

Improved classification rates under refined margin conditions

Ingrid Blaschzyk and Ingo Steinwart

*Institute for Stochastics and Applications
University of Stuttgart
Pfaffenwaldring 57
D-70569 Stuttgart*

e-mail: ingrid.blaschzyk@mathematik.uni-stuttgart.de
ingo.steinwart@mathematik.uni-stuttgart.de

Abstract: In this paper we present a simple partitioning based technique to refine the statistical analysis of classification algorithms. The core of our idea is to divide the input space into two parts such that the first part contains a suitable vicinity around the decision boundary, while the second part is sufficiently far away from the decision boundary. Using a set of margin conditions we are then able to control the classification error on both parts separately. By balancing out these two error terms we obtain a refined error analysis in a final step. We apply this general idea to the histogram rule and show that even for this simple method we obtain, under certain assumptions, better rates than the ones known for support vector machines, for certain plug-in classifiers, and for a recently analyzed tree based adaptive-partitioning ansatz. Moreover, we show that a margin condition which sets the critical noise in relation to the decision boundary makes it possible to improve the optimal rates proven for distributions without this margin condition.

MSC 2010 subject classifications: Primary 62H30; secondary 62G20, 68T05.

Keywords and phrases: Statistical learning, classification, excess risk, fast rates of convergence, histogram rule.

Received October 2016.

Contents

1	Introduction	794
2	General assumptions	795
3	Oracle inequality and learning rates	798
4	Comparison of rates	802
5	Proofs	807
A	Appendix	821
	Acknowledgement	823
	References	823

1. Introduction

Given a dataset $D := ((x_1, y_1), \dots, (x_n, y_n))$ of observations drawn in an i.i.d. fashion from a probability measure P on $X \times Y$, where $X \subset \mathbb{R}^d$ and $Y := \{-1, 1\}$, the learning goal of binary classification is to find a decision function $f_D: X \rightarrow \{-1, 1\}$ such that for new data (x, y) we have $f_D(x) = y$ with high probability.

The problem of classification is, apart from regression, one of the most considered problems in learning theory and many classical learning methods have been presented in the literature such as histogram rules, nearest neighbor methods or moving window rules. A general reference for these methods is [4]. Several more recent methods use trees to build a classifier, for example the random forest algorithm, introduced in [3], makes a prediction by a majority vote over a collection of random forest trees. Another example is the tree based adaptive-partitioning algorithm, presented in [2]. Here, a classifier is picked by empirical risk minimization over a nested sequence $(S_m)_{m \geq 1}$ of families of sets which is based on dyadic or decorated tree partitions. Examples of non-tree based algorithms are described in [1] and [7]. There, the final classifier is found by empirical risk minimization over a suitable grid of plug-in rules or is derived by plug-in kernel, partitioning or nearest neighbor classification rules. Another non-tree based algorithm is, for example, the support vector machine (SVM), which solves a regularized empirical risk minimization problem over a reproducing kernel Hilbert space H . For more details on statistical properties of SVM for classification we refer the reader to [10, Chapter 8].

In this paper we discuss a partitioning based technique to analyse the statistical properties of classification algorithms. In particular we show for the histogram rule that under certain assumptions this technique leads to rates, which are faster than the rates obtained in [1, 2, 7], and [10]. To be more precise, we divide the input space X into two overlapping regions that are adjustable by a parameter r in such a way that one set, which we will denote by N_r , contains points near the decision boundary, whereas the other set F_r contains those that are sufficiently far away from the decision boundary. We examine the excess risks over these two sets separately by applying an oracle inequality for empirical risk minimizers on both parts. It turns out that we have no approximation error on F_r and that we obtain, under a suitable assumption which relates critical noise to the decision boundary, an optimal variance bound on F_r , which in turn leads to an $\mathcal{O}(n^{-1})$ behavior of the excess risk on F_r . However, this bound still depends on the parameter r , namely it increases for $r \rightarrow 0$. In contrast, our bound on the risk on N_r decreases for $r \rightarrow 0$. By balancing out these two risks with respect to r we obtain a refined bound on X under additional assumptions describing the concentration of mass around the decision boundary.

A more detailed discussion on this technique and the statistical results, which include rate adaptivity, are presented in Section 3. Moreover, a comparison of the resulting learning rates to the ones known for the SVM, for certain plug-in classification rules and the tree based adaptive-partitioning algorithm described in [2] can be found at the end of Section 4. In particular, we show that the

above mentioned assumption that relates the location of critical noise to the decision boundary has an essential influence on our learning rates such that we outperform under a common set of assumptions the optimal rates obtained for the classifier in [1]. Furthermore, we show that if we omit the latter assumption, we obtain exactly the optimal rate of [1]. We note that all proofs are deferred to Section 5.

2. General assumptions

To describe our learning goal we consider in the following the classification loss $L := L_{\text{class}} : Y \times \mathbb{R} \rightarrow [0, \infty)$, defined by $L(y, t) := \mathbf{1}_{(-\infty, 0]}(y \cdot \text{sign} t)$ for $y \in Y, t \in \mathbb{R}$, where $\mathbf{1}_{(-\infty, 0]}$ denotes the indicator function on $(-\infty, 0]$. We define the risk of a measurable estimator $f : X \rightarrow \mathbb{R}$ by

$$\mathcal{R}_{L,P}(f) := \int_{X \times Y} L(y, f(x)) dP(x, y)$$

and the empirical risk by

$$\mathcal{R}_{L,D}(f) := \frac{1}{n} \sum_{i=1}^n L(y_i, f(x_i)),$$

where $D := \frac{1}{n} \sum_{i=1}^n \delta_{(x_i, y_i)}$ denotes the average of Dirac measures $\delta_{(x_i, y_i)}$ at (x_i, y_i) . The smallest possible risk

$$\mathcal{R}_{L,P}^* := \inf_{f: X \rightarrow \mathbb{R}} \mathcal{R}_{L,P}(f)$$

is called the Bayes risk, and a measurable function $f_{L,P}^* : X \rightarrow \mathbb{R}$ so that $\mathcal{R}_{L,P}(f_{L,P}^*) = \mathcal{R}_{L,P}^*$ holds is called Bayes decision function. Recall that the Bayes decision function $f_{L,P}^*$ for the classification loss is given by $\text{sign}(2P(y = 1|x) - 1)$ for $x \in X$, where $P(\cdot|x)$ is a regular conditional probability on Y given x .

Let us now briefly describe a particular histogram rule. To this end, let $\mathcal{A} = (A_j)_{j \geq 1}$ be a partition of \mathbb{R}^d into cubes of side length $s \in (0, 1]$ and $X := [-1, 1]^d$. For $x \in X$ we denote by $A(x)$ the unique cell of \mathcal{A} with $x \in A(x)$ and call the map $h_{P,s} : X \rightarrow Y$ defined by

$$h_{P,s}(x) := \begin{cases} -1 & \text{if } f_{P,s}(x) < 0, \\ 1 & \text{if } f_{P,s}(x) \geq 0, \end{cases} \tag{1}$$

where $f_{P,s}(x) := P(A(x) \times \{1\}) - P(A(x) \times \{-1\})$, infinite sample histogram rule. For a dataset D we further write

$$f_{D,s}(x) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{y_i=+1\}} \mathbf{1}_{A(x)}(x_i) - \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{y_i=-1\}} \mathbf{1}_{A(x)}(x_i).$$

Thus, the empirical histogram is defined by $h_{D,s} := \text{sign}f_{D,s}$. We define the set \mathcal{F} by

$$\mathcal{F} := \left\{ \sum_{A_j \cap [-1,1]^d \neq \emptyset} c_j \mathbf{1}_{A_j} : c_j \in \{-1, 1\} \right\}.$$

Then, it is easy to show that the empirical histogram rule $h_{D,s}$ is an empirical risk minimizer over \mathcal{F} for the classification loss, that means

$$\mathcal{R}_{L,D}(h_{D,s}) = \inf_{f \in \mathcal{F}} \mathcal{R}_{L,D}(f).$$

Since we aim in a further step to examine the risk on subsets of X consisting of cells, we have to specify the loss on those subsets. Therefore, we define for an arbitrary index set $J \subset \{1, \dots, m\}$ the set

$$T_J := \bigcup_{j \in J} A_j \tag{2}$$

and the related loss $L_{T_J} : X \times Y \times \mathbb{R} \rightarrow [0, \infty)$ by

$$L_{T_J}(x, y, t) := \mathbf{1}_{\bigcup_{j \in J} A_j}(x) L_{\text{class}}(y, t). \tag{3}$$

Furthermore, we define the risk over T_J by

$$\mathcal{R}_{L_{T_J}, P}(f) := \int_{X \times Y} L_{T_J}(x, y, f(x)) dP(x, y)$$

and define the shortcut $L_{T_J} \circ f := L_{T_J}(x, y, f(x))$.

We denote by P^n the product measure of the probability measure P . As mentioned in the introduction, we have to make assumptions on P to obtain rates. Therefore, we recall some notions from [10, Chapter 8] which describe the behavior of P in the vicinity of the decision boundary. To this end, let $\eta : X \rightarrow [0, 1]$, defined by $\eta(x) := P(y = 1|x)$ for $x \in X$, be a version of the posterior probability of P , that is, that the probability measures $P(\cdot|x)$ form a regular conditional probability of P . Clearly, if we have $\eta(x) = 0$ resp. $\eta(x) = 1$ for $x \in X$ we observe the label $y = -1$ resp. $y = 1$ with probability 1. Otherwise, if, e.g., $\eta(x) \in [1/2, 1)$ we observe the label $y = -1$ with the probability $1 - \eta(x) \in (0, 1/2]$ and we call the latter probability noise. Obviously, in the worst case this probability equals $1/2$ and we define the set containing those $x \in X$ by $X_0 := \{x \in X : \eta(x) = 1/2\}$. Furthermore, we write

$$\begin{aligned} X_1 &:= \{x \in X : \eta(x) > 1/2\}, \\ X_{-1} &:= \{x \in X : \eta(x) < 1/2\}. \end{aligned}$$

Then, the function $\Delta_\eta : X \rightarrow [0, \infty]$ defined by

$$\Delta_\eta(x) := \begin{cases} d(x, X_1) & \text{if } x \in X_{-1}, \\ d(x, X_{-1}) & \text{if } x \in X_1, \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

where $d(x, A) := \inf_{x' \in A} d(x, x')$, is called distance to the decision boundary. This helps us to describe the mass of the marginal distribution P_X of P around the decision boundary by the following exponents. We say that P has margin exponent (ME) $\alpha \in (0, \infty]$ if there exists a constant $c_{\text{ME}} > 0$ such that

$$P_X(\{\Delta_\eta(x) < t\}) \leq (c_{\text{ME}}t)^\alpha$$

for all $t > 0$. Descriptively, the ME α measures the amount of mass close to the decision boundary. Therefore, large values of α are better since they reflect a low concentration of mass in this region, which makes the classification easier. Furthermore, we say that P has margin-noise exponent (MNE) $\beta \in (0, \infty]$ if there exists a constant $c_{\text{MNE}} > 0$ such that

$$\int_{\{\Delta_\eta < t\}} |2\eta(x) - 1| dP_X(x) \leq (c_{\text{MNE}}t)^\beta$$

for all $t > 0$. The MNE β measures the mass and the noise, that means the amount of points $x \in X$ with $\eta(x) \approx 1/2$, around the decision boundary. That is, we have high MNE β if we have low mass and/or high noise around the decision boundary. Next, we say that the distance to the decision boundary Δ_η controls the noise from below by the exponent γ if there exist a $\gamma \in [0, \infty)$ and a constant $c_{\text{LC}} > 0$ with

$$\Delta_\eta^\gamma(x) \leq c_{\text{LC}}|2\eta(x) - 1| \tag{5}$$

for P_X -almost all $x \in X$. That means, if $\eta(x)$ is close to $1/2$ for some $x \in X$, this x is close to the decision boundary. For examples of typical values of these exponents and relations between them we refer the reader to [10, Chapter 8].

Finally, in order to describe the region of the decision boundary in a more geometrical way, we say according to [6, 3.2.14(1)] that a general set $T \subset X$ is m -rectifiable for an integer $m > 0$ if there exists a Lipschitzian function mapping some bounded subset of \mathbb{R}^m onto T . Furthermore, we denote by $\partial_X T$ the relative boundary of T in X . Moreover, we denote by \mathcal{H}^{d-1} the $(d - 1)$ -dimensional Hausdorff measure on \mathbb{R}^d , see [6, Introduction]. The following lemma, which is based on [9, Lemma A.10.4], describes the Lebesgue measure of the decision boundary in terms of the Hausdorff measure. Its result will be necessary for the analysis of the main theorem in Section 3.

Lemma 2.1. *Let $X := [-1, 1]^d$ and P be a probability measure on $X \times \{-1, 1\}$ with fixed version $\eta: X \rightarrow [0, 1]$ of its posterior probability. Moreover, let λ^d be the d -dimensional Lebesgue measure and \mathcal{H}^{d-1} be the $(d - 1)$ -dimensional Hausdorff measure on \mathbb{R}^d . Furthermore, let $X_0 = \partial_X X_1$ with $\mathcal{H}^{d-1}(X_0) > 0$ and let X_0 be $(d - 1)$ -rectifiable. Then, there exists a $\delta^* > 0$ such that for all $\delta \in (0, \delta^*]$ we have*

$$\lambda^d(\{x \in X \mid \Delta_\eta(x) \leq \delta\}) \leq 4\mathcal{H}^{d-1}(\{x \in X \mid \eta(x) = 1/2\}) \cdot \delta.$$

3. Oracle inequality and learning rates

Our goal is to find an upper bound for the excess risk $\mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^*$. The idea is to split X into two overlapping sets and to find a bound on the risks over these sets by using information on P . To this end, we denote the set of indices of cubes that intersect X by

$$J := \{j \geq 1 \mid A_j \cap [-1, 1]^d \neq \emptyset\}.$$

Next, we split this set into cubes that lie near the decision boundary and into cubes that are bounded away from the decision boundary. To be more precisely, we define, for $r > 0$ and a version η for which the assumptions at the end of Section 2 hold, the set of indices of cubes near the decision boundary by

$$J_N^r := \{j \in J \mid \forall x \in A_j : \Delta_\eta(x) \leq 3r\}$$

and the set of indices of cubes that are sufficiently bounded away by

$$J_F^r := \{j \in J \mid \forall x \in A_j : \Delta_\eta(x) \geq r\}.$$

Moreover, we write

$$N_r := \bigcup_{j \in J_N^r} A_j, \tag{6}$$

$$F_r := \bigcup_{j \in J_F^r} A_j. \tag{7}$$

The next lemma shows that we are able to assign all $x \in A_j$ with $j \in J_F^r$ either to the class X_{-1} or to X_1 . Furthermore, we need to set geometric requirements to ensure that $X \subset N_r \cup F_r$.

Lemma 3.1. *Let $\mathcal{A} = (A_j)_{j \geq 1}$ be a partition of \mathbb{R}^d into cubes of side length $s \in (0, 1]$ and let $X := [-1, 1]^d$. For $r \geq s/2$ define the sets N_r and F_r by (6) and (7). Then, the following statements are true:*

- i) We have either $A_j \cap X_1 = \emptyset$ or $A_j \cap X_{-1} = \emptyset$ for $j \in J_F^r$.*
- ii) If $X_0 = \partial_X X_1 = \partial_X X_{-1}$, we have $X \subset N_r \cup F_r$.*

Lemma 3.1 ii) leads to a helpful splitting of the excess risk as the following lemma shows.

Lemma 3.2. *Under the assumptions of Lemma 3.1 ii) we have*

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^* \\ & \leq \left(\mathcal{R}_{L_{N_r},P}(h_{D,s}) - \mathcal{R}_{L_{N_r},P}^* \right) + \left(\mathcal{R}_{L_{F_r},P}(h_{D,s}) - \mathcal{R}_{L_{F_r},P}^* \right). \end{aligned}$$

That means, we can bound the excess risk $\mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^*$ if we find bounds on the excess risks over the sets N_r and F_r . For that purpose, we use

an oracle inequality for empirical risk minimizer separately on both error terms, see [10, Theorem 7.2]. This is possible, since the following lemma shows that, considering the loss L_{T_J} for any set T_J constructed as in (2), the empirical histogram rule $h_{D,s}$ is still an empirical risk minimizer over \mathcal{F} .

Lemma 3.3. *Consider for an arbitrary index set $J \subset \{1, \dots, m\}$ the set $T_J := \bigcup_{j \in J} A_j$ and the related loss $L_{T_J} : X \times Y \times \mathbb{R} \rightarrow [0, \infty)$ defined in (3). Then, the empirical histogram rule $h_{D,s}$ is an empirical risk minimizer over \mathcal{F} for the loss L_{T_J} , that means*

$$\mathcal{R}_{L_{T_J},D}(h_{D,s}) = \inf_{f \in \mathcal{F}} \mathcal{R}_{L_{T_J},D}(f).$$

Before we state our oracle inequality we discuss in a more detailed way the improvement that we gained by our separation technique described above. First, we make no approximation error on the set F_r , which consists of cells that are sufficiently bounded away from the decision boundary. This follows from the circumstance that $h_{D,s}$ learns correctly on those cells. We refer the reader to Part 1 of the proof of Lemma 3.5 for details. Second, the main refinement arises from the fact that we achieve, under the condition that the decision boundary controls the noise from below, a bound on F_r of the form

$$\mathbb{E}_P(L \circ f - L \circ f_{L,P}^*)^2 \leq V \cdot \mathbb{E}_P(L \circ f - L \circ f_{L,P}^*)^\theta$$

with the best possible exponent, $\theta = 1$. Here, V is a positive constant. The latter bound is known in the literature as variance bound. This bound plays an important part in the analysis of the risk terms since we have small variance if the right-hand side of the latter inequality is small. This relation is shown in detail in the next lemma.

Lemma 3.4. *Let $X := [-1, 1]^d$ and P be a probability measure on $X \times \{-1, 1\}$ with fixed version $\eta : X \rightarrow [0, 1]$ of its posterior probability. Assume that the associated distance to the decision boundary Δ_η controls the noise from below by the exponent $\gamma \in [0, \infty)$ and consider for some fixed $r > 0$ the set F_r , defined in (7). Furthermore, let $L := L_{class}$ be the classification loss and let $f_{L,P}^*$ be a fixed Bayes decision function. Then, for all measurable $f : X \rightarrow \{-1, 1\}$ we have*

$$\mathbb{E}_P(L_{F_r} \circ f - L_{F_r} \circ f_{L,P}^*)^2 \leq \frac{c_{LC}}{r^\gamma} \mathbb{E}_P(L_{F_r} \circ f - L_{F_r} \circ f_{L,P}^*).$$

We remark that the right-hand side of the variance bound on F_r depends on the separation parameter r . This dependence is also reflected in the risk term on F_r . In particular, we show in Part 1 of the proof of Theorem 3.5 by applying [10, Theorem 7.2] on the risk term on the set F_r that the improvements mentioned above lead to

$$\mathcal{R}_{L_{F_r},P}(h_{D,s}) - \mathcal{R}_{L_{F_r},P}^* \leq \frac{32\tilde{c}(8^{d+1}s^{-d} + \tau)}{r^\gamma n}$$

with probability $P^n \geq 1 - e^{-\tau}$, where $\tau \geq 1$ and \tilde{c} is a positive constant. Whereas this error term increases for $r \rightarrow 0$, the error term on the set N_r behaves exactly

the opposite way, that is, it decreases for $r \rightarrow 0$. In fact, bounding the risk on N_r requires additional knowledge of the behavior of P in the vicinity of the decision boundary. By applying [10, Theorem 7.2] on the risk on the set N_r we show in Part 2 of the proof of Theorem 3.5 under the assumption that P has ME α and MNE β that

$$\mathcal{R}_{L_{N_r}, P}(h_{D,s}) - \mathcal{R}_{L_{N_r}, P}^* \leq 6(c_{\text{MNE}}s)^\beta + 4 \left(\frac{8V(\bar{c}rs^{-d} + \tau)}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}}$$

holds with probability $P^n \geq 1 - e^{-\tau}$. Here, \bar{c} is a positive constant, $\tau \geq 1$ and V is the prefactor of the variance bound on N_r , shown in the second part of the proof. We refer the reader to the proof of Theorem 3.5 for exact constants. If we balance the obtained risk terms over N_r and F_r with respect to r , we obtain the oracle inequality presented in the following theorem. For this purpose, we define the positive constant

$$\tilde{c}_{\alpha,\gamma,d} := \left(\frac{16\gamma(\alpha+2\gamma) \cdot 8^{d+1} \max\{c_{\text{LC}}, 2\gamma\} \cdot (\alpha+\gamma)^{-1}}{\hat{c}_{\frac{\alpha+\gamma}{\alpha+2\gamma}}} \right)^{\frac{\alpha+\gamma}{\alpha+\gamma+\gamma(\alpha+2\gamma)}}, \quad (8)$$

which depends on α, γ and d and where

$$\hat{c} := 32 \max\{12\mathcal{H}^{d-1}(\{\eta = 1/2\}), 1\} \cdot \max\left\{1, \frac{\alpha+\gamma}{\gamma} c_{\text{ME}}^{\frac{\alpha\gamma}{\alpha+\gamma}} \left(\frac{\gamma c_{\text{LC}}}{\alpha}\right)^{\frac{\alpha}{\alpha+\gamma}}\right\}.$$

Theorem 3.5. *Let $\mathcal{A} = (A_j)_{j \geq 1}$ be a partition of \mathbb{R}^d into cubes of side length $s \in (0, 1]$. Let $X := [-1, 1]^d$ and P be a probability measure on $X \times \{-1, 1\}$ with fixed version $\eta: X \rightarrow [0, 1]$ of its posterior probability. Assume that the associated distance to the decision boundary Δ_η controls the noise from below by the exponent $\gamma \in [0, \infty)$ and assume as well that P has MNE $\beta \in (0, \infty]$ and ME $\alpha \in (0, \infty]$. Furthermore, let $X_0 = \partial_X X_1 = \partial_X X_{-1}$ with $\mathcal{H}^{d-1}(X_0) > 0$ and let X_0 be $(d-1)$ -rectifiable. Let L be the classification loss and let for fixed $n \geq 1$ and $\tau \geq 1$ the bounds*

$$s \leq \tilde{c}_{\alpha,\gamma,d}^{\frac{(1+\gamma)(\alpha+\gamma)+\gamma^2}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d\gamma}} \left(\frac{\tau}{n}\right)^{\frac{\gamma}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d\gamma}}, \quad (9)$$

and

$$s^d n \geq \tau \left(\frac{\tilde{c}_{\alpha,\gamma,d}}{\min\{\frac{\delta^*}{3}, 1\}} \right)^{\frac{(1+\gamma)(\alpha+\gamma)+\gamma^2}{\gamma}} \quad (10)$$

be satisfied, where the constant $\tilde{c}_{\alpha,\gamma,d}$ is defined by (8) and the constant $\delta^* > 0$ is the one of Lemma 2.1. Then, there exists a constant $c_{\alpha,\gamma,d} > 0$ such that

$$\mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^* \leq 6(c_{\text{MNE}}s)^\beta + c_{\alpha,\gamma,d} \left(\frac{\tau}{s^d n}\right)^{\frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2}} \quad (11)$$

holds with probability $P^n \geq 1 - 2e^{-\tau}$, where the constant $c_{\alpha,\gamma,d}$ only depends on α, γ and d .

The proof shows that the constants $c_{\alpha,\gamma,d}$ is given by

$$c_{\alpha,\gamma,d} := 128 \cdot 8^{d+1} \max\{c_{LC}, 2^\gamma\} \cdot \max\left\{\frac{\gamma(\alpha + 2\gamma)}{\alpha + \gamma}, 1\right\} \cdot \tilde{c}_{\alpha,\gamma,d}^{-\gamma}. \quad (12)$$

By choosing an appropriate sequence of s_n in dependence of our data length n and setting a constraint on the MNE β we state learning rates in the next theorem. Prior to that, we define with $\kappa := (1 + \gamma)(\alpha + \gamma)$ the positive constant

$$\tilde{c}_{\alpha,\beta,\gamma,\tau,d} := \left(\frac{d \cdot \kappa \cdot c_{\alpha,\gamma,d} \cdot \tau^{\frac{(1+\gamma)(\alpha+\gamma)}{\kappa+\gamma^2}}}{6\beta c_{MNE}^\beta (\kappa + \gamma^2)} \right)^{\frac{\kappa+\gamma^2}{\beta(\kappa+\gamma^2)+d\kappa}}$$

that depends on $\alpha, \beta, \gamma, \tau$ and d and where $c_{\alpha,\gamma,d}$ is the constant from (12).

Theorem 3.6. *Assume that X and P satisfy the assumptions of Theorem 3.5 for $\beta \leq \gamma^{-1}\kappa$, where $\kappa := (1 + \gamma)(\alpha + \gamma)$. In addition, assume that the side length s_n in Theorem 3.5 is given by*

$$s_n = \tilde{c}_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\kappa}{\beta(\kappa+\gamma^2)+d\kappa}}.$$

Then, there exists a constant $c_{\alpha,\beta,\gamma,\tau,d} > 0$ such that for all $n \geq n_0$

$$\mathcal{R}_{L,P}(h_{D,s_n}) - \mathcal{R}_{L,P}^* \leq c_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}}$$

holds with probability $P^n \geq 1 - 2e^{-\tau}$, where n_0 and the constant $c_{\alpha,\beta,\gamma,\tau,d}$ only depend on $\tau, \alpha, \beta, \gamma$ and d .

The proof of the latter theorem shows that the constant $c_{\alpha,\beta,\gamma,\tau,d}$ is given by

$$c_{\alpha,\beta,\gamma,\tau,d} := 2 \max\left\{\frac{d \cdot \kappa}{\beta(\kappa + \gamma^2)}, 1\right\} c_{\alpha,\gamma,\delta} \cdot \tau^{\frac{\kappa}{\kappa+\gamma^2}} \cdot \tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{-\frac{d\kappa}{\kappa+\gamma^2}}.$$

Furthermore, we remark that the constraint on the MNE β in Theorem 3.6 is set to secure that the chosen side length s_n fulfils assumption (9). If we omit this constraint we have to choose another s_n . For this s_n we would not be able to balance the two terms in the right-hand side of the excess risk in Theorem 3.5. Since our examples in Section 4 fulfil this constraint we did not consider other choices of s_n .

To obtain the rates we have to know the parameters describing P . However, it is also possible to obtain the rates in Theorem 3.6 by the following data splitting ansatz, a hold-out procedure whose concept is similar to the one described in [10, Chapter 6.5]. Let (S_n) be a sequence of finite subsets $S_n \subset (0, 1]$. For a dataset $D := ((x_1, y_1), \dots, (x_n, y_n))$ we define the sets

$$\begin{aligned} D_1 &:= ((x_1, y_1), \dots, (x_k, y_k)), \\ D_2 &:= ((x_{k+1}, y_{k+1}), \dots, (x_n, y_n)), \end{aligned}$$

where $k := \lfloor \frac{n}{2} \rfloor + 1$ and $n \geq 4$. Then, we use D_1 as a training set and compute $h_{D_1, s}$ for $s \in S_n$ and use D_2 to determine $s_{D_2}^* \in S_n$ such that

$$s_{D_2}^* := \arg \min_{s \in S_n} \mathcal{R}_{L, D_2}(h_{D_1, s}).$$

The resulting decision function is $h_{D_1, s_{D_2}^*}$ and a learning method producing this decision function is called training validation histogram rule (TV-HR). The following lemma shows that the TV-HR learns with the same rate as in Theorem 3.6 without knowing the parameters describing P .

Theorem 3.7. *Assume that X and P satisfy the assumptions of Theorem 3.5 for $\beta \leq \gamma^{-1}\kappa$, where $\kappa := (1 + \gamma)(\alpha + \gamma)$. Let S_n be a finite subset of $(0, 1]$ such that S_n is a $n^{-1/d}$ -net of $(0, 1]$. Assume that the cardinality of S_n grows at most polynomially in n . Then, the TV-HR learns with rate*

$$n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}}.$$

4. Comparison of rates

In order to compare our rate obtained in Theorem 3.6 to the ones known from [1, 2, 7] and [10], we set in the following reasonable sets of common assumptions. Besides our geometric assumption on X , namely

- (i) X_0 is $(d - 1)$ -rectifiable with $\mathcal{H}^{d-1}(X_0) > 0$ and $X_0 = \partial_X X_1 = \partial_X X_{-1}$,

we make the following two assumptions on P :

- (ii) P has ME $\alpha \in (0, \infty]$,
 (iii) there exists a $\gamma \in [0, \infty)$ and constants $c_{\text{LC}}, c_{\text{UC}} > 0$ such that for all $x \in X$ we have
- a) $c_{\text{LC}}|2\eta(x) - 1| \geq \Delta_\eta^\gamma(x)$,
 - b) $c_{\text{UC}}|2\eta(x) - 1| \leq \Delta_\eta^\gamma(x)$.

Here, assumption $(iii)_a$ coincides with the definition in (5). Furthermore, assumption $(iii)_b$ shows that we have an upper control by Δ_η on the noise, which is up to a constant a kind of inverse to $(iii)_a$. Then, [10, Lemma 8.17] shows under the assumptions (ii) and $(iii)_b$ that P has MNE $\beta = \alpha + \gamma$. Hence, we find by Theorem 3.6 with $\kappa := (1 + \gamma)(\alpha + \gamma)$ and a suitable cell-width s_n that h_{D, s_n} learns with a rate with exponent

$$\begin{aligned} \frac{\beta(1+\gamma)(\alpha+\gamma)}{\beta[(1+\gamma)(\alpha+\gamma)+\gamma^2]+d(1+\gamma)(\alpha+\gamma)} &= \frac{(\alpha+\gamma)(1+\gamma)(\alpha+\gamma)}{(\alpha+\gamma)[(1+\gamma)(\alpha+\gamma)+\gamma^2]+d(1+\gamma)(\alpha+\gamma)} \\ &= \frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d(1+\gamma)}. \end{aligned}$$

A simple transformation shows that this exponent equals

$$\frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d(1+\gamma)} = \frac{\alpha+\gamma}{\alpha+\gamma+\frac{\gamma^2}{1+\gamma}+d} = \frac{\alpha+\gamma}{\alpha+2\gamma+\frac{\gamma^2}{1+\gamma}+d-\gamma} = \frac{\alpha+\gamma}{\alpha+2\gamma+d-\frac{\gamma}{1+\gamma}}. \quad (13)$$

First, we compare the rate with exponent (13) to the rate achieved by support vector machines (SVM) for the hinge loss by assuming that (i), (ii) and (iii) hold. For this purpose, [10, Chapter 8.3 (8.18)] shows that the best possible rate for SVMs using Gaussian kernels is obtained by

$$n^{-\frac{\alpha+\gamma}{\alpha+2\gamma+d}+\rho},$$

where $\rho > 0$ is an arbitrary small number. Hence, our rate in (13) is better by $-\frac{\gamma}{1+\gamma}$ in the denominator. For the typical value of $\gamma = 1$, indicating a moderate control of noise by the decision boundary, our rate is better by $-1/2$ in the denominator.

Second, we compare our rates to the ones for certain plug-in classifiers, see [1, 7], and to the rates obtained by the classification algorithms, described in [2]. In the cases of [1] and [2] the authors assume that P has a noise exponent (NE) $q \in [0, \infty]$, that is, that there exist a constant $c_{NE} > 0$ such that

$$P_X(\{x \in X : |2\eta(x) - 1| < \varepsilon\}) \leq (c\varepsilon)^q \tag{14}$$

for all $\varepsilon > 0$, c.f. [10, Definition 8.22]. Since (14) measures the amount of critical noise and does not locate noise we call this exponent noise exponent in contrast to [8] and the mentioned authors, who call this exponent margin exponent. The authors of [7] assume a weaker version of (14) on P , however, the latter implies this weak version, see [5, Section 2]. We compare our rates under a different assumption set as in the first comparison to SVMs. To this end, we impose in addition to (i), (ii) and (iii)_a that

(iv) η is Hölder-continuous for some $\gamma \in (0, 1]$.

Then, we find under condition (iv) with Lemma A.2 that assumption (iii)_b is fulfilled with exponent γ and thus we assume in the following that (iii)_a holds for the same γ . Note that in (iv) we have $\gamma \in (0, 1]$, whereas in the case of (iii) we have $\gamma \in [0, \infty)$. Moreover, under assumptions (ii) and (iii)_a we find with [10, Exercise 8.5] that the noise exponent in (14) holds with

$$q = \frac{\alpha}{\gamma}. \tag{15}$$

By assuming (i), (ii), (iii)_a and (iv) our rate yields the same exponent as in (13), that is

$$\frac{\alpha+\gamma}{\alpha+2\gamma+d-\frac{\gamma}{1+\gamma}}. \tag{16}$$

Furthermore, the plug-in classifiers based on kernel, partitioning or nearest neighbor regression estimates shown in [7, Theorem 1, 3 and 5] yield under these assumptions and thus in particular with (15) the rate

$$n^{-\frac{\alpha+\gamma}{\alpha+3\gamma+d}}, \tag{17}$$

such that our rate is better by $-\frac{\gamma(2+\gamma)}{1+\gamma}$ in the denominator. The authors were able to improve the rate given in (17) by making in addition the assumption that P_X has a density with respect to the Lebesgue measure, which is bounded away from zero, see [7, Theorem 2, 4 and 6]. Under this condition and (i), (ii), (iii)_a and (iv) the classifiers yield the rate

$$n^{-\frac{\alpha+\gamma}{2\gamma+d}}.$$

Hence, our rate with exponent (13) is better if our margin exponent α fulfils $\alpha < \frac{\gamma}{1+\gamma}$. We have small margin exponent α , for example, if we have much mass around the decision boundary, that is, the density is unbounded in this region. We remark that the authors obtained rates under the Hölder assumption (iv), a weak margin assumption, and improved them as discussed above by making the assumption that P_X has a density which is bounded away from zero.

Next, we compare our rates to the ones obtained by the classifier resulting from the classification method given in [2, Section 5]. Therefore, we consider in addition to (i), (iii)_a and (iv) for example that

(v) P_X is the uniform distribution.

Under the condition that (i) and (v) hold, we find with Lemma 2.1 that assumption (ii) is fulfilled for $\alpha = 1$. Then, we obtain in (15) that $q = \frac{1}{\gamma}$. Again, we find with Lemma A.2 that assumption (iii)_b is fulfilled with exponent γ and assume again that (iii)_a holds for the same γ . Hence, the conditions (i) and (iii)_a, (iv) and (v) yield in (16) a rate with exponent

$$\frac{1+\gamma}{1+2\gamma+d-\frac{\gamma}{1+\gamma}} \quad (18)$$

for our method. Furthermore, [2, Corollary 5.2(ii)] shows that the classifier mentioned in [2, Section 5] yields the rate

$$\left(\frac{(\log n)^{\frac{1}{2+d}}}{n} \right)^{\frac{1+\gamma}{2\gamma+d}}. \quad (19)$$

Hence, our rate is worse by $\frac{1}{1+\gamma}$. However, the rate given in (19) is also comparable under a more generic assumption set in which we do not fix an example of P_X . Indeed, if we assume the conditions (i), (ii), (iii)_a and (iv), then, our rate with exponent (16) holds and [2, Corollary 5.2(i)] shows that their classifier obtains the rate

$$\left(\frac{\log n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma+d}}. \quad (20)$$

Thus, our rate with exponent (16) is again better by $-\frac{\gamma}{1+\gamma}$ in the denominator.

Finally, we compare our rates to the ones obtained for the plug-in classifier defined by [1, (4.1) with $p = \infty$] under the conditions (i), (iii)_a, (iv) and

(vi) P_X has a uniformly bounded density.

Analogously as above one can show with (i) and (vi) with Lemma 2.1 that we have MNE $\alpha = 1$. Under these conditions our rate with exponent (18) holds. The classifier in [1, (4.1) with $p = \infty$] achieves the rate

$$n^{-\frac{1+\gamma}{1+2\gamma+d}} \tag{21}$$

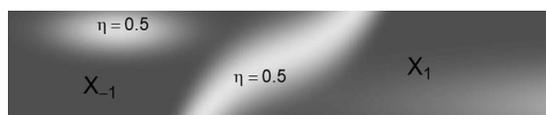
in expectation and we find that our rate is better by $-\frac{\gamma}{1+\gamma}$ in the denominator. We remark at this point that [1, Theorem 4.1 and 4.3] proved that the classifier achieves this rate under a different assumption set, namely under (iv), (vi) and the assumption that P has NE $q \in [0, \infty]$. The classifier then achieves the rate

$$n^{-\frac{\gamma(q+1)}{\gamma(q+2)+d}} \tag{22}$$

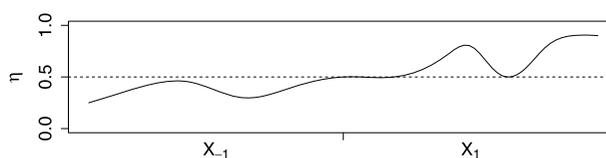
and for this set of assumptions the rate is optimal (in a minimax sense). Our assumptions, namely (i), (iii)_a, (iv) and (vi) imply the assumptions of [1], but, this is not a contradiction since our assumptions are a subset of the assumptions of [1].

Our improvement arises from assumption (iii)_a since it forces critical noise ($\eta \approx 1/2$) to be located close to the decision boundary and as we will see down below this assumption has an essential influence on the NE q . To be more precisely, there are two sources for slow learning rates. The first one is the approximation error around the decision boundary, the second one is the existence of critical noise. Assumption (iii)_a forces both to be in the same region such that both effects cannot independently occur, which in turn leads to better rates compared to [1]. In other words, with assumption (iii)_a we exclude distributions that have regions of critical noise that are far away from the decision boundary. Be aware that this does not mean that we consider only distributions without noisy regions bounded away from the decision boundary. In Fig. 1 we present two examples which make this situation more clear. Areas of noise that are, for example, located in the set X_1 in Fig. 1 (a) resp. in the set X_{-1} in Fig. 1 (b) are still allowed under (iii)_a whereas the areas of critical noise in the particular other set are permitted.

To make this heuristic argument more precise we take a look at Theorem 3.5 and its proof and show that if we omit assumption (iii)_a and thus consider the assumptions taken in [1], we match the optimal rate in (22). To this end, we consider in addition to (i) the above mentioned assumptions of [1], that is (iv), (vi) and the NE q . Since we do not assume (iii)_a we cannot use this assumption to obtain a variance bound on the set F_r , which is bounded away from the decision boundary (e.g., Lemma 3.4). Hence, the separation technique we used in the proof would make no sense any more, but we are able to bound the excess risk on the whole set by Part 2 in the proof of Theorem 3.5, where we bounded the excess risk on the set N_r that is close to the decision boundary. This situation corresponds to the fact that our set F_r is empty (letting go $r \rightarrow \infty$). There are two points that change in Part 2 of the proof. First, we have



(a) The brighter the region the more closer η is to $1/2$. The decision boundary (within the bright stripe) is located in the middle of the picture. Far away from the decision boundary we locate in both sets X_{-1} and X_1 noise (upper left and lower right corner), but only in the upper left corner the critical level $\eta = 1/2$ is reached.



(b) Far away from the decision boundary, in X_1 , we locate a region of critical noise. In X_{-1} we have noise too, but the critical level $\eta = 1/2$ is not reached.

FIG 1. Examples of η with regions of critical noise ($\eta \approx 1/2$) far away from the decision boundary. The size of these regions has an essential influence on the NE q . Since assumption $(iii)_a$ disallow areas of critical noise far away from the decision boundary we obtain a better noise exponent and hence, better rates. Note that areas of noise as in (a) in the lower right corner or in (b) in the lower left corner are still allowed.

no variance bound as in (38), but we can apply [10, Theorem 8.24], a general variance bound, and obtain

$$\mathbb{E}_P(h_{f_0}^N)^2 \leq \tilde{V} (\mathbb{E}_P h_{f_0}^N)^\theta,$$

where $\theta := \frac{q}{q+1}$ and where \tilde{V} is a positive constant. Second, we can bound the cardinality $|\mathcal{F}|$ in (40) as in Part 1 and yield with some calculations as in Part 3 for the overall excess risk

$$\mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^* \leq 6(c_{\text{MNE}}s)^\beta + c \left(\frac{\tau}{s^d \eta} \right)^{\frac{1}{2-\theta}},$$

where $c > 0$ is a constant and $\tau \geq 1$. Then, minimizing over s yields for our learning method the rate

$$n^{-\frac{\gamma(q+1)}{\gamma(q+2)+d}},$$

which matches the in [1, Theorem 4.1 and 4.3] proven optimal rate (22). We further remark that instead of the Hölder assumption (iv) the weaker assumption (iii)_b is sufficient for Theorem 3.5 and the above modified one.

If we consider now in addition to the assumptions in [1] that (i) and (iii)_a hold, we find that our rate improves immediately since we can directly apply Theorem 3.5 and thus obtain exactly the rate with exponent (18) which is better by $-\frac{\gamma}{1+\gamma}$. In summary, taking in consideration assumption (iii)_a influences the noise exponent in a good way since we exclude distributions that have critical noise far away from the decision boundary. This leads to better learning rates.

Finally, we remark that for our results as well as for the results from [1, 2, 7] and [10] less assumptions are sufficient and in the comparisons above we tried to formulate reasonable sets of common assumptions.

5. Proofs

Proof of Lemma 2.1: For a set $T \subset X$ and $\delta > 0$ we define as in [9] the sets

$$\begin{aligned} T^{+\delta} &:= \{x \in X \mid d(x, T) \leq \delta\}, \\ T^{-\delta} &:= X \setminus (X \setminus T)^{+\delta}. \end{aligned}$$

Since $X_1 := \{x \in X \mid \eta(x) \leq 1/2\}$ is bounded and measurable, we find with [9, Lemma A.10.3] and the proof of [9, Lemma A.10.4(ii)] that there exists a $\delta^* > 0$, such that for all $\delta \in (0, \delta^*]$ we have

$$\lambda^d(X_1^{+\delta} \setminus X_1^{-\delta}) \leq 4\mathcal{H}^{d-1}(\partial_X X_1) \cdot \delta = 4\mathcal{H}^{d-1}(X_0) \cdot \delta. \tag{23}$$

Next, we show that

$$\{x \in X \mid \Delta_\eta(x) \leq \delta\} \subset X_1^{+\delta} \setminus X_1^{-\delta} \cup X_0. \tag{24}$$

For this purpose, we remark that according to (4) we have

$$\begin{aligned} & \{x \in X \mid \Delta_\eta(x) \leq \delta\} \\ &= \{x \in X_1 \mid d(x, X_{-1}) \leq \delta\} \cup \{x \in X_{-1} \mid d(x, X_1) \leq \delta\} \cup X_0. \end{aligned}$$

Let us first show that $\{x \in X_1 \mid d(x, X_{-1}) \leq \delta\} \subset X_1^{+\delta} \setminus X_1^{-\delta}$. To this end, consider an $x \in X_1$ with $d(x, X_{-1}) \leq \delta$, where we check at once that $x \in X_1^{+\delta}$. Now, assume that $x \in X_1^{-\delta} = X \setminus (X \setminus X_1)^{+\delta}$. Then, we find that $x \notin (X \setminus X_1)^{+\delta}$ such that $d(x, X \setminus X_1) = d(x, X_{-1} \cup X_0) > \delta$. Hence, $x \notin X_1^{-\delta}$. Next, let us show that $\{x \in X_{-1} \mid d(x, X_1) \leq \delta\} \subset X_1^{+\delta} \setminus X_1^{-\delta}$. To this end, consider an $x \in X_{-1}$ with $d(x, X_1) \leq \delta$. Then, it is clear that $x \in X_1^{+\delta}$ by definition of $X_1^{+\delta}$. Furthermore, $x \notin X_1^{-\delta}$ since $X_1^{-\delta} = X \setminus (X_{-1})^{+\delta} \subset X_1$. Having showed (24), we find together with the fact that $\lambda^d(X_0) = 0$ since X_0 is $(d-1)$ -rectifiable that

$$\lambda^d(\{x \in X \mid \Delta_\eta(x) \leq \delta\}) \leq \lambda^d(X_1^{+\delta} \setminus X_1^{-\delta}).$$

Finally, with (23) and the fact that $X_0 = \partial_X X_1$ we find that

$$\begin{aligned} \lambda^d(\{x \in X \mid \Delta_\eta(x) \leq \delta\}) &\leq \lambda^d(X_1^{+\delta} \setminus X_1^{-\delta}) \\ &\leq 4\mathcal{H}^{d-1}(X_0) \cdot \delta \\ &= 4\mathcal{H}^{d-1}(\{x \in X \mid \eta(x) = 1/2\}) \cdot \delta \end{aligned}$$

for all $\delta \in (0, \delta^*]$. □

Proof of Lemma 3.1:

- i) We assume for A_j with $j \in J_F^r$ that we have an $x_1 \in A_j \cap X_1 \neq \emptyset$ and an $x_{-1} \in A_j \cap X_{-1} \neq \emptyset$. Then, the connecting line $\overline{x_{-1}x_1}$ from x_{-1} to x_1 is contained in A_j since A_j is convex and we have $\|x_{-1} - x_1\|_\infty \leq s$. Moreover, since $\Delta_\eta(x) \geq r$ for all $x \in F_r$ we have that $x \in X_1 \cup X_{-1}$. Next, pick an $m > 1$ such that

$$t_0 = 0, \quad t_m = 1, \quad t_i = \frac{i}{m}$$

and

$$x_i := t_i x_{-1} + (1 - t_i) x_1$$

for $i = 0, \dots, m$. Clearly, $x_i \in \overline{x_{-1}x_1}$ and $x_i \in X_{-1} \cup X_1$. Since $x_0 \in X_1$ and $x_m \in X_{-1}$, there exists an i with $x_i \in X_1$ and $x_{i+1} \in X_{-1}$ and we find that

$$\|x_i - x_{i+1}\|_\infty \geq \Delta_\eta(x_i) \geq r.$$

On the other hand,

$$\|x_i - x_{i+1}\|_\infty = \frac{1}{m} \|x_{-1} - x_1\|_\infty \leq \frac{s}{m} \leq \frac{2r}{m}$$

such that $r \leq \frac{2r}{m}$, which is not true for $m \geq 3$. Hence, we can not have an $x_1 \in A_j \cap X_1 \neq \emptyset$ and an $x_{-1} \in A_j \cap X_{-1} \neq \emptyset$ for $j \in J_F^r$.

ii) We define the set of indices

$$J_C^r := \{j \in J \mid \exists \tilde{x} \in A_j : \Delta_\eta(\tilde{x}) < r\}$$

and define the set

$$C_r := \bigcup_{j \in J_C^r} A_j.$$

Since $X \subset F_r \cup C_r$, it suffices to show that $C_r \subset N_r$. To show the latter we fix an $x \in C_r$. If $x \in X_0$ we immediately have $\Delta_\eta(x) = 0 < 3r$, hence we assume w.l.o.g. that $x \in X_1$. Then, there exists a $j \in J_C^r$ such that $x \in A_j$. Furthermore, there exists an $x^* \in A_j$ with $\Delta_\eta(x^*) < r$ and we find with $X_0 = \partial_X X_1 = \partial_X X_{-1}$ that

$$\begin{aligned} \Delta_\eta(x) &= \inf_{x' \in X_{-1}} \|x - x'\|_\infty \\ &\leq \inf_{x' \in X_{-1}} (\|x - x^*\|_\infty + \|x^* - x'\|_\infty) \\ &\leq s + \Delta_\eta(x^*) \\ &< s + r, \end{aligned}$$

where $\|\cdot\|_\infty$ is the supremum norm in \mathbb{R}^d . Since $s \leq 2r$, it follows that $\Delta_\eta(x) \leq 3r$ and therefore $x \in N_r$. \square

Proof of Lemma 3.2: Under the assumptions of Lemma 3.1 ii) we find that $X \subset N_r \cup F_r$. Since the excess risk is non-negative we then have

$$\begin{aligned} &\mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^* \\ &= \int_{X \times Y} L(y, f(x) - f^*(x)) dP(x, y) \\ &= \int_X \int_Y L(y, f(x) - f^*(x)) P(dy|x) dP_X(x) \\ &\leq \int_{N_r \cup F_r} \int_Y L(y, f(x) - f^*(x)) P(dy|x) dP_X(x) \\ &= \int_{N_r} \int_Y L(y, f(x) - f^*(x)) P(dy|x) dP_X(x) \\ &\quad + \int_{F_r} \int_Y L(y, f(x) - f^*(x)) P(dy|x) dP_X(x) \\ &= \left(\mathcal{R}_{L_{N_r},P}(h_{D,s}) - \mathcal{R}_{L_{N_r},P}^*\right) + \left(\mathcal{R}_{L_{F_r},P}(h_{D,s}) - \mathcal{R}_{L_{F_r},P}^*\right). \quad \square \end{aligned}$$

Proof of Lemma 3.3: For $f \in \mathcal{F}$ we have

$$\begin{aligned} \mathcal{R}_{L_{T_J}, D}(f) &= \int_{X \times Y} L_{T_J}(x, y, f(x)) dD(x, y) \\ &= \sum_{j \in J} \int_{A_j \times Y} L_{\text{class}}(y, f(x)) dD(x, y). \end{aligned}$$

Next, we take a closer look at the risk on a single cell A_j for $j \in J$. That is,

$$\int_{A_j \times Y} L_{\text{class}}(y, f(x)) dD(x, y) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{A_j}(x_i) \mathbf{1}_{y_i \neq c_j},$$

where $c_j \in \{-1, 1\}$ is the label of the cell A_j . The risk on a cell is the smaller the less often we have $y_i \neq c_j$ such that the best classifier on a cell is the one which decides by majority. This is true for the histogram rule by definition. Since the risk is zero on A_j with $j \notin J$, the histogram rule minimizes the risk with respect to L_{T_J} . \square

Proof of Lemma 3.4: We define $h_f := L_{F_r} \circ f - L_{F_r} \circ f_{L, P}^*$ for a measurable $f: X \rightarrow \{-1, 1\}$. Since $(L_{F_r} \circ f - L_{F_r} \circ f_{L, P}^*)^2 = \mathbf{1}_{F_r} \frac{|f - f_{L, P}^*|}{2}$ we obtain

$$\begin{aligned} &\mathbb{E}_P(h_f - \mathbb{E}_P h_f)^2 \\ &\leq \mathbb{E}_P(h_f)^2 \\ &= \mathbb{E}_P(L_{F_r} \circ f - L_{F_r} \circ f_{L, P}^*)^2 \\ &= \frac{1}{2} \mathbb{E}_P \mathbf{1}_{F_r} |f - f_{L, P}^*|. \end{aligned}$$

For $x \in F_r$ we have $\Delta_\eta(x) \geq r$ and thus we find with our lower-control assumption that

$$r^\gamma \leq \Delta_\eta^\gamma(x) \leq c_{LC} |2\eta(x) - 1|$$

and therefore

$$1 \leq c_{LC} r^{-\gamma} |2\eta(x) - 1|.$$

By using $\mathbf{1}_{F_r} \frac{|f - f^*|}{2} = \mathbf{1}_{(X_{-1} \Delta \{f < 0\}) \cap F_r}$, where Δ denotes the symmetric difference defined by $C \Delta D := (C \setminus D) \cup (D \setminus C)$ for sets $C, D \subset X$ and by using Lemma A.1 we obtain for the variance bound

$$\begin{aligned} \mathbb{E}_P(h_f - \mathbb{E}_P h_f)^2 &\leq \frac{1}{2} \int \mathbf{1}_{F_r}(x) |f(x) - f_{L, P}^*(x)| dP_X(x) \\ &\leq \frac{c_{LC}}{2r^\gamma} \int \mathbf{1}_{F_r}(x) |2\eta(x) - 1| |f(x) - f_{L, P}^*(x)| dP_X(x) \end{aligned}$$

$$\begin{aligned}
&= \frac{c_{LC}}{r^\gamma} \int_{(X_{-1} \Delta \{f < 0\}) \cap F_r} |2\eta(x) - 1| dP_X(x) \\
&= \frac{c_{LC}}{r^\gamma} (\mathcal{R}_{L_{F_r}, P}(f) - \mathcal{R}_{L_{F_r}, P}^*) \\
&= \frac{c_{LC}}{r^\gamma} \mathbb{E}_P h_f. \quad \square
\end{aligned}$$

Proof of Theorem 3.5: We define the set of cubes N_r and F_r as in (6), (7) for the choice of

$$r := \tilde{c}_{\alpha, \gamma, d} \left(\frac{\tau}{s^d n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}}, \quad (25)$$

where

$$\theta := \frac{\alpha}{\alpha + \gamma}. \quad (26)$$

With (25) we find that $s \leq r$. To see the latter, we remark that

$$\begin{aligned}
s \leq \tilde{c}_{\alpha, \gamma, d} \left(\frac{\tau}{s^d n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} &\iff s^{\frac{1+\gamma(2-\theta)+d(1-\theta)}{1+\gamma(2-\theta)}} \leq \tilde{c}_{\alpha, \gamma, d} \left(\frac{\tau}{n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} \\
&\iff s \leq \left(\tilde{c}_{\alpha, \gamma, d} \left(\frac{\tau}{n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} \right)^{\frac{1+\gamma(2-\theta)}{1+\gamma(2-\theta)+d(1-\theta)}}
\end{aligned}$$

and conclude by replacing θ by (26) that $s \leq r$ holds if

$$s \leq \tilde{c}_{\alpha, \gamma, d}^{\frac{(1+\gamma)(\alpha+\gamma)+\gamma^2}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d\gamma}} \left(\frac{\tau}{n} \right)^{\frac{\gamma}{(1+\gamma)(\alpha+\gamma)+\gamma^2+d\gamma}},$$

which equals (9). Hence, we are able to split the excess risk $\mathcal{R}_{L, P}(h_{D, s}) - \mathcal{R}_{L, P}^*$ according to Lemma 3.2 by

$$\begin{aligned}
&\mathcal{R}_{L, P}(h_{D, s}) - \mathcal{R}_{L, P}^* \\
&\leq \left(\mathcal{R}_{L_{N_r}, P}(h_{D, s}) - \mathcal{R}_{L_{N_r}, P}^* \right) + \left(\mathcal{R}_{L_{F_r}, P}(h_{D, s}) - \mathcal{R}_{L_{F_r}, P}^* \right). \quad (27)
\end{aligned}$$

The rest of the proof is structured in three parts, where we establish error bounds on N_r and F_r in the first two parts and combine the results obtained in the third and last part of the proof. In the following we write $N := N_r$ and $F := F_r$ and keep in mind, that these sets depend on a parameter r . Furthermore, we write $h_D := h_{D, s}$.

Part 1: In the first part we establish an oracle inequality for $\mathcal{R}_{L_{F, P}}(h_{D, s}) - \mathcal{R}_{L_{F, P}}^*$. Therefore we define $h_f^F := L_F \circ f - L_F \circ f_{L_{F, P}}^*$ and find that

$$\|h_f^F\|_\infty = \|L_F \circ f - L_F \circ f_{L_{F, P}}^*\|_\infty \leq 1$$

for all $f \in \mathcal{F}$. Furthermore, with Lemma 3.4 we obtain

$$\mathbb{E}_P (h_f^F)^2 \leq \frac{c_{LC}}{r^\gamma} \mathbb{E}_P h_f^F \leq \frac{c_1}{r^\gamma} \mathbb{E}_P h_f^F, \quad (28)$$

where $c_1 := \max\{c_{\text{LC}}, 2^\gamma\}$. We observe that $r^\gamma \leq c_1$, since with assumption (10), where we rewrite the exponent by $\frac{(1+\gamma)(\alpha+\gamma)+\gamma^2}{\gamma} = \frac{1+\gamma(2-\theta)}{1-\theta}$, we find

$$\begin{aligned} r &= \tilde{c}_{\alpha,\gamma,d} \left(\frac{\tau}{s^d n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} \\ &\leq \tilde{c}_{\alpha,\gamma,d} \left(\left(\frac{\min\{\frac{\delta^*}{3}, 1\}}{\tilde{c}_{\alpha,\gamma,d}} \right)^{\frac{1+\gamma(2-\theta)}{1-\theta}} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} \\ &= \min \left\{ \frac{\delta^*}{3}, 1 \right\} \\ &\leq 1 \end{aligned}$$

and therefore $r^\gamma \leq 2^\gamma \leq c_1$. As we conclude from Lemma 3.3 that h_D is an empirical risk minimizer over \mathcal{F} for the loss L_F , we are able to use [10, Theorem 7.2], an improved oracle inequality for ERM. We obtain for all fixed $\tau \geq 1$ and $n \geq 1$ that

$$\mathcal{R}_{L_F,P}(h_D) - \mathcal{R}_{L_F,P}^* < 6(\mathcal{R}_{L_F,P,\mathcal{F}}^* - \mathcal{R}_{L_F,P}^*) + \frac{32c_1(\log(|\mathcal{F}| + 1) + \tau)}{r^\gamma n}$$

holds with probability $P^n \geq 1 - e^{-\tau}$, where $\mathcal{R}_{L_F,P,\mathcal{F}}^* := \inf_{f \in \mathcal{F}} \mathcal{R}_{L_F,P}(f)$. Next, we refine the right-hand side of this oracle inequality. Obviously we have $|\mathcal{F}| \leq 2^{|J|}$. We bound the cardinality $|J|$ by using a volume comparison argument. To this end, we define the set $\tilde{J} := \{j \geq 1 \mid A_j \cap 2[-1, 1]^d \neq \emptyset\}$ and observe that $\bigcup_{j \in J} A_j \subset \bigcup_{j \in \tilde{J}} A_j \subset 4B_{\ell_\infty^d}$. Then,

$$|J|s^d = \lambda^d \left(\bigcup_{j \in J} A_j \right) \leq \lambda^d \left(\bigcup_{j \in \tilde{J}} A_j \right) \leq \lambda^d (4B_{\ell_\infty^d}) = 8^d,$$

such that we deduce with $|J| \leq 8^d s^{-d}$ that

$$\begin{aligned} \log(|\mathcal{F}| + 1) &\leq \log(2^{8^d s^{-d}} + 1) \\ &\leq \log(2 \cdot 2^{8^d s^{-d}}) \\ &= \log(2^{8^d s^{-d} + 1}) \\ &= (8^d s^{-d} + 1)\log(2) \\ &\leq 8^d s^{-d} + 1 \\ &\leq 8^{d+1} s^{-d}. \end{aligned}$$

Thus,

$$\mathcal{R}_{L_F,P}(h_D) - \mathcal{R}_{L_F,P}^* < 6(\mathcal{R}_{L_F,P,\mathcal{F}}^* - \mathcal{R}_{L_F,P}^*) + \frac{32c_1(8^{d+1}s^{-d} + \tau)}{r^\gamma n} \quad (29)$$

holds with probability $P^n \geq 1 - e^{-\tau}$.

Finally, we have to bound the *approximation error* $\mathcal{R}_{L_F, P, \mathcal{F}}^* - \mathcal{R}_{L_F, P}^* = \inf_{f \in \mathcal{F}} \mathcal{R}_{L_F, P}(f) - \mathcal{R}_{L_F, P}^*$. We find with $h_{P,s} \in \mathcal{F}$ and Lemma A.1 that

$$\begin{aligned} \mathcal{R}_{L_F, P, \mathcal{F}}^* - \mathcal{R}_{L_F, P}^* &\leq \mathcal{R}_{L_F, P}(h_{P,s}) - \mathcal{R}_{L_F, P}^* \\ &= \int_{(X_1 \Delta \{h_{P,s} \geq 0\}) \cap F} |2\eta - 1| dP_X \\ &= \sum_{j \in J_F^r} \int_{(X_1 \Delta \{h_{P,s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X \\ &= 0, \end{aligned} \tag{30}$$

since $P_X((X_1 \Delta \{h_{P,s} \geq 0\}) \cap A_j) = 0$ for each $j \in J_F^r$. To see the latter, we first remark that the latter set contains those $x \in A_j$ for that either $h_{P,s}(x) \geq 0$ and $\eta(x) \leq 1/2$ or $h_{P,s}(x) < 0$ and $\eta(x) > 1/2$. Since we have $A_j \subset X_{-1} \cup X_1$ we can ignore the case $\eta(x) = 1/2$. Furthermore, we know by Lemma 3.1 i) that either $A_j \cap X_{-1} = \emptyset$ or $A_j \cap X_1 = \emptyset$. Let us first consider the case $A_j \cap X_{-1} = \emptyset$ and thus $A_j \subset X_1$. According to the definition of the histogram rule, cf. (1), we find for all $x \in A_j$ that $h_{P,s}(x) = 1$, since

$$\begin{aligned} &f_{P,s}(x) \\ &= P(A_j(x) \times \{1\}) - P(A_j(x) \times \{-1\}) \\ &= \int_{A_j} \int_Y \mathbf{1}_{A_j \times \{1\}}(x, y) P(dy|x) dP_X(x) \\ &\quad - \int_{A_j} \int_Y \mathbf{1}_{A_j \times \{-1\}}(x, y) P(dy|x) dP_X(x) \\ &= \int_{A_j} \mathbf{1}_{A_j \times \{1\}}(x, 1) \eta(x) dP_X(x) - \int_{A_j} \mathbf{1}_{A_j \times \{-1\}}(x, -1) (1 - \eta(x)) dP_X(x) \\ &= \int_{A_j} 2\eta(x) - 1 dP_X(x) \\ &\geq 0. \end{aligned}$$

Obviously we have $\eta(x) \geq 1/2$ and $h_{P,s}(x) = 1$ for all $x \in A_j$. Analogously we can show for cells with $A_j \cap X_1 = \emptyset$ for $j \in J_F^r$ that $\eta(x) \leq 1/2$ and $h_{P,s}(x) = -1$ for all $x \in A_j$. Hence, $P_X((X_1 \Delta \{h_{P,s} \geq 0\}) \cap A_j) = 0$ for all $j \in J_F^r$ and the approximation error vanishes on the set F .

Altogether, for the oracle inequality on F we obtain with (29) and (30) that

$$\mathcal{R}_{L_F, P}(h_D) - \mathcal{R}_{L_F, P}^* < \frac{32c_1(8^{d+1}s^{-d} + \tau)}{r^\gamma \eta} \tag{31}$$

holds with probability $P^n \geq 1 - e^{-\tau}$.

Part 2: In the second part we establish an oracle inequality for $\mathcal{R}_{L_N, P}(h_D) - \mathcal{R}_{L_N, P}^*$, again by using [10, Theorem 7.2]. Analogously to Part 1 we define $h_f^N := L_N \circ f - L_N \circ f_{L_N, P}^*$ for $f \in \mathcal{F}$ and find $\|h_f^N\|_\infty \leq 1$. Since $(h_{f_0}^N)^2 =$

$\mathbf{1}_N \frac{|f - f_{L,P}^*|}{2} = \mathbf{1}_{(X_{-1} \Delta \{f < 0\}) \cap N}$ we find with [Appendix, Lemma A.1] that

$$\begin{aligned}
& \mathbb{E}_P(h_{f_0}^N)^2 \\
&= \frac{1}{2} \int_N |f_0(x) - f_{L_N,P}^*(x)| dP_X(x) \\
&= \frac{1}{2} \int_{N \cap \{|2\eta-1| \geq t\}} |f_0(x) - f_{L_N,P}^*(x)| dP_X(x) \\
&\quad + \frac{1}{2} \int_{N \cap \{|2\eta-1| < t\}} |f_0(x) - f_{L_N,P}^*(x)| dP_X(x) \\
&\leq \frac{1}{2t} \int_{N \cap \{|2\eta-1| \geq t\}} |2\eta(x) - 1| |f_0(x) - f_{L_N,P}^*(x)| dP_X(x) \\
&\quad + P_X(\{x \in N : |2\eta(x) - 1| < t\}) \\
&\leq \frac{1}{2t} \int_N |2\eta(x) - 1| |f_0(x) - f_{L_N,P}^*(x)| dP_X(x) \\
&\quad + P_X(\{x \in N : |2\eta(x) - 1| < t\}) \\
&\leq t^{-1} \mathbb{E}_P h_{f_0}^N + \min\{P_X(N), P_X(\{x \in X : |2\eta(x) - 1| < t\})\}
\end{aligned} \tag{32}$$

for all $t > 0$. We turn our attention to the minimum and note that by the definition of N we have

$$P_X(N) \leq P_X(\{\Delta_\eta(x) \leq 3r\}). \tag{33}$$

For $x \in X$ with $|2\eta(x) - 1| < t$ by the definition of the lower control we conclude from

$$\frac{\Delta_\eta^\gamma(x)}{c_{LC}} \leq |2\eta(x) - 1| < t.$$

that

$$\Delta_\eta(x) \leq (c_{LC}t)^{\frac{1}{\gamma}}$$

and consequently

$$\{x \in X : |2\eta(x) - 1| < t\} \subset \{x \in X : \Delta_\eta(x) \leq (c_{LC}t)^{\frac{1}{\gamma}}\}. \tag{34}$$

Then, we find by (33), (34) and by the definition of the margin exponent that

$$\begin{aligned}
& \min\{P_X(N), P_X(\{x \in X : |2\eta(x) - 1| < t\})\} \\
&\leq \min\{P_X(\{\Delta_\eta(x) \leq 3r\}), P_X(\{x \in X : \Delta_\eta(x) \leq (c_{LC}t)^{\frac{1}{\gamma}}\})\} \\
&\leq \min\{(c_{ME}3r)^\alpha, c_{ME}^\alpha (c_{LC}t)^{\frac{\alpha}{\gamma}}\}.
\end{aligned} \tag{35}$$

Combining (35) with (32) we obtain

$$\begin{aligned}
\mathbb{E}_P(h_{f_0}^N - \mathbb{E}_P h_{f_0}^N)^2 &\leq t^{-1} \mathbb{E}_P h_{f_0}^N + \min\{(c_{ME}3r)^\alpha, c_{ME}^\alpha (c_{LC}t)^{\frac{\alpha}{\gamma}}\} \\
&\leq t^{-1} \mathbb{E}_P h_{f_0}^N + c_{ME}^\alpha (c_{LC}t)^{\frac{\alpha}{\gamma}}.
\end{aligned} \tag{36}$$

Minimizing the right-hand side of (36) yields

$$\min_{t>0} \left(t^{-1} \mathbb{E}_P h_{f_0}^N + c_{\text{ME}}^\alpha (c_{\text{LC}} t)^{\frac{\alpha}{\gamma}} \right) = c_2 \left(\mathbb{E}_P h_{f_0}^N \right)^{\frac{\alpha}{\alpha+\gamma}},$$

where $c_2 := \frac{\alpha+\gamma}{\gamma} c_{\text{ME}}^{\frac{\alpha\gamma}{\alpha+\gamma}} \left(\frac{\gamma c_{\text{LC}}}{\alpha} \right)^{\frac{\alpha}{\alpha+\gamma}}$, such that with

$$V := \max\{1, c_2\} \quad (37)$$

and (26) we have

$$\mathbb{E}_P (h_{f_0}^N)^2 \leq t^{-1} \mathbb{E}_P h_{f_0}^N + c_{\text{ME}} (c_\gamma t^{\frac{1}{\gamma}})^\alpha = c_2 \left(\mathbb{E}_P h_{f_0}^N \right)^{\frac{\alpha}{\alpha+\gamma}} \leq V \left(\mathbb{E}_P h_{f_0}^N \right)^\theta. \quad (38)$$

Note, that the definition of V yields $V^{\frac{1}{2-\theta}} \geq 1$. Since h_D is an ERM over \mathcal{F} for the loss L_N due to Lemma 3.3, by using [10, Theorem 7.2] we obtain for fixed $\tau \geq 1$ and $n \geq 1$ that

$$\begin{aligned} & \mathcal{R}_{L_N, P}(h_D) - \mathcal{R}_{L_N, P}^* \\ & < 6(\mathcal{R}_{L_N, P, \mathcal{F}}^* - \mathcal{R}_{L_N, P}^*) + 4 \left(\frac{8V(\log(|\mathcal{F}| + 1) + \tau)}{n} \right)^{\frac{1}{2-\theta}} \end{aligned} \quad (39)$$

holds with probability $P^n \geq 1 - e^{-\tau}$. In order to refine the right-hand side in (39), we establish a bound on the cardinality $|\mathcal{F}| = 2^{|J_A|}$ and on the approximation error. To bound the mentioned cardinality we use the fact that N is contained in a tube around the decision line, that is $\bigcup_{j \in J_N} A_j \subset \{\Delta_\eta(x) \leq 3r\}$, see (6). We remark that $3r \leq \delta^*$ holds, where δ^* is the constant from Lemma 2.1, since with assumption (10) we have

$$3r = 3\tilde{c}_{\alpha, \gamma, d} \left(\frac{\tau}{s^d n} \right)^{\frac{1-\theta}{1+\gamma(2-\theta)}} \leq 3 \min \left\{ \frac{\delta^*}{3}, 1 \right\} \leq \delta^*.$$

Then, with Lemma 2.1 we find that

$$\lambda^d(\{\Delta_\eta(x) \leq 3r\}) \leq 12\mathcal{H}^{d-1}(\{\eta = 1/2\})r$$

and we obtain

$$|J_A|s^d = \lambda^d \left(\bigcup_{j \in J_A} A_j \right) \leq \lambda^d(\{\Delta_\eta(x) \leq 3r\}) \leq 12\mathcal{H}^{d-1}(\{\eta = 1/2\})r.$$

This yields to

$$|J_A| \leq 12\mathcal{H}^{d-1}(\{\eta = 1/2\})rs^{-d} = c_3rs^{-d},$$

where $c_3 := 12\mathcal{H}^{d-1}(\{\eta = 1/2\})$. By $r \geq s \geq s^d$ we hence conclude that

$$\begin{aligned}
\log(|\mathcal{F}| + 1) &\leq \log(2^{c_3 r s^{-d}} + 1) \\
&\leq \log(2 \cdot 2^{c_3 r s^{-d}}) \\
&= \log(2^{c_3 r s^{-d} + 1}) \\
&= (c_3 r s^{-d} + 1) \log(2) \\
&\leq c_3 r s^{-d} + r s^{-d} \\
&\leq c_4 r s^{-d},
\end{aligned} \tag{40}$$

where $c_4 := 2 \max\{12\mathcal{H}^{d-1}(\{\eta = 1/2\}), 1\}$. Thus, (39) changes to

$$\mathcal{R}_{L_N, P}(h_D) - \mathcal{R}_{L_N, P}^* \leq 6(\mathcal{R}_{L_N, P, \mathcal{F}}^* - \mathcal{R}_{L_N, P}^*) + 4 \left(\frac{8V(c_4 r s^{-d} + \tau)}{n} \right)^{\frac{1}{2-\theta}} \tag{41}$$

with probability $P^n \geq 1 - e^{-\tau}$.

Finally, we have to bound the *approximation error* $\mathcal{R}_{L_N, P, \mathcal{F}}^* - \mathcal{R}_{L_N, P}^*$ in (41). For $f_0 = h_{P, s}$ we have with Lemma A.1 that

$$\begin{aligned}
\mathcal{R}_{L_N, P}(h_{P, s}) - \mathcal{R}_{L_N, P}^* &= \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap N} |2\eta - 1| dP_X \\
&= \sum_{j \in J_N^r} \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X.
\end{aligned}$$

We split J_N^r in indices where cells do not intersect the decision line and those which do by

$$\begin{aligned}
J_{N_1}^r &:= \{j \in J_N^r \mid P_X(A_j \cap X_1) = 0 \vee P_X(A_j \cap X_{-1}) = 0\} \\
J_{N_2}^r &:= \{j \in J_N^r \mid P_X(A_j \cap X_1) > 0 \wedge P_X(A_j \cap X_{-1}) > 0\}.
\end{aligned}$$

such that

$$\begin{aligned}
&\sum_{j \in J_N^r} \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X \\
&= \sum_{j \in J_{N_1}^r} \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X \\
&\quad + \sum_{j \in J_{N_2}^r} \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X.
\end{aligned}$$

We notice that, as in the calculation of the approximation error in Part 1, the first sum vanishes since $P_X((X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j) = 0$ for all $j \in J_{N_1}^r$. Moreover, we remark that $J_{N_2}^r$ only contains cells of width s that intersect the

decision boundary. Hence, by using the margin-noise assumption we find

$$\begin{aligned} \mathcal{R}_{L_N, P}(h_{P, s}) - \mathcal{R}_{L_N, P}^* &= \sum_{j \in \mathcal{J}_{N_2}^r} \int_{(X_1 \Delta \{h_{P, s} \geq 0\}) \cap A_j} |2\eta - 1| dP_X \\ &\leq \int_{\{\Delta_\eta(x) \leq s\}} |2\eta - 1| dP_X \\ &\leq (c_{\text{MNES}})^\beta. \end{aligned} \quad (42)$$

Altogether for the oracle inequality on N with (41) we find that

$$\mathcal{R}_{L_N, P}(h_D) - \mathcal{R}_{L_N, P}^* \leq 6(c_{\text{MNES}})^\beta + 4 \left(\frac{8V(c_4 r s^{-d} + \tau)}{n} \right)^{\frac{1}{2-\theta}} \quad (43)$$

holds with probability $P^n \geq 1 - e^{-\tau}$.

Part 3: In the last part we combine the results obtained in Part 1, the oracle inequality on F and Part 2, the oracle inequality on N . That means, with the separation in (27) we obtain with (31) and (43) for the oracle inequality on X that

$$\begin{aligned} \mathcal{R}_{L, P}(h_{D, s}) - \mathcal{R}_{L, P}^* &\leq (\mathcal{R}_{L_N, P}(h_{D, s}) - \mathcal{R}_{L_N, P}^*) + (\mathcal{R}_{L_F, P}(h_{D, s}) - \mathcal{R}_{L_F, P}^*) \\ &\leq 6(c_{\text{MNES}})^\beta + 4 \left(\frac{8V(c_4 r s^{-d} + \tau)}{n} \right)^{\frac{1}{2-\theta}} + \frac{32c_1(8^{d+1}s^{-d} + \tau)}{r^\gamma n} \end{aligned} \quad (44)$$

holds with probability $P^n \geq 1 - 2e^{-\tau}$. Since $s \in (0, 1]$ and $r \geq s$, we find that $r s^{-d} \geq 1$. Together with the fact $s^{-d}, \tau \geq 1$ and $c_4 \geq 1$ it follows that

$$\begin{aligned} \mathcal{R}_{L, P}(h_{D, s}) - \mathcal{R}_{L, P}^* &\leq 6(c_{\text{MNES}})^\beta + 4 \left(\frac{8V(c_4 r s^{-d} + \tau)}{n} \right)^{\frac{1}{2-\theta}} + \frac{32c_1(8^{d+1}s^{-d} + \tau)}{r^\gamma n} \\ &\leq 6(c_{\text{MNES}})^\beta + 4 \left(\frac{8V(c_4 \tau r s^{-d} + c_4 \tau r s^{-d})}{n} \right)^{\frac{1}{2-\theta}} + \frac{32c_1(8^{d+1}\tau s^{-d} + \tau s^{-d})}{r^\gamma n} \\ &\leq 6(c_{\text{MNES}})^\beta + 4 \left(\frac{c_5 \tau r s^{-d}}{n} \right)^{\frac{1}{2-\theta}} + \frac{c_6 \tau s^{-d}}{r^\gamma n} \\ &\leq 6(c_{\text{MNES}})^\beta + r^{\frac{1}{2-\theta}} 4 \left(\frac{c_5 \tau}{s^d n} \right)^{\frac{1}{2-\theta}} + \frac{c_6 \tau}{r^\gamma s^d n}, \end{aligned}$$

where $c_5 := 32V \max\{12\mathcal{H}^{d-1}(\{\eta = 1/2\}), 1\}$ and $c_6 := 64 \cdot 8^{d+1} \max\{c_{LC}, 2^\gamma\}$. Thus, inserting r , defined in (25), with the choice of $\tilde{c}_{\alpha, \gamma, d} :=$

$\left(\frac{(\gamma(2-\theta)c_6)^{2-\theta}}{4^{2-\theta}c_5}\right)^{\frac{1}{1+\gamma(2-\theta)}}$ minimizes the right-hand side and yields

$$\begin{aligned}
& \mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^* \\
& \leq 6(c_{\text{MNE}}s)^\beta + r^{\frac{1}{2-\theta}} 4 \left(\frac{c_5\tau}{s^d n}\right)^{\frac{1}{2-\theta}} + \frac{c_6\tau}{r^\gamma s^d n} \\
& = 6(c_{\text{MNE}}s)^\beta + 4(\tilde{c}_{\alpha,\gamma,d}c_5)^{\frac{1}{2-\theta}} \left(\frac{\tau}{s^d n}\right)^{\frac{2-\theta+\gamma(2-\theta)}{(1+\gamma(2-\theta))(2-\theta)}} + \frac{c_6}{\tilde{c}_{\alpha,\gamma,d}^\gamma} \left(\frac{\tau}{s^d n}\right)^{\frac{1+\gamma}{1+\gamma(2-\theta)}} \\
& = 6(c_{\text{MNE}}s)^\beta + \left(\frac{\tilde{c}_{\alpha,\gamma,d}^{\frac{1+\gamma(2-\theta)}{2-\theta}} 4c_5^{\frac{1}{2-\theta}} + c_6}{\tilde{c}_{\alpha,\gamma,d}^\gamma}\right) \left(\frac{\tau}{s^d n}\right)^{\frac{1+\gamma}{1+\gamma(2-\theta)}} \\
& = 6(c_{\text{MNE}}s)^\beta + \left(\frac{\gamma(2-\theta)c_6 + c_6}{\tilde{c}_{\alpha,\gamma,d}^\gamma}\right) \left(\frac{\tau}{s^d n}\right)^{\frac{1+\gamma}{1+\gamma(2-\theta)}} \\
& \leq 6(c_{\text{MNE}}s)^\beta + \left(\frac{2c_6 \max\{\gamma(2-\theta), 1\}}{\tilde{c}_{\alpha,\gamma,d}^\gamma}\right) \left(\frac{\tau}{s^d n}\right)^{\frac{1+\gamma}{1+\gamma(2-\theta)}}
\end{aligned}$$

and we find again by inserting θ that

$$\mathcal{R}_{L,P}(h_{D,s}) - \mathcal{R}_{L,P}^* \leq 6(c_{\text{MNE}}s)^\beta + c_{\alpha,\gamma,d} \left(\frac{\tau}{s^d n}\right)^{\frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2}} \quad (45)$$

holds with probability $P^n \geq 1 - 2e^{-\tau}$, where $c_{\alpha,\gamma,d} := \frac{2c_6 \max\{\gamma(2-\theta), 1\}}{\tilde{c}_{\alpha,\gamma,d}^\gamma} = \frac{2c_6 \max\{\frac{\gamma(\alpha+2\gamma)}{\alpha+\gamma}, 1\}}{\tilde{c}_{\alpha,\gamma,d}^\gamma}$. \square

Proof of Theorem 3.6: We begin by proving that the chosen sequence s_n satisfies assumptions (9) and (10). To this end, we define $n_{\tau,\alpha,\beta,\gamma,d} := \left(\frac{\tilde{c}_{\alpha,\beta,\gamma,\tau,d}}{c_1}\right)^{\frac{1}{\zeta_1}}$ with $c_1 := \tilde{c}_{\alpha,\gamma,d}^{\frac{\kappa+\gamma^2}{\kappa+\gamma^2+d\gamma}} \tau^{\frac{\gamma}{\kappa+\gamma^2+d\gamma}}$, where $\tilde{c}_{\alpha,\gamma,d}$ is the constant from Theorem 3.5, and $\zeta_1 := \frac{\kappa(\kappa+\gamma^2+d\gamma)-\gamma(\beta(\kappa+\gamma^2)+d\kappa)}{(\beta(\kappa+\gamma^2)+d\kappa)(\kappa+\gamma^2+d\gamma)}$. We remark that $\zeta_1 \geq 0$ since we find by $\beta \leq \gamma^{-1}(1+\gamma)(\alpha+\gamma)$ that

$$\begin{aligned}
\kappa(\kappa+\gamma^2+d\gamma) - \gamma(\beta(\kappa+\gamma^2)+d\kappa) &= \kappa^2 + \kappa\gamma^2 - \gamma\beta\kappa - \beta\gamma^3 \\
&\geq \kappa^2 + \kappa\gamma^2 - \kappa^2 - \kappa\gamma^2 \\
&= 0.
\end{aligned}$$

Then, for $n \geq n_{\tau,\alpha,\beta,\gamma,d}$ a simple calculation shows that the latter is equivalent to

$$c_1 n^{-\frac{\gamma}{\kappa+\gamma^2+d\gamma}} \geq \tilde{c}_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\kappa}{\beta(\kappa+\gamma^2)+d\kappa}},$$

which equals assumption (9) with $s_n := \tilde{c}_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\kappa}{\beta(\kappa+\gamma^2)+d\kappa}}$. To see that assumption (10) is satisfied we define $\tilde{n}_{\tau,\alpha,\beta,\gamma,d} := \left(\frac{c_2}{\tilde{c}_{\alpha,\beta,\gamma,\tau,d}}\right)^{\frac{1}{\zeta_2}}$ with $c_2 :=$

$\tau^{\frac{1}{d}} \left(\frac{\tilde{c}_{\alpha,\gamma,d}}{\min\{\frac{\delta^*}{3}, 1\}} \right)^{\frac{\kappa+\gamma^2}{d\gamma}}$, where $\tilde{c}_{\alpha,\gamma,d}$ is the constant from Theorem 3.5, δ^* the one from Lemma 2.1 and where $\zeta_2 := \frac{\beta(\kappa+\gamma^2)}{d(\beta(\kappa+\gamma^2)+d\kappa)}$. Then, a simple transformation shows again that for all $n \geq \tilde{n}_{\tau,\alpha,\beta,\gamma,d}$ we find

$$\tilde{c}_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\kappa}{(\beta(\kappa+\gamma^2)+d\kappa)}} \geq c_2 n^{-\frac{1}{d}},$$

which equals assumption (10) with $s_n := \tilde{c}_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\kappa}{(\beta(\kappa+\gamma^2)+d\kappa)}}$.

Finally, we obtain for all $n \geq n_0 := \lceil \max\{n_{\tau,\alpha,\beta,\gamma,d}, \tilde{n}_{\tau,\alpha,\beta,\gamma,d}\} \rceil$ by inserting our chosen sequence s_n , satisfying (9) and (10), in (11) that

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D,s_n}) - \mathcal{R}_{L,P}^* \\ & \leq 6(c_{\text{MNE}} s_n)^\beta + c_{\alpha,\gamma,d} \left(\frac{\tau}{s_n^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \\ & = 6c_{\text{MNE}}^\beta \tilde{c}_{\alpha,\beta,\gamma,\tau,d}^\beta n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} + c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}} \tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{-\frac{d\kappa}{\kappa+\gamma^2}} n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \\ & = \left(\frac{6c_{\text{MNE}}^\beta \tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{\frac{\beta(\kappa+\gamma^2)+d\kappa}{\kappa+\gamma^2}} + c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}}}{\tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{\frac{d\kappa}{\kappa+\gamma^2}}} \right) n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \\ & = \left(\frac{\frac{d\kappa}{\beta(\kappa+\gamma^2)} c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}} + c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}}}{\tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{\frac{d\kappa}{\kappa+\gamma^2}}} \right) n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \\ & \leq \left(\frac{2 \max\left\{ \frac{d\kappa}{\beta(\kappa+\gamma^2)}, 1 \right\} c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}}}{\tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{\frac{d\kappa}{\kappa+\gamma^2}}} \right) n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \\ & = c_{\alpha,\beta,\gamma,\tau,d} n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \end{aligned}$$

holds with probability $P^n \geq 1 - 2e^{-\tau}$, where $c_{\alpha,\beta,\gamma,\tau,d} :=$

$$2 \max\left\{ \frac{d\kappa}{\beta(\kappa+\gamma^2)}, 1 \right\} c_{\alpha,\gamma,d} \tau^{\frac{\kappa}{\kappa+\gamma^2}} \cdot \tilde{c}_{\alpha,\beta,\gamma,\tau,d}^{\frac{d\kappa}{\kappa+\gamma^2}}. \quad \square$$

Proof of Theorem 3.7: Let s_n^* behave as s_n in Theorem 3.6, that is $s_n^* \sim n^{-\frac{\kappa}{\beta(\kappa+\gamma^2)+d\kappa}}$. We assume that $S_n := \{s_1^{(n)}, \dots, s_l^{(n)}\}$ and $s_{i-1}^{(n)} < s_i^{(n)}$ for $i \in \{2, \dots, l\}$. Since S_n is a $n^{-1/d}$ -net we have

$$s_i^{(n)} - s_{i-1}^{(n)} \leq 2n^{-1/d}. \quad (46)$$

Furthermore, there exists indices $i \in \{1, \dots, l\}$ such that $s_{i-1}^{(n)} \leq s_n^* \leq s_i^{(n)}$. An analogous calculation as at the beginning of the proof of Theorem 3.6 shows then that $s_{i-1}^{(n)}$ and $s_i^{(n)}$ satisfy assumption (9) and (10) for sufficiently large n . Hence, we find for $s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}$ with Theorem 3.5 that

$$\mathcal{R}_{L,P}(h_{D_1,s}) - \mathcal{R}_{L,P}^* \leq 6(c_{\text{MNE}} s)^\beta + c_{\alpha,\gamma,d} \left(\frac{\tau}{s^d k} \right)^{\frac{\kappa}{\kappa+\gamma^2}}. \quad (47)$$

holds with $P^k \geq 1 - 4e^{-\tau}$. Since $h_{D_1, s_{D_2}^*}$ is an ERM we find with [10, Theorem 7.2, Theorem 8.24 and Exercise 8.5] and $\tau_n := \tau + \log(1 + 2|S_n|)$ that

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D_1, s_{D_2}^*}) - \mathcal{R}_{L,P}^* \\ & < 6 \inf_{s \in S_n} (\mathcal{R}_{L,P}(h_{D_1, s}) - \mathcal{R}_{L,P}^*) + 4 \left(\frac{8c_{\alpha, \gamma} \tau_n}{n-k} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} \\ & \leq \inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} (\mathcal{R}_{L,P}(h_{D_1, s}) - \mathcal{R}_{L,P}^*) + 4 \left(\frac{8c_{\alpha, \gamma} \tau_n}{n-k} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} \end{aligned} \quad (48)$$

holds with $P^{n-k} \geq 1 - e^{-\tau}$. Combining (47) and (48) we obtain with $k \geq n/2$ and $n-k = n/2 + n/2 - k \geq n/4$ that

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D_1, s_{D_2}^*}) - \mathcal{R}_{L,P}^* \\ & \leq 6 \inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} \left(6(c_{\text{MNES}})^\beta + c_{\alpha, \gamma, d} \left(\frac{2\tau}{s^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \right) + 4 \left(\frac{32c_{\alpha, \gamma} \tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} \\ & \leq c_1 \left(\inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} \left(s^\beta + \left(\frac{2\tau}{s^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \right) + \left(\frac{\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} \right) \end{aligned}$$

holds with $P^n \geq 1 - (1+4)e^{-\tau}$. With Lemma A.3 we find that

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D_1, s_{D_2}^*}) - \mathcal{R}_{L,P}^* \\ & \leq c_1 \left(\inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} \left(s^\beta + \left(\frac{2\tau}{s^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \right) + \left(\frac{\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} \right) \\ & \leq c_2 \left((s_n^*)^\beta + \left(\frac{2\tau}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + \left(\frac{\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} + n^{-\frac{\beta}{d}} \right) \\ & \leq c_2 \left((s_n^*)^\beta + \left(\frac{2\tau_n}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + \left(\frac{2\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} + n^{-\frac{\beta}{d}} \right) \end{aligned} \quad (49)$$

holds with $P^n \geq 1 - 5e^{-\tau}$. Next, it is easy to verify with $\beta \leq \gamma^{-1}\kappa$ that

$$\left(\frac{2\tau_n}{\tilde{s}_n^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \geq \left(\frac{2\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}}$$

such that we can omit the latter right-hand side term in (49). Hence, we have that

$$\begin{aligned} & \mathcal{R}_{L,P}(h_{D_1, s_{D_2}^*}) - \mathcal{R}_{L,P}^* \\ & \leq c_2 \left((s_n^*)^\beta + \left(\frac{2\tau_n}{(s_n^*)^d n} \right)^{\frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2}} + \left(\frac{2\tau_n}{n} \right)^{\frac{\alpha+\gamma}{\alpha+2\gamma}} + n^{-\frac{\beta}{d}} \right) \end{aligned}$$

$$\begin{aligned} &\leq c_3 \left((s_n^*)^\beta + \left(\frac{2\tau_n}{(s_n^*)^d n} \right)^{\frac{(1+\gamma)(\alpha+\gamma)}{(1+\gamma)(\alpha+\gamma)+\gamma^2}} + n^{-\frac{\beta}{d}} \right) \\ &\leq c_3 \left(n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} + n^{-\frac{\beta}{d}} \right) \\ &\leq c_4 \cdot n^{-\frac{\beta\kappa}{\beta(\kappa+\gamma^2)+d\kappa}} \end{aligned}$$

holds with $P^n \geq 1 - 5e^{-\tau}$, where in the last step we used that

$$\frac{\beta\kappa}{\beta(\kappa + \gamma^2) + d\kappa} \leq \frac{\beta\kappa}{d\kappa} = \frac{\beta}{d}. \quad \square$$

Appendix A: Appendix

Lemma A.1. *Let $Y := \{-1, 1\}$ and P be a probability measure on $X \times Y$. For $\eta(x) := P(y = 1|x), x \in X$ define the set $X_1 := \{x \in X \mid \eta(x) > 1/2\}$. Let L be the classification loss and consider for $A \subset X$ the loss $L_A(x, y, t) := \mathbf{1}_A(x)L(y, t)$, where $y \in Y, t \in \mathbb{R}$. For a measurable $f: X \rightarrow \mathbb{R}$ we then have*

$$\mathcal{R}_{L_A, P}(f) - \mathcal{R}_{L_A, P}^* = \int_{(X_1 \Delta \{f \geq 0\}) \cap A} |2\eta(x) - 1| dP_X(x),$$

where Δ denotes the symmetric difference.

Proof of Lemma A.1: It is well known, e.g., [10, Example 3.8], that

$$\begin{aligned} &\mathcal{R}_{L_A, P}(f) - \mathcal{R}_{L_A, P}^* \\ &= \int_A |2\eta(x) - 1| \cdot \mathbf{1}_{(-\infty, 0)}((2\eta(x) - 1)\text{sign}f(x)) dP_X(x). \end{aligned} \tag{50}$$

Next, for P_X -almost all $x \in A$ we have

$$\mathbf{1}_{(-\infty, 0]}((2\eta(x) - 1)\text{sign}f(x)) = 1 \Leftrightarrow (2\eta(x) - 1)\text{sign}f(x) \leq 0.$$

The latter is true if for $x \in A$ holds that $f(x) < 0$ and $\eta(x) > 1/2$ or that $f(x) \geq 0$ and $\eta(x) \leq 1/2$ or that $\eta(x) = 1/2$. However, for $\eta(x) = 1/2$ we have $|2\eta(x) - 1| = 0$ and hence this case can be ignored. Then, the latter obviously equals the set $(X_1 \Delta \{f \geq 0\}) \cap A$ and we obtain in (50)

$$\mathcal{R}_{L_A, P}(f) - \mathcal{R}_{L_A, P}^* = \int_{(X_1 \Delta \{f \geq 0\}) \cap A} |2\eta(x) - 1| dP_X(x). \quad \square$$

Lemma A.2. *Let $X := [-1, 1]^d$ and P be a probability measure on $X \times \{-1, 1\}$ with fixed version $\eta: X \rightarrow [0, 1]$ of its posterior probability. Then, if η is Hölder-continuous with exponent γ , we have that Δ_η controls the noise from above with exponent γ , that means there exists a constant $c_{UC} > 0$ such that*

$$|2\eta(x) - 1| \leq c_{UC} \Delta_\eta^\gamma(x)$$

for P_X -almost all $x \in X$.

Proof of Lemma A.2: Fix w.l.o.g. an $x \in X_1$. Then, $\eta(x) > 1/2$. Since η is Hölder-continuous with exponent γ , there exists a constant $c > 0$ such that we have

$$|2\eta(x) - 1| = 2|\eta(x) - 1/2| \leq 2|\eta(x) - \eta(x')| \leq 2c(d(x, x'))^\gamma$$

for all $x' \in X_{-1}$ and hence

$$|2\eta(x) - 1| \leq 2c \inf_{\tilde{x} \in X_{-1}} (d(x, \tilde{x}))^\gamma = c_{\text{UC}} \Delta_\eta^\gamma(x),$$

where $c_{\text{UC}} := 2c$. Obviously, the last inequality holds immediately for $x \in X$ with $\eta(x) = 1/2$. \square

Lemma A.3. Let $\beta, \gamma, \kappa, \tau$ be as in Theorem 3.6 and let $S_n := \{s_1^{(n)}, \dots, s_l^{(n)}\}$ and $s_{i-1}^{(n)}, s_i^{(n)} \in S_n$ for $i \in \{1, \dots, l\}$ be given as in the beginning of the proof of Theorem 3.7. Furthermore, let $c_1, c_2 > 0$ be constants. Then, we have

$$\begin{aligned} & \inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} \left(6(c_1 s)^\beta + c_2 \left(\frac{2\tau}{s^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \right) \\ & \leq 6(c_1 s_n^*)^\beta + c_2 \left(\frac{2\tau}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + 6 \left(\frac{2c_1}{n^{1/d}} \right)^\beta. \end{aligned}$$

Proof of Lemma A.3: For $\delta > 0$ we fix $\tilde{s} \in (0, 1]$ such that

$$6(c_1 \tilde{s})^\beta + c_2 \left(\frac{2\tau}{(\tilde{s})^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \leq 6(c_1 s_n^*)^\beta + c_2 \left(\frac{2\tau}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + \delta,$$

where s_n^* is given as in the proof of Theorem 3.7. Then, we find that $s_{i-1}^{(n)} \leq \tilde{s} \leq s_i^{(n)}$ and with (46) that

$$\tilde{s} \leq s_i^{(n)} \leq \tilde{s} + 2n^{-1/d}.$$

Hence,

$$\begin{aligned} & 6(c_1 s_i^{(n)})^\beta + c_2 \left(\frac{2\tau}{(s_i^{(n)})^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \\ & \leq 6(c_1 (\tilde{s} + 2n^{-1/d}))^\beta + c_2 \left(\frac{2\tau}{(\tilde{s})^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \\ & \leq 6(c_1 \tilde{s})^\beta + c_2 \left(\frac{2\tau}{(\tilde{s})^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \\ & \leq (c_1 s_n^*)^\beta + c_2 \left(\frac{2\tau}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + \delta + 6 \left(\frac{2c_1}{n^{1/d}} \right)^\beta \end{aligned}$$

and finally

$$\begin{aligned} & \inf_{s \in \{s_{i-1}^{(n)}, s_i^{(n)}\}} \left(6 (c_1 s)^\beta + c_2 \left(\frac{2\tau}{s^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} \right) \\ & \leq 6 (c_1 s_n^*)^\beta + c_2 \left(\frac{2\tau}{(s_n^*)^d n} \right)^{\frac{\kappa}{\kappa+\gamma^2}} + 6 \left(\frac{2c_1}{n^{1/d}} \right)^\beta. \quad \square \end{aligned}$$

Acknowledgement

We thank the reviewers for their valuable remarks, which led to a better understanding of the influence and interplay of our margin conditions.

References

- [1] AUDIBERT, J. Y. and TSYBAKOV, A. (2007). Fast learning rates for plug-in classifiers. *Ann. Statist.* **35** 608–633. [MR2336861](#)
- [2] BINEV, P., COHEN, A., DAHMEN, W. and DEVORE, R. (2014). Classification algorithms using adaptive partitioning. *Ann. Statist.* **42** 2141–2163. [MR3269976](#)
- [3] BREIMAN, L. (2001). Random Forests. *Machine Learning* **45** 5–32.
- [4] DEVROYE, L., GYÖRFI, L. and LUGOSI, L. (1996). *A Probabilistic Theory of Pattern Recognition*. Springer. [MR1383093](#)
- [5] DÖRING, M., GYÖRFI, L. and WALK, H. (2015). *Exact rate of convergence of kernel-based classification rule* In *Challenges in Computational Statistics and Data Mining* **605** 71–91. Springer International Publishing. [MR3408541](#)
- [6] FEDERER, H. (1969). *Geometric measure theory*. Springer. [MR0257325](#)
- [7] KOHLER, M. and KRZYŻAK, A. (2007). On the rate of convergence of local averaging plug-in classification rules under a margin condition. *IEEE Trans. Inf. Theor.* **53** 1735–1742. [MR2317135](#)
- [8] MASSART, P. and NEDELEC, E. (2006). Risk bounds for statistical learning. *Ann. Statist.* **34** 2326–2366. [MR2291502](#)
- [9] STEINWART, I. (2015). Fully adaptive density-based clustering. *Ann. Statist.* **43** 2132–2167. [MR3396981](#)
- [10] STEINWART, I. and CHRISTMANN, A. (2008). *Support Vector Machines*. Springer. [MR2450103](#)