then a compromise is required. We illustrate this aspect with the `airfoil` dataset, and the k NN predictor (see Section 5) in the contiguous Figure 1. The two curves correspond to the evaluation of the error $\text{Err}(\hat{\Gamma}_\lambda)$ (blue-solid line) and the rejection rate $r(\hat{\Gamma}_\lambda)$ (red-dashed line) as a function of λ . In general any choice of the parameter λ is difficult to interpret. Indeed, one of the major drawbacks of this approach is that any fixed λ (or even an ‘‘optimal’’ value of this parameter) does not allow to control neither of the two parts of the risk function. Especially, the rejection rate can be arbitrary large.

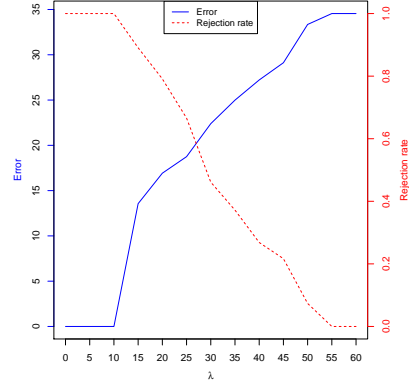


Figure 1: $\widehat{\text{Err}}(\hat{\Gamma}_\lambda)$ and $\hat{r}(\hat{\Gamma}_\lambda)$ vs. λ .

For this reason, we investigate in Section 2.2 the setting where the rejection rate is fixed. We understand this rejection rate as a budget one has beforehand.

2.2 Optimal predictor with fixed rejection rate

In this section, we introduce the framework where the rejection rate is fixed or at least bounded. That is to say, for a given predictor with reject option Γ_f and a given rejection rate $\varepsilon \in (0, 1)$, we ask that Γ_f satisfies following constraint $r(\Gamma_f) \leq \varepsilon$. This kind of constraint has also been considered by [7] in the classification setting. Our objective becomes to solve the constrained problem²:

$$\Gamma_\varepsilon^* \in \arg \min \{ \text{Err}(\Gamma_f) : r(\Gamma_f) \leq \varepsilon \} \quad . \quad (2)$$

²By abuse of notation, we refer to Γ_λ^* as the solution of the penalized problem and to Γ_ε^* as the solution of the constrained problem.

In the same vein as Proposition 2.1, we aim at writing an explicit expression of Γ_ε^* , referred in what follows to as ε -predictor. However, this expression is not well identified in the general case. Therefore, we make the following mild assumption on the distribution of $\sigma^2(X)$, which translates the fact that the function σ^2 is not constant on any set with non-zero measure *w.r.t.* \mathbb{P}_X .

Assumption 2.2. *The cumulative distribution function F_{σ^2} of $\sigma^2(X)$ is continuous.*

Let us denote by $F_{\sigma^2}^{-1}$ the generalized inverse of the cumulative distribution F_{σ^2} defined for all $u \in (0, 1)$ as $F_{\sigma^2}^{-1}(u) = \inf\{t \in \mathbb{R} : F_{\sigma^2}(t) \geq u\}$. Under Assumption 2.2 and from Proposition 2.1, we derive an explicit expression of the ε -predictor Γ_ε^* given by (2).

Proposition 2.3. *Let $\varepsilon \in (0, 1)$, and let $\lambda_\varepsilon = F_{\sigma^2}^{-1}(1 - \varepsilon)$. Under Assumption 2.2, we have $\Gamma_\varepsilon^* = \Gamma_{\lambda_\varepsilon}^*$.*

As an immediate consequence of the above proposition and properties on quantile functions is that

$$r(\Gamma_\varepsilon^*) = \mathbb{P}(|\Gamma_\varepsilon^*(X)| = 0) = \mathbb{P}(\sigma^2(X) \geq \lambda_\varepsilon) = \mathbb{P}(F_{\sigma^2}(\sigma^2(X)) \geq 1 - \varepsilon) = \varepsilon ,$$

and then the ε -predictor has rejection rate exactly ε . The continuity Assumption 2.2 is a sufficient condition to ensure that this property holds true. Besides, from this assumption, the ε -predictor can be expressed as follows

$$\Gamma_\varepsilon^*(x) = \begin{cases} \{f^*(x)\} & \text{if } F_{\sigma^2}(\sigma^2(x)) \leq 1 - \varepsilon \\ \emptyset & \text{otherwise .} \end{cases} \quad (3)$$

Finally, as suggested by Proposition 2.1 and 2.3, the performance of a given predictor with reject option Γ_f is measured through the risk \mathcal{R}_λ when $\lambda = \lambda_\varepsilon$. Then, its excess risk is given by

$$\mathcal{E}_{\lambda_\varepsilon}(\Gamma_f) = \mathcal{R}_{\lambda_\varepsilon}(\Gamma_f) - \mathcal{R}_{\lambda_\varepsilon}(\Gamma_\varepsilon^*) ,$$

for which the following result provides a closed formula.

Proposition 2.4. *Let $\varepsilon \in (0, 1)$. For any predictor Γ_f , we have*

$$\mathcal{E}_{\lambda_\varepsilon}(\Gamma_f) = \mathbb{E}_X [(f^*(X) - f(X))^2 \mathbb{1}_{\{|\Gamma_f(X)|=1\}}] + \mathbb{E}_X [|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\Gamma_f(X)| \neq |\Gamma_\varepsilon^*(X)|\}}] .$$

The above excess risk consists of two terms that translates two different aspect of the regression with reject option problem. The first one is related to the L_2 risk of the prediction function f and is rather classical in the regression setting. On contrast, the second is related to the reject option problem. It is dictated by the behavior of the conditional variance σ^2 around the threshold λ_ε .

3 Plug-in ε -predictor with reject option

We devote this section to the study of a data-driven predictor with reject option based on the *plug-in* principle that mimics the optimal rule derived in Proposition 2.3.

3.1 Estimation strategy

Equation (3) indicates that a possible way to estimate Γ_ε^* relies on the plug-in principle. To be more specific, Eq. (3) suggests that estimating f^* and σ^2 , as well as the cumulative distribution F_{σ^2} would be enough to get an estimator of Γ_ε^* . To build such predictor, we first introduce a learning sample $\mathcal{D}_n = \{(X_i, Y_i), i = 1, \dots, n\}$ which consists of n independent copies of (X, Y) . This dataset helps us to construct estimators \hat{f} and $\hat{\sigma}^2$ of the regression function f^* and the conditional variance function σ^2 respectively. In this paper, we focus on estimator $\hat{\sigma}^2$ which relies on the residual-based methods [11]. Based on \mathcal{D}_n , the estimator $\hat{\sigma}^2$ is obtained by solving the regression problem of the output variable $(Y - \hat{f}(X))^2$ on the input variable X . Estimating the last quantity F_{σ^2} is rather simple by replacing cumulative distribution function by its empirical version. Since this term only depends on the marginal distribution \mathbb{P}_X , we estimate it using a second *unlabeled* dataset $\mathcal{D}_N = \{X_{n+1}, \dots, X_{n+N}\}$ composed of N independent copies of X . This is an important feature of our methodology since unlabeled data are usually easy to get. The dataset \mathcal{D}_N is assumed to be independent of \mathcal{D}_n . We set

$$\hat{F}_{\hat{\sigma}^2}(\cdot) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{\hat{\sigma}^2(X_{n+i}) \leq \cdot\}} ,$$

as an estimator for $F_{\hat{\sigma}^2}$. With this notation, the *plug-in ε -predictor* is the predictor with reject option defined for each $x \in \mathbb{R}^d$ as

$$\hat{\Gamma}_\varepsilon(x) = \begin{cases} \hat{f}(x) & \text{if } \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(x)) \leq 1 - \varepsilon \\ \emptyset & \text{otherwise} \end{cases} \quad (4)$$

It is worth noting that the proposed methodology is flexible enough to rely upon any off-the-shelf estimators of the regression function f^* and the conditional variance function σ^2 .

3.2 Consistency of plug-in ε -predictors

In this part, we investigate the statistical properties of the plug-in ε -predictors with reject option. This analysis requires an additional assumption on the following quantity

$$F_{\hat{\sigma}^2}(\cdot) = \mathbb{P}_X(\hat{\sigma}^2(X) \leq \cdot | \mathcal{D}_n) \quad .$$

Assumption 3.1. *The cumulative distribution function $F_{\hat{\sigma}^2}$ of $\hat{\sigma}^2(X)$ is continuous.*

This condition is analogous to Assumption 2.2 but deals with the estimator $\hat{\sigma}^2(X)$ instead of the true conditional variance $\sigma^2(X)$. This difference makes Assumption 3.1 rather weak as the estimator $\hat{\sigma}^2(X)$ is chosen by the practitioner. Moreover, we can make any estimator satisfy this condition by providing a smoothed version of it. We illustrate this strategy with k NN algorithm in Section 4. Next theorem is the main result of this section, it establishes the consistency of the predictor $\hat{\Gamma}_\varepsilon$ to the optimal one.

Theorem 3.2. *Let $\varepsilon \in (0, 1)$. Assume that σ^2 is bounded, \hat{f} is a consistent estimator of f^* w.r.t. the L_2 risk, and $\hat{\sigma}^2$ is a consistent estimator of σ^2 w.r.t. the L_1 risk. Under Assumptions 2.2- 3.1, the followings hold*

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) \right]_{n, N \rightarrow +\infty} \longrightarrow 0, \quad \text{and} \quad \mathbb{E} \left[|r(\hat{\Gamma}_\varepsilon) - \varepsilon| \right] \leq CN^{-1/2} \quad ,$$

where $C > 0$ is an absolute constant.

This theorem establishes the fact that the plug-in ε -predictor behaves asymptotically as well as the optimal ε -predictor both in terms of risk and rejection rate. The convergence of the rejection rate requires only Assumption 3.1 which is rather weak and can even be removed following the process detailed in Section 4.2. In particular, the theorem shows that the rejection rate of the plug-in ε -predictor is of level ε up to a term of order $O(N^{-1/2})$. This rate is similar to the one obtained in the classification setting [7]. It relies on the difference between the cumulative distribution $F_{\hat{\sigma}^2}$ and its empirical counterpart $\hat{F}_{\hat{\sigma}^2}$ that is controlled using Dvoretzky-Kiefer-Wolfowitz Inequality [17]. Interestingly, this result applies to any consistent estimators of f^* and σ^2 .

The estimation of regression function f^* is widely studied and suitable algorithm such as random forests, kernel procedures, or k NN estimators can be used, see [3, 10, 20, 22, 23]. The estimation of the conditional variance function which relies on the residual-based methods has also been extensively studied based on kernel procedures, see for instance [9, 11, 12, 14, 21]. In the next section, we derive rates of convergence in the case where both estimators \hat{f} and $\hat{\sigma}^2$ rely on the k NN algorithm. In particular, we establish rates of convergence for $\hat{\sigma}^2$ in sup norm (see Proposition D.5 in the supplementary material).

4 Application to k NN algorithm: rates of convergence

The plug-in ε -predictor $\hat{\Gamma}_\varepsilon$ relies on estimators of the regression and the conditional variance functions. In this section, we consider the specific case of k NN based estimations. We refer to the resulting predictor as *k NN predictor with reject option*. Specifically, we establish rates of convergence for this procedure. In addition, since k NN estimator of σ^2 violates Assumption 3.1, applying our methodology to k NN has the benefit of illustrating the smoothing technique to make this condition be satisfied.

4.1 Assumptions

To study the performance of the k NN predictor with reject option in the finite sample regime, we assume that X belongs to a regular compact set $\mathcal{C} \subset \mathbb{R}^d$, see [1]. Besides, we make the following assumptions.

Assumption 4.1. *The functions f^* and σ^2 are Lipschitz.*

Assumption 4.2 (Strong density assumption). *We assume that the marginal distribution \mathbb{P}_X admits a density μ w.r.t to the Lebesgue measure such that for all $x \in \mathcal{C}$, we have $0 < \mu_{\min} \leq \mu(x) \leq \mu_{\max}$.*

These two assumptions are rather classical when we deal with rate of convergence and we refer the reader to the baseline books [10, 23]. In particular, we point out that the strong density assumption has been introduced in the context of binary classification for instance in [1]. The last assumption that we require highlights the behavior of σ^2 around the threshold λ_ε .

Assumption 4.3 (α -exponent assumption). *We say that σ^2 has exponent $\alpha \geq 0$ (at level λ_ε) with respect to \mathbb{P}_X if there exists $c^* > 0$ such that for all $t > 0$*

$$\mathbb{P}_X (0 < |\sigma^2(X) - \lambda_\varepsilon| \leq t) \leq c^* t^\alpha .$$

This assumption has been first introduced in [19] and is also referred as Margin assumption in the binary classification setting, see [16]. For $\alpha > 0$, Assumption 4.3 ensures that the random variable $\sigma^2(X)$ can not concentrate too much around the threshold λ_ε . It allows to derive faster rates of convergence. Note that, if $\alpha = 0$ there is no assumption.

4.2 k NN predictor with reject option

For any $x \in \mathbb{R}^d$, we denote by $(X_{(i,n)}(x), Y_{(i,n)}(x)), i = 1, \dots, n$ the reordered data according to the ℓ_2 distance in \mathbb{R}^d , meaning that $\|X_{(i,n)}(x) - x\| < \|X_{(j,n)}(x) - x\|$ for all $i < j$ in $\{1, \dots, n\}$. Note that Assumption 4.2 ensures that ties occur with probability 0 (see [10] for more details). Let $k = k_n$ be an integer. The k NN estimator of f^* and σ^2 are then defined, for all x , as follows

$$\hat{f}(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x) \quad \text{and} \quad \hat{\sigma}^2(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} \left(Y_{(i,n)}(x) - \hat{f}(X_{(i,n)}(x)) \right)^2 .$$

Conditional on \mathcal{D}_n , the cumulative distribution function $F_{\hat{\sigma}^2}$ is not continuous and then Assumption 3.1 does not hold. To avoid this issue, we introduce a random perturbation ζ distributed according to the Uniform distribution on $[0, u]$ that is independent from every other random variable where $u > 0$ is a (small) fixed real number that will be specified later. Then, we define the random variable $\bar{\sigma}^2(X, \zeta) := \hat{\sigma}^2(X) + \zeta$. It is not difficult to see that, conditional on \mathcal{D}_n the cumulative distribution $F_{\bar{\sigma}^2}$ of $\bar{\sigma}^2(X, \zeta)$ is continuous. Furthermore, by the triangle inequality, the consistency of $\hat{\sigma}^2$ implies the consistency of $\bar{\sigma}^2$ provided that u tends to 0. Therefore, we naturally define the k NN predictor with reject option as follows.

Let $(\zeta_1, \dots, \zeta_N)$ be independent copies of ζ and independent of every other random variable. We set

$$\hat{F}_{\bar{\sigma}^2}(\cdot) = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{\hat{\sigma}^2(X_{n+i}) + \zeta_i \leq \cdot\}} ,$$

and the k NN ε -predictor with reject option is then defined for all x and ζ as

$$\hat{\Gamma}_\varepsilon(x, \zeta) = \begin{cases} \left\{ \hat{f}(x) \right\} & \text{if } \hat{F}_{\bar{\sigma}^2}(\bar{\sigma}^2(x, \zeta)) \leq 1 - \varepsilon \\ \emptyset & \text{otherwise} . \end{cases}$$

4.3 Rates of convergence

In this section, we derive the rates of convergence of the k NN ε -predictor in the following framework. We assume that Y is bounded or that Y satisfies

$$Y = f^*(X) + \sigma(X)\xi , \quad (5)$$

where ξ is independent of X and distributed according to a standard normal distribution. Note that these assumptions covers a broad class of applications. Under these assumptions, we can state the following result.

Theorem 4.4. *Grant Assumptions 2.2, 4.1, 4.2, and 4.3. Let $\varepsilon \in (0, 1)$, if $k_n \propto n^{2/(d+2)}$, and $u \leq n^{-1/(d+2)}$, then the k NN ε -predictor $\hat{\Gamma}_\varepsilon$ satisfies*

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\hat{\Gamma}_\varepsilon \right) \right] \leq C \left(n^{-2/(d+2)} + \log(n)^{(\alpha+1)} n^{-(\alpha+1)/(d+2)} + N^{-1/2} \right),$$

where $C > 0$ is a constant which depends on f^* , σ^2 , c_0 , c^* , α , \mathcal{C} , and on the dimension d .

Each part of the above rate describes a given feature of the problem. The first one relies on the estimation error of the regression function. The second one, which depends in part on the parameter α from Assumption 4.3, is due to the estimation error in sup norm of the conditional variance $\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| \right) \right] \leq C \log(n) n^{-1/(d+2)}$ stated in Proposition D.5 in the supplementary material. Notice that for $\alpha = 1$, the second term is of the same order (up to logarithmic factor) as the term corresponding to the estimation of the regression function. The last term is directly linked to the estimation of the threshold λ_ε . Lastly, for $\alpha > 1$, we observe, provided that the size of the unlabeled sample N is sufficiently large, that this rate is the same as the rate of \hat{f} in L_2 norm which is then the best situation that we can expect for the rejection setting.

5 Numerical experiments

In this section, we present numerical experiments to illustrate the performance of the plug-in ε -predictor. The construction process of this predictor is described in Section 3.1 and relies on estimators of the regression and the conditional variance functions. The code used for the implementation of the plug-in ε -predictor can be found at https://github.com/ZaouiAmed/Neurips2020_RejectOption. For this experimental study, we consider the same algorithm for both estimation tasks and build three plug-in ε -predictors based respectively on support vector machines (svm), random forests (rf), and k NN (knn) algorithms. Besides, to avoid non continuity issues, we add the random perturbation $\zeta \sim \mathcal{U}[0, 10^{-10}]$ to all of the considered methods as described in Section 4.2. The performance is evaluated on two benchmark datasets: *QSAR aquatic toxicity* and *Airfoil Self-Noise* coming from the UCI database. We refer to these two datasets as *aquatic* and *airfoil* respectively. For all datasets, we split the data into three parts (50 % train labeled, 20 % train unlabeled, 30 % test). The first part is used to estimate both regression and variance functions, while the second part is used to compute the empirical cumulative distribution function. Finally, for each $\varepsilon \in \{i/10, i = 0, \dots, 9\}$ and each plug-in ε -predictor, we compute the empirical rejection rate \hat{r} and the empirical error $\widehat{\text{Err}}$ on the test set. This procedure is repeated 100 times and we report the average performance on the test set alongside its standard deviation. We employ the 10-fold cross-validation to select the parameter $k \in \{5, 10, 15, 20, 30, 50, 70, 100, 150\}$ of the k NN algorithm. For random forests and svm procedures, we used respectively the R packages `randomForest` and `e1071` with default parameters.

5.1 Datasets

The datasets used for the experiments are briefly described bellow:

QSAR aquatic toxicity has been used to develop quantitative regression QSAR models to predict acute aquatic toxicity towards the fish *Pimephales promelas*. This dataset is composed of $n = 546$ observations for which 8 numerical features are measured. The output takes its values in $[0.12, 10.05]$. *Airfoil Self-Noise* is composed of $n = 1503$ observations for which 5 features are measured. This dataset is obtained from a series of aerodynamic and acoustic tests. The output is the scaled sound pressure level, in decibels. It takes its values in $[103, 140]$.

Since the variance function plays a key role in the construction of the plug-in ε -predictor, we display in Figure 2 the histogram of an estimate of $\sigma^2(X)$ produced by the random forest algorithm. More specifically, for each $i = 1, \dots, n$, we evaluate $\hat{\sigma}^2(X_i)$ by 10-fold cross-validation and build the histogram of $(\hat{\sigma}^2(X_i))_{i=1, \dots, n}$ thereafter. Left and right panels of Figure 2 deal respectively with the *aquatic* and *airfoil* datasets and reflect two different situations where the use of reject option is relevant. The estimated variance in the *airfoil* dataset is typically large (about 40% of the values are larger than 10) and then we may have some doubts in the associated prediction. According to the *aquatic* dataset, main part of the estimated values $\hat{\sigma}^2$ is smaller than 1 and then the use of the reject option may seem less significant. However, in this case, the predictions produced by the plug-in ε -predictors would be very accurate.

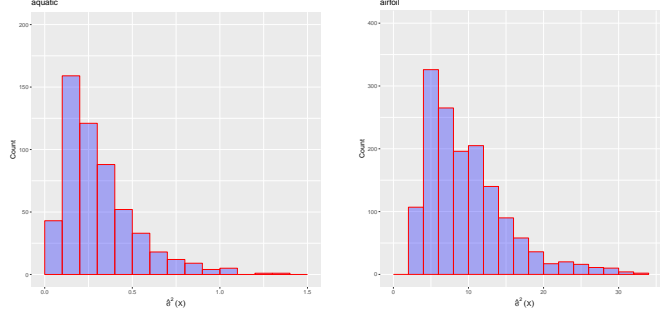


Figure 2: Histogram of the estimates of $\sigma^2(X)$

5.2 Results

We present the obtained results in Figure 3 and Table 1. We make a focus on the values of $\varepsilon \in \{0, 0.2, 0.5, 0.8\}$. As a general picture, the results are reflecting our theory: the empirical errors of the plug-in ε -predictors are decreasing *w.r.t.* ε for both datasets and their empirical rejection rates are very close to their expected values. Indeed, Table 1 displays how precise is the estimation of the rejection rate whatever the method used. This is in accordance with our theoretical findings. Moreover, the empirical errors of the plug-in ε -predictors based on the random forests and k NN algorithms are decreasing *w.r.t.* ε for both datasets. As expected, the use of the reject option improves the prediction precision. As an illustration, for *airfoil* dataset and the predictor based on random forests, the error is divided by 2 if we reject 50% of the data. However, we discover that the decrease for the prediction error is not systematic. In the case of plug-in ε -predictor based on the svm algorithm and with the *aquatic* dataset, we observe a strange curve for the error rate (see Figure 3-left). We conjecture that this phenomenon is due to a poor estimation of the variance. Indeed, in Figure 4, we present the performance of some kind of hybrid plug-in ε -predictors: we still use the svm algorithm to estimate the regression function; the variance function estimation is done based on svm (dashed line), random forests (dotted line), and k NN (dash-dotted line). From Figure 4, we observe that the empirical error $\widehat{\text{Err}}$ is now decreasing *w.r.t.* ε for the hybrid predictors based on svm and random forests, and that the performance is quite good.

Table 1: Performances of the three plug-in ε -predictors on the real datasets *aquatic*, and *airfoil*.

$1-\varepsilon$	aquatic						airfoil					
	svm		rf		knn		svm		rf		knn	
	$\widehat{\text{Err}}$	$1-\hat{r}$	$\widehat{\text{Err}}$	$1-\hat{r}$	$\widehat{\text{Err}}$	$1-\hat{r}$	$\widehat{\text{Err}}$	$1-\hat{r}$	$\widehat{\text{Err}}$	$1-\hat{r}$	$\widehat{\text{Err}}$	$1-\hat{r}$
1	1.38 (0.18)	1.00 (0.00)	1.34 (0.18)	1.00 (0.00)	2.29 (0.27)	1.00 (0.00)	11.81 (1.03)	1.00 (0.00)	14.40 (1.04)	1.00 (0.00)	35.40 (2.05)	1.00 (0.00)
0.8	1.08 (0.17)	0.81 (0.05)	1.04 (0.16)	0.80 (0.05)	1.98 (0.26)	0.80 (0.04)	8.27(0.86)	0.80 (0.03)	10.26 (0.95)	0.80 (0.03)	31.13 (1.96)	0.80 (0.03)
0.5	0.91 (0.18)	0.50 (0.06)	0.81 (0.18)	0.50 (0.06)	1.51 (0.30)	0.50 (0.06)	5.15 (0.92)	0.50 (0.04)	7.22 (0.92)	0.50 (0.3)	22.42 (2.13)	0.50 (0.03)
0.2	1.01 (0.32)	0.19 (0.05)	0.55 (0.21)	0.20 (0.05)	0.75 (0.37)	0.19 (0.05)	2.6 (0.64)	0.20 (0.03)	4.00 (0.74)	0.20 (0.03)	17.27 (3.00)	0.19 (0.03)

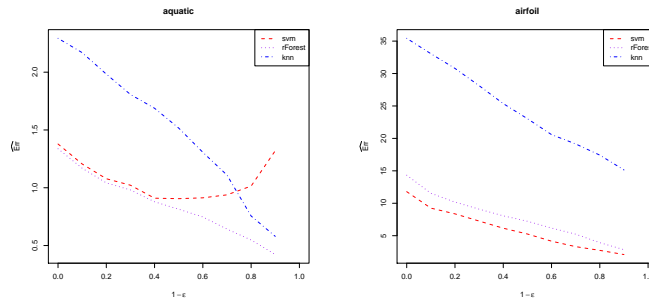


Figure 3: Visual description of the performance of three plug-in ε -predictors on the *aquatic*, and *airfoil* datasets.

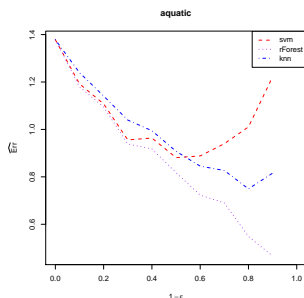


Figure 4: Additional description of the performance of the plug-in procedure on aquatic dataset.

6 Conclusion

We generalized the use of the reject option to the regression setting. We investigated the particular setting where the rejection rate is bounded. In this framework, an optimal rule is derived, it relies on thresholding of the variance function. Based on the *plug-in* principle, we derived a semi-supervised algorithm that can be applied on top of any off-the-shelf estimators of both regression and variance functions. One of the main features of the proposed procedure is that it precisely controls the probability of rejection. We derived general consistency results on rejection rate and on excess risk. We also established rates of convergence for the predictor with reject option when the regression and the variance functions are estimated by k NN algorithm. In future work, we plan to apply our methodology to the high-dimensional setting, taking advantage of sparsity structure of the data.

Broader impact

Approaches based on reject option may be helpful at least from two perspectives. First, when human action is limited by time or any other constraint, reject option is an efficient tool to prioritize the human action. On the other hand, in a world where automatic decisions need to be balanced and considered with caution, abstaining from prediction is one way to prevent from damageable decision-making. In particular, human is more likely able to detect anomalies such as bias in data. In a manner of speaking, the use of the reject option compromises between human and machines! Our numerical and theoretical analyses support this idea, in particular because our estimation of the rejection rate is accurate.

While the rejection rate has to be fixed according to the considered problem, it appears that the main drawback of our approach is that border instances may be automatically treated while they would have deserved a human consideration. From a general perspective, this is a weakness of all methods based on reject option. This inconvenience is even stronger when the conditional variance function is poorly estimated.

References

- [1] J.-Y. Audibert and A. Tsybakov. Fast learning rates for plug-in classifiers. *Annals of Statistics*, 35(2):608–633, 2007.
- [2] P. Bartlett and M. Wegkamp. Classification with a reject option using a hinge loss. *Journal of Machine Learning Research*, 9:1823–1840, 2008.
- [3] G. Biau and L. Devroye. *Lectures on the Nearest Neighbor Method*. Springer Series in the Data Sciences. Springer New York, 2015.
- [4] S. Bobkov and M. Ledoux. One-dimensional empirical measures, order statistics and Kantorovich transport distances. 2016. to appear in the *Memoirs of the American Mathematical Society*.
- [5] C. Chow. An optimum character recognition system using decision functions. *IRE Transactions on Electronic Computers*, (4):247–254, 1957.

- [6] C. Chow. On optimum error and reject trade-off. *IEEE Transactions on Information Theory*, 16:41–46, 1970.
- [7] C. Denis and M. Hebiri. Consistency of plug-in confidence sets for classification in semi-supervised learning. *Journal of Nonparametric Statistics*, 2019.
- [8] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1996.
- [9] J. Fan and Q. Yao. Efficient estimation of conditional variance functions in stochastic regression. *Biometrika*, 85(3):645–660, 1998.
- [10] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer Series in Statistics. Springer-Verlag, New York, 2002.
- [11] P. Hall and R.J. Carroll. Variance function estimation in regression: The mean effect of estimating the mean. *Journal of the Royal Statistical Society: Series B (Methodological)*, 51(1):3–14, 1989.
- [12] W. Härdle and A. Tsybakov. Local polynomial estimators of the volatility function in nonparametric autoregression. *Journal of Econometrics*, 81(1):223–242, 1997.
- [13] R. Herbei and M. Wegkamp. Classification with reject option. *The Canadian Journal of Statistics*, 34(4):709–721, 2006.
- [14] R. Kulik and C. Wichelhaus. Nonparametric conditional variance and error density estimation in regression models with dependent errors and predictors. *Electronic Journal of Statistics*, 5:856–898, 2011.
- [15] J. Lei. Classification with confidence. *Biometrika*, 101(4):755–769, 2014.
- [16] E. Mammen and A. Tsybakov. Smooth discrimination analysis. *Annals of Statistics*, 27(6):1808–1829, 1999.
- [17] P. Massart. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *Annals of Probability*, 18(3):1269–1283, 1990.
- [18] M. Naadeem, J.D. Zucker, and B. Hanczar. Accuracy-rejection curves (ARCs) for comparing classification methods with a reject option. In *MLSB*, pages 65–81, 2010.
- [19] W. Polonik. Measuring mass concentrations and estimating density contour clusters-an excess mass approach. *Annals of Statistics*, 23(3):855–881, 1995.
- [20] E. Scornet, G. Biau, and J.-P. Vert. Consistency of random forests. *Annals of Statistics*, 43(4):1716–1741, 08 2015.
- [21] Y. Shen, G. Gao, D. Witten, and F. Han. Optimal estimation of variance in nonparametric regression with random design. 2019.
- [22] C. Stone. Consistent nonparametric regression. *Annals of Statistics*, pages 595–620, 1977.
- [23] A. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer New York, 2008.
- [24] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- [25] V. Vovk, A. Gammerman, and G. Shafer. *Algorithmic learning in a random world*. Springer, New York, 2005.
- [26] Y. Wiener and R. El-Yaniv. Pointwise tracking the optimal regression function. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 2042–2050. Curran Associates, Inc., 2012.

Supplementary material

This supplementary material is organized as follows. Section A provides all proofs of results related to the optimal predictors (that is, Propositions 2.1, 2.3 2.4). In Sections B and C we prove Theorem 3.2 that establishes the consistency and Theorem 4.4 that states the rates of convergence of the plug-in ε -predictor $\hat{\Gamma}_\varepsilon$ respectively. We further establish several finite sample guarantees on k NN estimator in Section D. To help readability of the paper, we provide in Section E some technical tools that are used for the proofs.

A Proofs for optimal predictors

Proof of Proposition 2.1. By definition of \mathcal{R}_λ , we have for any predictor with reject option Γ_f

$$\begin{aligned}\mathcal{R}_\lambda(\Gamma_f) &= \mathbb{E}[(Y - f(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda \mathbb{P}(|\Gamma_f(X)| = 0) \\ &= \mathbb{E}[(Y - f^*(X) + f^*(X) - f(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda \mathbb{P}(|\Gamma_f(X)| = 0) \\ &= \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \mathbb{E}[(f^*(X) - f(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] \\ &\quad + 2\mathbb{E}[(Y - f^*(X))(f^*(X) - f(X)) \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda \mathbb{P}(|\Gamma_f(X)| = 0) .\end{aligned}$$

Since

$$\mathbb{E}[(Y - f^*(X))(f^*(X) - f(X)) \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] = 0 ,$$

and

$$\mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] = \mathbb{E}[\sigma^2(X) \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] ,$$

we deduce,

$$\begin{aligned}\mathcal{R}_\lambda(\Gamma_f) &= \mathbb{E}[(f^*(X) - f(X))^2 \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \mathbb{E}[\sigma^2(X) \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda(1 - \mathbb{1}_{\{|\Gamma_f(X)|=1\}}) \\ &= \mathbb{E}[\{(f^*(X) - f(X))^2 + (\sigma^2(X) - \lambda)\} \mathbf{1}_{\{|\Gamma_f(X)|=1\}}] + \lambda .\end{aligned}\tag{6}$$

Clearly, on the event $\{|\Gamma_f(X)| = 1\}$, the mapping $f \mapsto (f^*(X) - f(X))^2 + (\sigma^2(X) - \lambda)$ achieves its minimum at $f = f^*$. Then, it remains to consider the minimization of

$$\Gamma \mapsto \mathbb{E}[\{(\sigma^2(X) - \lambda)\} \mathbf{1}_{\{|\Gamma(X)|=1\}}] + \lambda ,$$

on the set Υ_{f^*} , which leads to $\{|\Gamma(X)| = 1\} = \{\sigma^2(X) \leq \lambda\}$. Putting all together, we get

$$\{|\Gamma_\lambda^*(X)| = 1\} = \{\sigma^2(X) \leq \lambda\} \text{ and on this event } \Gamma_\lambda^*(X) = \{f^*(X)\} ,$$

and point 1. of Proposition 2.1 is proven. For the second point, we observe that for $0 < \lambda < \lambda'$,

$$\{|\Gamma_\lambda^*(X)| = 1\} = \{\sigma^2(X) \leq \lambda\} \subset \{\sigma^2(X) \leq \lambda'\} = \{|\Gamma_{\lambda'}^*(X)| = 1\} .$$

From this inclusion, we deduce $r(\Gamma_{\lambda'}^*) \leq r(\Gamma_\lambda^*)$. Furthermore, using the relation $\{|\Gamma_\lambda^*(X)| = 1\} = \{\sigma^2(X) \leq \lambda\}$ and if we denote by $a_\lambda = \mathbb{P}(|\Gamma_\lambda^*(X)| = 1)$ we have

$$\begin{aligned}\text{Err}(\Gamma_\lambda^*) - \text{Err}(\Gamma_{\lambda'}^*) &= \frac{1}{a_\lambda} \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\sigma^2(X) \leq \lambda\}}] - \frac{1}{a_{\lambda'}} \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\sigma^2(X) \leq \lambda'\}}] \\ &= \left(\frac{1}{a_\lambda} - \frac{1}{a_{\lambda'}} \right) \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\sigma^2(X) \leq \lambda\}}] \\ &\quad - \frac{1}{a_{\lambda'}} \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\lambda < \sigma^2(X) \leq \lambda'\}}] .\end{aligned}\tag{7}$$

By definition of $\sigma^2(X)$, we can write

$$\begin{aligned}\mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\sigma^2(X) \leq \lambda\}}] &= \mathbb{E}[\mathbf{1}_{\{\sigma^2(X) \leq \lambda\}} \mathbb{E}[(Y - f^*(X))^2 | X]] \\ &= \mathbb{E}[\mathbf{1}_{\{\sigma^2(X) \leq \lambda\}} \sigma^2(X)] \leq \lambda a_\lambda ,\end{aligned}$$

and then

$$\left(\frac{1}{a_\lambda} - \frac{1}{a_{\lambda'}} \right) \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\sigma^2(X) \leq \lambda\}}] \leq \lambda \left(1 - \frac{a_\lambda}{a_{\lambda'}} \right) .$$

In the same way, we obtain

$$\frac{1}{a_{\lambda'}} \mathbb{E}[(Y - f^*(X))^2 \mathbf{1}_{\{\lambda < \sigma^2(X) \leq \lambda'\}}] \geq \frac{\lambda}{a_{\lambda'}} (a_{\lambda'} - a_\lambda) = \lambda \left(1 - \frac{a_\lambda}{a_{\lambda'}} \right) .$$

From Equation (7), we then get $\text{Err}(\Gamma_\lambda^*) \leq \text{Err}(\Gamma_{\lambda'}^*)$. \square

Proof of Proposition 2.3. First of all, observe that for any $\varepsilon \in (0, 1)$, if we set $\lambda_\varepsilon = F_{\sigma^2}^{-1}(1 - \varepsilon)$, then the optimal predictor Γ_λ^* given by (1) with $\lambda = \lambda_\varepsilon$ is such that,

$$r(\Gamma_{\lambda_\varepsilon}^*) = \mathbb{P}(\sigma^2(X) \geq \lambda_\varepsilon) = \mathbb{P}(F_{\sigma^2}(\sigma^2(X)) \geq 1 - \varepsilon) = \varepsilon .$$

We need to prove that any predictor Γ_f such that $r(\Gamma_f) = \varepsilon'$ with $\varepsilon' \leq \varepsilon$, satisfies $\text{Err}(\Gamma_f) \geq \text{Err}(\Gamma_{\lambda_\varepsilon}^*)$. To this end, consider $\Gamma_{\lambda_{\varepsilon'}}^*$ with $\lambda_{\varepsilon'} = F_{\sigma^2}^{-1}(1 - \varepsilon')$. On one hand, by optimality of $\Gamma_{\lambda_{\varepsilon'}}^*$ (cf. point 1. of Proposition 2.1), we have

$$\text{Err}(\Gamma_f) - \text{Err}(\Gamma_{\lambda_{\varepsilon'}}^*) = \frac{1}{1 - \varepsilon'} \left(\mathcal{R}_{\lambda_{\varepsilon'}}(\Gamma_f) - \mathcal{R}_{\lambda_{\varepsilon'}}(\Gamma_{\lambda_{\varepsilon'}}^*) \right) \geq 0 .$$

On the other hand, since $\varepsilon' \leq \varepsilon$ implies $\lambda_\varepsilon \leq \lambda_{\varepsilon'}$, point 2. of Proposition 2.1 reads as

$$\text{Err}(\Gamma_{\lambda_\varepsilon}^*) \leq \text{Err}(\Gamma_{\lambda_{\varepsilon'}}^*) .$$

Combining these two facts gives the desired result. \square

Proof of Proposition 2.4. First, from Equation (6), we have the following decomposition

$$\begin{aligned} \mathcal{R}_{\lambda_\varepsilon}(\Gamma_f) &= \mathbb{E} \left[\{(f^*(X) - f(X))^2 + \sigma^2(X) - \lambda_\varepsilon\} \mathbb{1}_{\{|\Gamma_f(X)|=1\}} \right] + \lambda_\varepsilon \\ &= \mathbb{E} \left[(f^*(X) - f(X))^2 \mathbb{1}_{\{|\Gamma_f(X)|=1\}} \right] + \mathbb{E} \left[(\sigma^2(X) - \lambda_\varepsilon) \mathbb{1}_{\{|\Gamma_f(X)|=1\}} \right] + \lambda_\varepsilon . \end{aligned}$$

Therefore, we deduce

$$\mathcal{E}(\Gamma_f) = \mathbb{E} \left[(f^*(X) - f(X))^2 \mathbb{1}_{\{|\Gamma_f(X)|=1\}} \right] + \mathbb{E} \left[(\sigma^2(X) - \lambda_\varepsilon) \left\{ \mathbb{1}_{\{|\Gamma_f(X)|=1\}} - \mathbb{1}_{\{|\Gamma_\varepsilon^*(X)|=1\}} \right\} \right],$$

and the result follows from the fact that all non zero values of $\mathbb{1}_{\{|\Gamma_f(X)|=1\}} - \mathbb{1}_{\{|\Gamma_\varepsilon^*(X)|=1\}}$ equal the sign of $(\sigma^2(X) - \lambda_\varepsilon)$ due to the fact that $\{|\Gamma_\varepsilon^*(X)|=1\} = \{\sigma^2(X) - \lambda_\varepsilon \leq 0\}$. \square

B Proof of the consistency results: Theorem 3.2

The consistency of $\hat{\Gamma}_\varepsilon$ consists in the introduction of a pseudo oracle ε -predictor $\tilde{\Gamma}_\varepsilon$ defined for all $x \in \mathbb{R}^d$ by

$$\tilde{\Gamma}_\varepsilon(x) = \begin{cases} \hat{f}(x) & \text{if } \hat{\sigma}^2(x) \leq F_{\hat{\sigma}^2}^{-1}(1 - \varepsilon) \\ \emptyset & \text{otherwise} . \end{cases} \quad (8)$$

This predictor differs from $\hat{\Gamma}_\varepsilon$ in that it knows the marginal distribution \mathbb{P}_X and then it has rejection rate exactly ε . Then, we consider the following decomposition

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) \right] = \mathbb{E} \left[\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] + \mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] , \quad (9)$$

and show that both terms in the r.h.s. go to zero.

• **Step 1.** $\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] \rightarrow 0$. We use Proposition 2.4 and get the following result.

Proposition B.1. *Let $\varepsilon \in (0, 1)$. Under Assumptions 2.2 and 3.1, the following holds*

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] \leq \mathbb{E} \left[(\hat{f}(X) - f^*(X))^2 \right] + \mathbb{E} \left[|\hat{\sigma}^2(X) - \sigma^2(X)| \right] + C \mathbb{E} \left[|F_{\hat{\sigma}^2}(\lambda_\varepsilon) - F_{\sigma^2}(\lambda_\varepsilon)| \right],$$

where $C > 0$ is constant which depends on the upper bounds of σ^2 and λ_ε .

Proof of Proposition B.1. Let $\varepsilon \in (0, 1)$. First, we recall our notation $F_{\hat{\sigma}^2}(\cdot) = \mathbb{P}_X(\hat{\sigma}^2(X) \leq \cdot | \mathcal{D}_n)$ and $\lambda_\varepsilon = F_{\sigma^2}^{-1}(1 - \varepsilon)$. We also introduce $\tilde{\lambda}_\varepsilon = F_{\hat{\sigma}^2}^{-1}(1 - \varepsilon)$ for the pseudo-oracle counterpart of λ_ε . A direct application of Proposition 2.4 yields

$$\mathcal{E}(\tilde{\Gamma}_\varepsilon) \leq \mathbb{E}_X \left[(\hat{f}(X) - f^*(X))^2 \right] + \mathbb{E}_X \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)| \neq |\Gamma_\varepsilon^*(X)|\}} \right]. \quad (10)$$

We first observe that if $\sigma^2(X) \leq \lambda_\varepsilon$ and $\hat{\sigma}^2(X) \geq \tilde{\lambda}_\varepsilon$, we have either of the two cases

- $\tilde{\lambda}_\varepsilon \geq \lambda_\varepsilon$ and then $|\sigma^2(X) - \lambda_\varepsilon| \leq |\hat{\sigma}^2(X) - \sigma^2(X)|$;
- $\tilde{\lambda}_\varepsilon \leq \lambda_\varepsilon$ and then either $|\sigma^2(X) - \lambda_\varepsilon| \leq |\hat{\sigma}^2(X) - \sigma^2(X)|$ or $\hat{\sigma}^2(X) \in (\tilde{\lambda}_\varepsilon, \lambda_\varepsilon)$.

Similar reasoning holds in the case where $\sigma^2(X) \geq \lambda_\varepsilon$ and $\hat{\sigma}^2(X) \leq \tilde{\lambda}_\varepsilon$. Therefore

$$\begin{aligned} & \mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)| \neq |\Gamma_\varepsilon^*(X)|\}} | \mathcal{D}_n \right] \\ & \leq \mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\sigma^2(X) - \lambda_\varepsilon| \leq |\hat{\sigma}^2(X) - \sigma^2(X)|\}} | \mathcal{D}_n \right] \\ & \quad + \mathbb{1}_{\{\lambda_\varepsilon \leq \tilde{\lambda}_\varepsilon\}} \mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{\lambda_\varepsilon \leq \hat{\sigma}^2(X) \leq \tilde{\lambda}_\varepsilon\}} | \mathcal{D}_n \right] \\ & \quad + \mathbb{1}_{\{\tilde{\lambda}_\varepsilon \leq \lambda_\varepsilon\}} \mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{\tilde{\lambda}_\varepsilon \leq \hat{\sigma}^2(X) \leq \lambda_\varepsilon\}} | \mathcal{D}_n \right] . \end{aligned}$$

From the above inequality, since σ^2 is bounded, there exists a constant $C > 0$ such that

$$\mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)| \neq |\Gamma_\varepsilon^*(X)|\}} \right] \leq \mathbb{E} \left[|\hat{\sigma}^2(X) - \sigma^2(X)| \right] + C \mathbb{E} \left[|F_{\hat{\sigma}^2}(\tilde{\lambda}_\varepsilon) - F_{\hat{\sigma}^2}(\lambda_\varepsilon)| \right] .$$

Now, from Assumptions 2.2 and 3.1, we have that $F_{\sigma^2}(\lambda_\varepsilon) = 1 - \varepsilon = F_{\hat{\sigma}^2}(\tilde{\lambda}_\varepsilon)$. Therefore, we deduce that

$$\mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)| \neq |\Gamma_\varepsilon^*(X)|\}} \right] \leq \mathbb{E} \left[|\hat{\sigma}^2(X) - \sigma^2(X)| \right] + C \mathbb{E} \left[|F_{\sigma^2}(\lambda_\varepsilon) - F_{\hat{\sigma}^2}(\lambda_\varepsilon)| \right] .$$

Putting this into Equation (10) gives the result in Proposition B.1. \square

Since \hat{f} and $\hat{\sigma}^2$ are consistent *w.r.t.* the L_2 and L_1 risks respectively, the first two terms in the bound of Proposition B.1 converge to zero. It remains to study the convergence of the last term. To this end, we prove that

$$\begin{aligned} \mathbb{E} \left[|F_{\sigma^2}(\lambda_\varepsilon) - F_{\hat{\sigma}^2}(\lambda_\varepsilon)| \right] &= \mathbb{E} \left[|\mathbb{1}_{\{\sigma^2(X) \leq \lambda_\varepsilon\}} - \mathbb{1}_{\{\hat{\sigma}^2(X) \leq \lambda_\varepsilon\}}| \right] \\ &\leq \mathbb{P} \left(|\sigma^2(X) - \hat{\sigma}^2(X)| \geq |\sigma^2(X) - \lambda_\varepsilon| \right) . \end{aligned}$$

Hence, for any $\beta > 0$, using Markov's Inequality we have

$$\begin{aligned} \mathbb{E} \left[|F_{\sigma^2}(\lambda_\varepsilon) - F_{\hat{\sigma}^2}(\lambda_\varepsilon)| \right] &\leq \mathbb{P} \left(|\sigma^2(X) - \lambda_\varepsilon| \leq \beta \right) + \mathbb{P} \left(|\sigma^2(X) - \hat{\sigma}^2(X)| \geq \beta \right) \\ &\leq \mathbb{P} \left(|\sigma^2(X) - \lambda_\varepsilon| \leq \beta \right) + \frac{\mathbb{E} \left[|\hat{\sigma}^2(X) - \sigma^2(X)| \right]}{\beta} . \end{aligned}$$

Combining this last inequality with Proposition B.1 and the consistency of \hat{f} and $\hat{\sigma}^2$ *w.r.t.* the L_2 and L_1 risks respectively implies that for all $\beta > 0$

$$\limsup_{n, N \rightarrow +\infty} \mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] \leq C \mathbb{P} \left(|\sigma^2(X) - \lambda_\varepsilon| \leq \beta \right) .$$

Since the above inequality holds for all $\beta > 0$, under Assumption 2.2, we deduce that

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] \rightarrow 0 ,$$

and then this step of the proof is complete.

• **Step 2.** $\mathbb{E} \left[\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] \rightarrow 0$. Thanks to Equation (6), we have that

$$\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) = \mathbb{E}_X \left[\left\{ (f^*(X) - \hat{f}(X))^2 + (\sigma^2(X) - \lambda_\varepsilon) \right\} \left(\mathbb{1}_{\{|\hat{\Gamma}_\varepsilon(X)|=1\}} - \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)|=1\}} \right) \right] .$$

Therefore, since σ^2 is bounded, there exists a constant $C > 0$ such that

$$\begin{aligned} \mathbb{E} \left[|\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon)| \right] &\leq 2 \mathbb{E} \left[(f^*(X) - \hat{f}(X))^2 \right] + C \mathbb{E} \left[|\mathbb{1}_{\{|\hat{\Gamma}_\varepsilon(X)|=1\}} - \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)|=1\}}| \right] \\ &\leq 2 \mathbb{E} \left[(f^*(X) - \hat{f}(X))^2 \right] + C A_\varepsilon , \end{aligned} \tag{11}$$

where

$$A_\varepsilon = \mathbb{E} \left[\left| \mathbb{1}_{\{|\hat{\Gamma}_\varepsilon(X)|=1\}} - \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X)|=1\}} \right| \right] = \mathbb{E} \left[\left| \mathbb{1}_{\{\hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \geq 1-\varepsilon\}} - \mathbb{1}_{\{F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \geq 1-\varepsilon\}} \right| \right] . \quad (12)$$

Considering the fact that \hat{f} is consistent *w.r.t.* the L_2 risk, it remains to treat the term A_ε . We have

$$A_\varepsilon \leq \mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \right| \geq |F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - (1-\varepsilon)| \right) ,$$

and then, for all $\beta > 0$, the following holds

$$A_\varepsilon \leq \mathbb{P} \left(|F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - (1-\varepsilon)| < \beta \right) + \mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \right| \geq \beta \right) . \quad (13)$$

Under Assumption 3.1, the random variable $F_{\hat{\sigma}^2}(\hat{\sigma}^2(X))$ is uniformly distributed on $[0, 1]$ conditionally on \mathcal{D}_n . Therefore, we deduce that

$$\begin{aligned} \mathbb{P} \left(|F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - (1-\varepsilon)| < \beta \right) &= \mathbb{E} \left[\mathbb{P}_X \left(|F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - (1-\varepsilon)| < \beta \mid \mathcal{D}_n \right) \right] \\ &= \mathbb{E} [2\beta \mid \mathcal{D}_n] = 2\beta . \end{aligned} \quad (14)$$

According to the second term in the r.h.s. of Equation (13), we have that

$$\begin{aligned} \mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \right| \geq \beta \right) &\leq \mathbb{P} \left(\sup_{x \in \mathbb{R}} \left| \hat{F}_{\hat{\sigma}^2}(x) - F_{\hat{\sigma}^2}(x) \right| \geq \beta \right) \\ &= \mathbb{E} \left[\mathbb{P}_{\mathcal{D}_N} \left(\sup_{x \in \mathbb{R}} \left| \hat{F}_{\hat{\sigma}^2}(x) - F_{\hat{\sigma}^2}(x) \right| \geq \beta \mid \mathcal{D}_n \right) \right] , \end{aligned}$$

where $\mathbb{P}_{\mathcal{D}_N}$ is the probability measure *w.r.t.* the dataset \mathcal{D}_N . Since, conditionally on \mathcal{D}_n , $\hat{F}_{\hat{\sigma}^2}$ is the empirical counterpart of the continuous cumulative distribution function $F_{\hat{\sigma}^2}$, applying the Dvoretzky-Kiefer-Wolfowitz Inequality [17], we deduce that

$$\mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \right| \geq \beta \right) \leq 2 \exp(-2N\beta^2) . \quad (15)$$

Putting (14) and (15) into Eq. (13), we have that for all $\beta > 0$

$$A_\varepsilon \leq 2 \left(\beta + \exp(-2N\beta^2) \right) . \quad (16)$$

Since Equation (16) holds for all $\beta > 0$, we have that $A_\varepsilon \rightarrow 0$ as $N, n \rightarrow +\infty$. Hence, from the above inequality we get the desired result in **Step 2**:

$$\mathbb{E} \left[\left| \mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right| \right] \rightarrow 0 .$$

Combining **Step 1** and **Step 2** yields the convergence: $\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) \right] \rightarrow 0$ as $N, n \rightarrow +\infty$.

• **Bound on $\mathbb{E} \left[r(\hat{\Gamma}_\varepsilon) \right]$.** To finish the proof of Theorem 3.2, it remains to control the rejection rate $\mathbb{E} \left[r(\hat{\Gamma}_\varepsilon) \right]$ and show that it satisfies $\mathbb{E} \left[\left| r(\hat{\Gamma}_\varepsilon) - \varepsilon \right| \right] \leq CN^{-1/2}$ for some constant $C > 0$. We observe that

$$\mathbb{E} \left[\left| r(\hat{\Gamma}_\varepsilon) - \varepsilon \right| \right] = \mathbb{E} \left[\left| r(\hat{\Gamma}_\varepsilon) - r(\tilde{\Gamma}_\varepsilon) \right| \right] \leq A_\varepsilon ,$$

where A_ε is given by Eq. (12). Repeating the same reasoning as in **Step 2** above, we bound A_ε as in Eq. (13), and get from Dvoretzky-Kiefer-Wolfowitz Inequality (see Equation (15)), that for all $\beta > 0$,

$$\mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) \right| \geq \beta \right) \leq 2 \exp(-2N\beta^2) ,$$

and from Equation (14),

$$\mathbb{P} \left(|F_{\hat{\sigma}^2}(\hat{\sigma}^2(X)) - (1-\varepsilon)| < \beta \right) = 2\beta .$$

These two bounds combined the classical peeling argument of [1] (see Lemma E.1 below) imply the desired result:

$$A_\varepsilon \leq CN^{-1/2} . \quad (17)$$

C Proof of rates of convergence: Theorem 4.4

In this section, we follow the same strategy as in Section B but here we care about rates of convergence. Moreover, we have to pay attention to the randomness we introduced in the predictor because of the use of k NN. As in Section B, we introduce some pseudo-oracle predictor. However, this one needs to depend on the randomness we introduced in the definition of $\hat{\Gamma}_\varepsilon(x, \zeta)$. Define the pseudo-oracle ε -predictor $\tilde{\Gamma}_\varepsilon$ for all $x \in \mathbb{R}^d$ and $\zeta \in [0, u]$ as³

$$\tilde{\Gamma}_\varepsilon(x, \zeta) = \begin{cases} \hat{f}(x) & \text{if } \bar{\sigma}^2(x, \zeta) \leq F_{\bar{\sigma}^2}^{-1}(1 - \varepsilon) \\ \emptyset & \text{otherwise .} \end{cases}$$

To study the excess risk $\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\hat{\Gamma}_\varepsilon \right) \right]$ of our predictor, we also consider a similar decomposition as in Eq. (9) and treat each of the two terms separately.

• **Step 1.** Study of $\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right]$. We establish the following result.

Proposition C.1. *Assume that Assumptions 4.2 and 4.3 are fulfilled for some $\alpha \geq 0$, then the following inequality holds*

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] \leq \mathbb{E} \left[\left(f^*(X) - \hat{f}(X) \right)^2 \right] + C \left(\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| \right)^{1+\alpha} \right] + u^{1+\alpha} \right) ,$$

where $C > 0$ depends only on c^* and α .

Proof. Let $\varepsilon \in (0, 1)$. We recall that $\lambda_\varepsilon = F_{\sigma^2}^{-1}(1 - \varepsilon)$ and $\tilde{\lambda}_\varepsilon = F_{\bar{\sigma}^2}^{-1}(1 - \varepsilon)$. Since ζ is distributed according to a Uniform distribution on $[0, u]$, we observe that

$$|\bar{\sigma}^2(X, \zeta) - \sigma^2(X)| \leq \sup_{x \in \mathcal{C}} |\sigma^2(x) - \hat{\sigma}^2(x)| + u := \hat{h}_u .$$

Hence, according to Theorem 2.12 in [4] (recalled in Lemma E.2), we have that conditionally on \mathcal{D}_n

$$|\tilde{\lambda}_\varepsilon - \lambda_\varepsilon| \leq \hat{h}_u .$$

Furthermore, since X and ζ are independent, we can use Proposition 2.4 and get

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] \leq \mathbb{E} \left[\left(\hat{f}(X) - f^*(X) \right)^2 \right] + \mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X, \zeta)| \neq |\Gamma_\varepsilon^*(X)|\}} \right] .$$

On the event $\left\{ |\tilde{\Gamma}_\varepsilon(X, \zeta)| \neq |\Gamma_\varepsilon^*(X)| \right\}$, we note that

$$|\sigma^2(X) - \lambda_\varepsilon| \leq |\bar{\sigma}^2(X, \zeta) - \sigma^2(X)| + |\tilde{\lambda}_\varepsilon - \lambda_\varepsilon| .$$

Therefore, conditional on \mathcal{D}_n , we deduce the following

$$\begin{aligned} \mathbb{E}_{(X, \zeta)} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X, \zeta)| \neq |\Gamma_\varepsilon^*(X)|\}} \right] &\leq \mathbb{E}_{(X, \zeta)} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\sigma^2(X) - \lambda_\varepsilon| \leq |\bar{\sigma}^2(X, \zeta) - \sigma^2(X)| + |\tilde{\lambda}_\varepsilon - \lambda_\varepsilon|\}} \right] \\ &\leq \mathbb{E}_X \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\sigma^2(X) - \lambda_\varepsilon| \leq 2\hat{h}_u\}} \right] \\ &\leq 2\hat{h}_u \mathbb{P}_X \left(|\sigma^2(X) - \lambda_\varepsilon| \leq 2\hat{h}_u \right) . \end{aligned}$$

Finally, applying Assumption 4.3, we deduce that there exists a constant $C > 0$ such that

$$\mathbb{E} \left[|\sigma^2(X) - \lambda_\varepsilon| \mathbb{1}_{\{|\tilde{\Gamma}_\varepsilon(X, \zeta)| \neq |\Gamma_\varepsilon^*(X)|\}} \right] \leq C \left(\mathbb{E} \left[\sup_{x \in \mathcal{C}} |\sigma^2(x) - \hat{\sigma}^2(x)|^{1+\alpha} \right] + u^{1+\alpha} \right) ,$$

which ends the proof. \square

³The only difference between $\tilde{\Gamma}_\varepsilon(x, \zeta)$ and $\tilde{\Gamma}_\varepsilon(x)$ given in (8) is the dependency in ζ that is hidden inside $\bar{\sigma}^2$. To avoid useless additional notation, we write $\tilde{\Gamma}_\varepsilon$ for both pseudo-oracles.

Based on Proposition C.1, the control of $\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right]$ requires a bound on $\mathbb{E} \left[\left(f^*(X) - \hat{f}(X) \right)^2 \right]$ and on $\mathbb{E} \left[\sup_{x \in \mathcal{C}} \left| \sigma^2(x) - \hat{\sigma}^2(x) \right|^{1+\alpha} \right]$. The first of these two terms relies on estimation of the regression function with k NN algorithm and is rather well studied. In particular, thanks to Proposition D.4 we have with the choice $k_n \propto n^{2/(d+2)}$

$$\mathbb{E} \left[\left(\hat{f}(X) - f^*(X) \right)^2 \right] \leq C n^{-2/(d+2)} , \quad (18)$$

where $C > 0$ is a constant which depends on f^* , c_0 , \mathcal{C} , and d . Then it remains to bound the second term which is the purpose of Proposition D.5 that relies on the rate of convergence of the k NN estimator of the conditional variance $\hat{\sigma}^2$ in supremum norm. This result says that under our assumptions and for the choice $k_n \propto n^{2/(d+2)}$, we have that

$$\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \hat{\sigma}^2(x) - \sigma^2(x) \right| \right)^{1+\alpha} \right] \leq C \log(n)^{(\alpha+1)} n^{-(\alpha+1)/(d+2)} ,$$

for a constant $C > 0$ that depends on f^* , σ^2 , c_0 , \mathcal{C} , and on the dimension d . Putting this last inequality and Eq. (18) into the upper bound on the excess risk of $\tilde{\Gamma}_\varepsilon$ from Proposition C.1 we show that when we set $u = u_n \leq n^{-1/(d+2)}$ we can write

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] \leq C \left(n^{-2/(d+2)} + \log(n)^{(\alpha+1)} n^{-(\alpha+1)/(d+2)} \right) ,$$

where $C > 0$ is a constant which depends on f^* , σ^2 , c_0 , c^* , α , \mathcal{C} , and on the dimension d . This ends the first step of the proof.

• **Step 2.** Study of $\mathbb{E} \left[\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right]$. Since X and ζ are independent, as in **Step 2** of the proof of Theorem 3.2 (cf. Eq. (11)), we get

$$\mathbb{E} \left[\left| \mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right| \right] \leq 2 \mathbb{E} \left[\left(f^*(X) - \hat{f}(X) \right)^2 \right] + C A_\varepsilon ,$$

where A_ε is defined similarly as in Equation (12) with a small modification due to the random perturbation we made on $\hat{\sigma}^2$. Similarly we have

$$A_\varepsilon \leq \mathbb{P} \left(\left| \hat{F}_{\hat{\sigma}^2}(\hat{\sigma}^2(X, \zeta)) - \{F_{\hat{\sigma}^2}(\hat{\sigma}^2(X, \zeta))\} \right| \geq \left| F_{\hat{\sigma}^2}(\hat{\sigma}^2(X, \zeta)) - (1 - \varepsilon) \right| \right) .$$

Therefore using the same arguments as in **Step 2** of the proof of Theorem 3.2 to get (17), it is easy to see that there exists $C > 0$ such that $A_\varepsilon \leq C N^{-1/2}$. Then, we deduce

$$\mathbb{E} \left[\left| \mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right| \right] \leq 2 \mathbb{E} \left[\left(f^*(X) - \hat{f}(X) \right)^2 \right] + C N^{-1/2} .$$

Finally, an application of Theorem D.4 yields

$$\mathbb{E} \left[\left| \mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right| \right] \leq C \left(n^{-2/(d+2)} + N^{-1/2} \right) ,$$

where $C > 0$ is a constant which depends on f^* , c_0 , \mathcal{C} , and d . This ends **Step 2** of the proof.

Lastly, we combine the results in **Step 1** and **Step 2**, together with the decomposition

$$\mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\hat{\Gamma}_\varepsilon \right) \right] = \mathbb{E} \left[\mathcal{R}_{\lambda_\varepsilon}(\hat{\Gamma}_\varepsilon) - \mathcal{R}_{\lambda_\varepsilon}(\tilde{\Gamma}_\varepsilon) \right] + \mathbb{E} \left[\mathcal{E}_{\lambda_\varepsilon} \left(\tilde{\Gamma}_\varepsilon \right) \right] ,$$

and get the desired bound on the excess risk.

D Rate of convergence for k NN estimator

In this section, we focus on rates of convergence of k NN for the estimation of the regression function f^* and the conditional variance function σ^2 . The proofs techniques are largely inspired by those in [3, 10], though we provide some additional steps to build for instance finite sample bounds for the sup norm in the problem of conditional variance estimation.

D.1 Regression function estimation

We provide the rate of convergence of the k NN estimator of f^* in the regression model for which we make the following assumptions. We assume that f^* is Lipschitz (Assumption 4.1) and that Assumption 4.2 are fulfilled. We recall that from Assumption 4.2, we have that \mathbb{P}_X is supported on a compact set \mathcal{C} . Furthermore, we also assume that $Y - f^*(X)$ satisfies a uniform noise condition: there exists $c_0 > 0$ such that

$$\sup_{x \in \mathcal{C}} \mathbb{E} [\exp(\lambda (Y - f^*(X))) \mid X = x] \leq \exp(c_0^2 \lambda^2), \quad \text{for } |\lambda| \leq \frac{1}{c_0}. \quad (19)$$

This assumption is rather weak and requires that conditional on X is sub-exponential uniformly over \mathcal{C} (see [24]). Using the same notation as in Section 3, we recall that the k NN estimator \hat{f} of f is defined as follows

$$\hat{f}(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x).$$

The purpose of the appendix is to provide rates of convergence for the k NN estimator \hat{f} under the above assumption. To this end we require two auxiliary lemmata, which provide a control respectively with high probability and in expectation on the distance between a feature point and its neighbors uniformly over \mathcal{C} .

Lemma D.1. *Assume Assumptions 4.1-4.2 hold. Then there exist $C_1 > 0$, which depends only on \mathcal{C} , μ_{\min} , and on d and $C_2 > 0$, which depends on \mathcal{C} and on d , such that for all $t \geq \left(\frac{\log(n)k_n}{C_1 n}\right)^{1/d}$, we have*

$$\mathbb{P} \left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \geq t \right) \leq C_2 \exp(\log(n) - C_1 t^d n / k_n).$$

Proof. For any $a \in \mathbb{R}$, let us denote by $\lfloor a \rfloor$ the largest integer which is smaller or equal to a . Consider some $x \in \mathcal{C}$. Following the same arguments as in proof of Theorem 6.2 in [10], we split the data X_1, \dots, X_n into $k_n + 1$ folds such that the first k_n folds have the same size $\lfloor \frac{n}{k_n} \rfloor$ and the last fold contains the remaining data if there are. We denote \tilde{X}_j^x the nearest neighbor of x in the j th fold and then obviously

$$\sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \leq \sum_{j=1}^{k_n} \|\tilde{X}_j^x - x\|.$$

Let $\bar{B}_2(a, r)$ be the closed Euclidean ball in \mathbb{R}^d centered in a with radius $r > 0$. Since \mathcal{C} is compact, we have $\mathcal{C} \subset \bar{B}_2(0, R)$ for some $R > 0$, and therefore there exists an ε -net \mathcal{C}_ε of \mathcal{C} w.r.t. $\|\cdot\|$ such that $|\mathcal{C}_\varepsilon| \leq \left(\frac{3R}{\varepsilon}\right)^d$. In particular, for all $x \in \mathcal{C}$ there exists $x_\varepsilon \in \mathcal{C}_\varepsilon$ such that $\|x - x_\varepsilon\| \leq \varepsilon$. Then, for all $x \in \mathcal{C}$ and all $j \in \{1, \dots, k_n\}$, there exists $x_\varepsilon \in \mathcal{C}_\varepsilon$ such that

$$\|\tilde{X}_j^x - x\| \leq \|\tilde{X}_j^{x_\varepsilon} - x_\varepsilon\| + \varepsilon. \quad (20)$$

Besides, we observe that

$$\|\tilde{X}_j^{x_\varepsilon} - x_\varepsilon\| \leq \|\tilde{X}_j^{x_\varepsilon} - x_\varepsilon\| + 2\varepsilon. \quad (21)$$

Indeed, if $\|\tilde{X}_j^{x_\varepsilon} - x_\varepsilon\| + 2\varepsilon < \|\tilde{X}_j^x - x_\varepsilon\|$ we can write

$$\begin{aligned} \|\tilde{X}_j^{x_\varepsilon} - x\| + \varepsilon &\leq \|\tilde{X}_j^{x_\varepsilon} - x_\varepsilon\| + 2\varepsilon \\ &< \|\tilde{X}_j^x - x_\varepsilon\| \leq \|\tilde{X}_j^x - x\| + \varepsilon, \end{aligned}$$

which contradicts the fact that \tilde{X}_j^x is the nearest neighbor of x in the j th fold. Hence, from Equations (20) and (21), we deduce that

$$\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \leq 3\varepsilon + \sup_{x \in \mathcal{C}_\varepsilon} \frac{1}{k_n} \sum_{j=1}^{k_n} \|\tilde{X}_j^x - x\|.$$

From the above inequality, we obtain that for $t > 6\varepsilon$,

$$\mathbb{P} \left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \geq t \right) \leq \mathbb{P} \left(\sup_{x \in \mathcal{C}_\varepsilon} \frac{1}{k_n} \sum_{j=1}^{k_n} \|\tilde{X}_j^x - x\| \geq t/2 \right). \quad (22)$$

Our goal becomes to bound r.h.s. of the above inequality. Using union bound, we deduce that for all $t > 6\varepsilon$

$$\begin{aligned} \mathbb{P} \left(\sup_{x \in \mathcal{C}_\varepsilon} \frac{1}{k_n} \sum_{j=1}^{k_n} \|\tilde{X}_j^x - x\| \geq t/2 \right) &\leq \sum_{x \in \mathcal{C}_\varepsilon} \mathbb{P} \left(\frac{1}{k_n} \sum_{j=1}^{k_n} \|\tilde{X}_j^x - x\| \geq t/2 \right) \\ &\leq \sum_{x \in \mathcal{C}_\varepsilon} \sum_{j=1}^{k_n} \mathbb{P} \left(\|\tilde{X}_j^x - x\| \geq t/2 \right). \end{aligned} \quad (23)$$

For each $x \in \mathcal{C}_\varepsilon$ and $j \in \{1, \dots, k_n\}$, by definition of \tilde{X}_j^x and since $(X_i)_{i=1, \dots, n}$ are i.i.d., we have

$$\mathbb{P} \left(\|\tilde{X}_j^x - x\| \geq t/2 \right) = \mathbb{P} \left(\|X_{(1, \lfloor n/k_n \rfloor)} - x\| \geq t/2 \right) = (\mathbb{P}(\|X_1 - x\| \geq t/2))^{\lfloor n/k_n \rfloor}. \quad (24)$$

On one hand, observe that for $t \geq 4R$, $(\mathbb{P}(\|X_1 - x\| \geq t/2) = 0$. On the other hand for $t \leq 4R$, using the elementary inequality $\log(1 - a) \leq -a$ for all $a \in [0, 1)$, we have that

$$(\mathbb{P}(\|X_1 - x\| \geq t/2))^{\lfloor n/k_n \rfloor} \leq \exp \left(- \left\lfloor \frac{n}{k_n} \right\rfloor \mathbb{P}(\|X_1 - x\| \leq t/2) \right),$$

which yields, thanks to Assumption 4.2, there exists $C > 0$ which depends on μ_{\min} and d such that

$$(\mathbb{P}(\|X_1 - x\| \geq t/2))^{\lfloor n/k_n \rfloor} \leq \exp(-Ct^d \lfloor n/k_n \rfloor).$$

We finally deduce from Equation (22), (23), and (24), that for all $t \geq 6\varepsilon$

$$\mathbb{P} \left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \geq t \right) \leq k_n |\mathcal{C}_\varepsilon| \exp(-Ct^d \lfloor n/k_n \rfloor).$$

Choosing $\varepsilon = (\frac{k_n}{6^d C n})^{1/d}$, we get that for $t \geq (\frac{k_n}{C n})^{1/d}$,

$$\mathbb{P} \left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \geq t \right) \leq C_2 \exp(\log(n) - C_1 t^d n/k_n),$$

which yields the expected result. \square

The second lemma establishes a control in expectation of the uniform distance.

Lemma D.2. *Under Assumption 4.2, there exist $C > 0$, which depends only on \mathcal{C} , μ_{\min} , and on d such that*

$$\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \right)^p \right] \leq C \left(\frac{k_n \log(n)}{n} \right)^{p/d}.$$

Proof. Since Assumption 4.2 holds, we can use Lemma D.1. Then there exist two non negative constants C_1 and C_2 such that for all $t \geq (\frac{\log(n)k_n}{C_1 n})^{1/d}$, we have

$$\mathbb{P} \left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \geq t \right) \leq C_2 \exp(\log(n) - C_1 t^d n/k_n).$$

Therefore an application of Lemma E.3 implies directly the result. \square

Below, we state the main result of this section related to the rate of convergence in sup norm of the k NN estimator of the regression function.

Theorem D.3. Assume Assumption 4.2 is satisfied. Moreover, let $p \geq 1$ and $k_n \propto n^{2/(d+2)}$. Then

$$\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \hat{f}(x) - f^*(x) \right| \right)^p \right] \leq C \log(n) p n^{-p/(d+2)},$$

where $C > 0$ is a constant which depends on f^* , c_0 , \mathcal{C} , μ_{\min} and d .

Proof. First, we have that

$$\hat{f}(x) - f^*(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} (Y_{(i,n)}(x) - f^*(X_{(i,n)}(x))) + \frac{1}{k_n} \sum_{i=1}^{k_n} (f^*(X_{(i,n)}(x)) - f^*(x)) .$$

Therefore, since f^* is L -Lipschitz, we then deduce that

$$\sup_{x \in \mathcal{C}} \left| \hat{f}(x) - f^*(x) \right| \leq \sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x) - f^*(X_{(i,n)}(x)) \right| + L \sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| ,$$

which implies that

$$\begin{aligned} \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \hat{f}(x) - f^*(x) \right| \right)^p \right] &\leq 2^{p-1} \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x) - f^*(X_{(i,n)}(x)) \right| \right)^p \right] \\ &\quad + 2^{p-1} L^p \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \frac{1}{k_n} \sum_{i=1}^{k_n} \|X_{(i,n)}(x) - x\| \right)^p \right] . \end{aligned} \quad (25)$$

Lemma D.2 provides a bound on the second term in the r.h.s. of the above inequality. Then it remains to study the first term in the r.h.s. of Eq. (25). Let $x \in \mathcal{C}$, and denote by $\mathcal{N}_{k_n}(x) = \{X_{(1,n)}(x), \dots, X_{(k_n,n)}(x)\}$ the set of the k_n -nearest neighbors of x among $\{X_1, \dots, X_n\}$. We denote by \mathcal{B} the set of all closed balls in \mathbb{R}^d . We observe that there exists $\rho_x > 0$ such that $\mathcal{N}_{k_n}(x) \subset \{\bar{B}(x, \rho_x) \cap \{X_1, \dots, X_n\}\}$, where $\bar{B}(x, \rho_x)$ is the closed ball centered on x with radius ρ_x . Therefore

$$\{\mathcal{N}_{k_n}(x), x \in \mathcal{C}\} \subset \{\{X_1, \dots, X_n\} \cap B, B \in \mathcal{B}\} .$$

Besides, since the VC-dimension of the class of balls in \mathbb{R}^d is upper bounded by $d + 2$ (see for instance Corollary 13.2 in [8]), Sauer Lemma implies that

$$|\{\{X_1, \dots, X_n\} \cap B, B \in \mathcal{B}\}| \leq \mathcal{S}(\mathcal{B}, n) \leq (n + 1)^{d+2} ,$$

where $\mathcal{S}(\mathcal{B}, n)$ denotes the shatter coefficient of \mathcal{B} by n points from \mathcal{C} . We then deduce that $|\{\mathcal{N}_{k_n}(x), x \in \mathcal{C}\}| \leq (n + 1)^{d+2}$, which implies in turn that there exists $\{x_1, \dots, x_J\}$, with $J \leq (n + 1)^{d+2}$ such that

$$\begin{aligned} \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x) - f^*(X_{(i,n)}(x)) \right| \right)^p \right] \\ \leq \mathbb{E} \left[\left(\max_{j \in \{1, \dots, J\}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \right)^p \right] . \end{aligned}$$

Notice that conditional on X_1, \dots, X_n the random variables $(Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)))_{i=1, \dots, k_n}$ are independent with zero mean (see Proposition 8.1 in [3]). Besides from Equation (19) they are uniformly sub-exponential over \mathcal{C} , then we deduce from the Bernstein Inequality (see [24]) that for all $t \geq 0$ and $j = 1, \dots, J$,

$$\mathbb{P} \left(\left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \geq t \right) \leq \exp \left(-ck_n \min \left(\frac{t^2}{K^2}, \frac{t}{K} \right) \right) ,$$

where $c > 0$ is an absolute constant and $K > 0$ depends on c_0 in Eq. (19). Set $v_n = \sqrt{\frac{(d+2) \log(n+1)}{ck_n}}$. Our choice of k_n ensures that $v_n \leq 1$, and then we deduce from the union bound that for $t \in (Kv_n, K)$,

$$\mathbb{P} \left(\max_{j \in \{1, \dots, J\}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \geq t \right) \leq \exp \left((d + 2) \log(n + 1) - ck_n t^2 / K^2 \right) ,$$

and for $t > K$,

$$\mathbb{P} \left(\max_{j \in \{1, \dots, J\}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \geq t \right) \leq \exp((d+2) \log(n+1) - ck_n t/K) .$$

Considering these two cases, we can derive an exponential bound on the term $\mathbb{P} \left(\max_{j \in \{1, \dots, J\}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \geq t \right)$ for all $t \geq K v_n$, therefore we can use similar arguments as in Lemma E.3 and conclude that

$$\begin{aligned} & \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x) - f^*(X_{(i,n)}(x)) \right| \right)^p \right] \\ & \leq \mathbb{E} \left[\left(\max_{j \in \{1, \dots, J\}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} Y_{(i,n)}(x_j) - f^*(X_{(i,n)}(x_j)) \right| \right)^p \right] \leq C \left(\frac{\log(n)}{k_n} \right)^{p/2} . \end{aligned} \quad (26)$$

Combining the above inequality, Equation (25), and Lemma D.2, gives the desired result. \square

To conclude this section, we also provide the rate of convergence of the k NN estimator in L_2 -norm

Theorem D.4. *Assume Assumption 4.2 is satisfied and let $k_n \propto n^{2/(d+2)}$, then*

$$\mathbb{E} \left[\left(\hat{f}(X) - f^*(X) \right)^2 \right] \leq C n^{-2/(d+2)} ,$$

where $C > 0$ is a constant which depends on f^* , c_0 , \mathcal{C} , and d .

The proof of this result is provided in [10] for $d \geq 3$ (see Theorem 6.2). However, a small change implies that the same proof holds for all d under Assumption 4.2.

D.2 Conditional variance function estimation

We provide the rate of convergence of the k NN estimator of σ^2 . This proof is largely inspired by [3], though we are interested here in finite sample bounds.

Proposition D.5. *Grant Assumptions 4.1 and 4.2. Let $k_n \propto n^{2/(d+2)}$, the following holds*

$$\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| \right)^{1+\alpha} \right] \leq C \log(n)^{(\alpha+1)} n^{-(\alpha+1)/(d+2)} ,$$

for all $\alpha \geq 0$, where $C > 0$ is a constant which depends on f^* , σ^2 , c_0 , \mathcal{C} , and on the dimension d .

Proof. First, we define the function $\tilde{\sigma}^2$ by

$$\tilde{\sigma}^2(x) = \frac{1}{k_n} \sum_{i=1}^{k_n} (Y_{(i,n)}(x) - f^*(X_{(i,n)}(x)))^2, \quad \forall x \in \mathbb{R}^d .$$

The function $\tilde{\sigma}^2$ is the pseudo-estimator of σ^2 that would be used in the case where the function f^* is known. By the triangle inequality, we have that for all $x \in \mathcal{C}$,

$$|\hat{\sigma}^2(x) - \sigma^2(x)| \leq |\hat{\sigma}^2(x) - \tilde{\sigma}^2(x)| + |\tilde{\sigma}^2(x) - \sigma^2(x)| .$$

Now, we observe that

$$\begin{aligned} & \hat{\sigma}^2(x) - \tilde{\sigma}^2(x) = \\ & \frac{1}{k_n} \sum_{i=1}^{k_n} \left(f^*(X_{(i,n)}(x)) - \hat{f}(X_{(i,n)}(x)) \right) \left(2(Y_{(i,n)}(x) - f^*(X_{(i,n)}(x))) + f^*(X_{(i,n)}(x)) - \hat{f}(X_{(i,n)}(x)) \right) . \end{aligned}$$

Therefore, we deduce

$$\begin{aligned} \sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| &\leq \sup_{x \in \mathcal{C}} |\tilde{\sigma}^2(x) - \sigma^2(x)| + \left(\sup_{x \in \mathcal{C}} |\hat{f}(x) - f^*(x)| \right)^2 + \\ &\quad 2 \sup_{x \in \mathcal{C}} |\hat{f}(x) - f^*(x)| \sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} (Y_{(i,n)}(x) - f^*(X_{(i,n)}(x))) \right|. \end{aligned}$$

From the above inequality, using the fact that $(a+b+c)^p \leq 3^{p-1}(a^p+b^p+c^p)$ for $p \geq 1$, $a, b, c \in \mathbb{R}$ and applying the Cauchy-Schwartz Inequality, we obtain

$$\begin{aligned} \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| \right)^{1+\alpha} \right] &\leq \\ &\quad C_1 \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\tilde{\sigma}^2(x) - \sigma^2(x)| \right)^{1+\alpha} \right] + C_2 \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{f}(x) - f^*(x)| \right)^{2(1+\alpha)} \right] \\ &+ C_3 \left\{ \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{f}(x) - f^*(x)| \right)^{2(1+\alpha)} \right] \right\}^{1/2} \left\{ \mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} \left| \frac{1}{k_n} \sum_{i=1}^{k_n} (Y_{(i,n)}(x) - f^*(X_{(i,n)}(x))) \right| \right)^{2(1+\alpha)} \right] \right\}^{1/2}, \end{aligned}$$

where C_1 , C_2 and C_3 are non negative reals. We finish the proof of the proposition by bounded the above l.h.s. This relies on controls of estimation error of k NN for the regression function f^* and the conditional variance function σ^2 . Observe that when Y is either bounded or satisfies the model conditions in Eq. (5), we have that the random variables $Y - f^*(X)$ and $(Y - f^*(X))^2 - \sigma^2(X)$ satisfy the uniform noise condition (19). Indeed, while this fact is clear for $Y - f^*(X)$, it also holds true for $(Y - f^*(X))^2 - \sigma^2(X)$ since, conditionally on X , this random variable is either bounded (since σ^2 is bounded as well) or sub-exponential. Therefore, the result of Theorem D.3 applies for the k NN estimators $\tilde{\sigma}^2$ and \hat{f} . Furthermore, using the result in Eq. (26), we deduce from the above inequality

$$\mathbb{E} \left[\left(\sup_{x \in \mathcal{C}} |\hat{\sigma}^2(x) - \sigma^2(x)| \right)^{1+\alpha} \right] \leq C \log(n)^{(\alpha+1)} n^{-(\alpha+1)/(d+2)},$$

where $C > 0$ is a constant which depends on f^* , σ^2 , c_0 , \mathcal{C} , and the dimension d . \square

E Technical tools

In this section, we state several results that may help for readability of the paper. The first result is a direct application of the classical peeling argument of [1].

Lemma E.1 (Lemma 1 in [7]). *Let X be a real random variable, $(X_n)_{n \geq 1}$ be a sequence of real random variables and $t_0 \in \mathbb{R}$. Assume that there exist $C_1 > 0$ and $\gamma_0 > 0$ such that*

$$\mathbb{P}_X (|X - t_0| \leq \delta) \leq C_1 \delta^{\gamma_0}, \quad \forall \delta > 0,$$

and a sequence of positive numbers a_n tends towards infinity, C_2, C_3 some positive constants such that

$$\mathbb{P}_{X_n} (|X_n - X| \geq \delta | X) \leq C_2 \exp(-C_3 a_n \delta^2), \quad \forall \delta > 0, \quad \forall n \in \mathbb{N}.$$

Then, there exists $C > 0$ depending only on C_1, C_2 and C_3 , such that

$$|\mathbb{E} [\mathbb{1}_{X_n \geq t_0} - \mathbb{1}_{X \geq t_0}]| \leq C a_n^{-\gamma_0/2}.$$

The next result describes the representation of ∞ -Wasserstein distance (W_∞) on the real line. Let $Z_\infty(\mathbb{R})$ be the collection of all compactly supported probability measures on \mathbb{R} .

Lemma E.2 (Theorem 2.12 in [4]). *Let μ and ν be probability measures in $Z_\infty(\mathbb{R})$ with respective distribution functions F and G . Then, $W_\infty(\mu, \nu) := \sup_{0 < t < 1} |F^{-1}(t) - G^{-1}(t)|$ is the infimum over all $h \geq 0$ such that*

$$G(x - h) \leq F(x) \leq G(x + h) \text{ for all } x \in \mathbb{R}.$$

The following result provides a bound on moments of a positive random variable provided a tail control.

Lemma E.3. *Let $a \geq 1$, let b, c be two non negative real numbers, and let $m \in \mathbb{N}$. Consider Z a positive random variable such that*

$$\mathbb{P}(Z \geq t) \leq c \exp(a - bt^m) ,$$

for all $t \geq (a/b)^{1/m}$. Then for all $p \geq 1$, there exists a constant $C > 0$ such that

$$\mathbb{E}[Z^p] \leq C(a/b)^{p/m} .$$

Proof. Using the following equality which holds for any positive random variable Z , and any $p \geq 1$

$$\mathbb{E}[Z^p] = \int_0^{+\infty} \mathbb{P}(Z \geq t) pt^{p-1} dt , \quad (27)$$

and the condition in Lemma E.3, we deduce

$$\mathbb{E}[Z^p] \leq \int_0^u pt^{p-1} dt + c \int_u^{+\infty} \exp(a - bt^m) pt^{p-1} dt , \quad (28)$$

where $u = (a/b)^{1/m}$ and where we used the trivial inequality $\mathbb{P}(Z \geq t) \leq 1$ to bound the first term in the r.h.s. Since $(a')^m - (b')^m \geq (a' - b')^m$ for all $a', b' \in \mathbb{R}$ such that $a' \geq b' \geq 0$, we can write that

$$\exp(a - bt^m) \leq \exp(-(t - u)^m b) ,$$

which yields

$$\begin{aligned} \int_u^{+\infty} \exp(a - bt^m) pt^{p-1} dt &\leq \int_u^{+\infty} \exp(-(t - u)^m b) pt^{p-1} dt \\ &\leq \frac{1}{u} \int_u^{+\infty} \exp(-(t - u)^m b) pt^p dt \\ &= \frac{p}{u} \left(\frac{1}{b}\right)^{1/m} \int_0^{+\infty} e^{-v^m} \left(v \left(\frac{1}{b}\right)^{1/m} + u\right)^p dv , \end{aligned}$$

where we consider the changing of variable $v = ((t - u)^m b)^{1/m}$ in the last equality. Finally, using that $(a' + b')^p \leq 2^{p-1}((a')^p + (b')^p)$ for all $p \geq 1$, $a', b' \in \mathbb{R}$ and given that $u \geq (1/b)^{1/m}$, we show from the above inequality that

$$\begin{aligned} \int_u^{+\infty} \exp(a - bt^m) pt^{p-1} dt &\leq C_1 \left(\frac{1}{b}\right)^{p/m} \int_0^{+\infty} v^p e^{-v^m} dv + C_2 u^p \int_0^{+\infty} e^{-v^m} dv \\ &\leq C_3 u^p , \end{aligned}$$

for positive constants C_1, C_2, C_3 . Inject this into Eq.(28) leads to the result. \square