Learning theory estimates with observations from general stationary stochastic processes

Hanyuan Hang[1], Yunlong Feng[1], Ingo Steinwart[2], and Johan A.K. Suykens[1]

[1]Department of Electrical Engineering, KU Leuven, Leuven, Belgium
[2]Institute for Stochastics and Applications, University of Stuttgart, Stuttgart, Germany

Abstract

This paper investigates the supervised learning problem with observations drawn from certain general stationary stochastic processes. Here by *general*, we mean that many stationary stochastic processes can be included. We show that when the stochastic processes satisfy a generalized Bernstein-type inequality, a unified treatment on analyzing the learning schemes with various mixing processes can be conducted and a sharp oracle inequality for generic regularized empirical risk minimization schemes can be established. The obtained oracle inequality is then applied to derive convergence rates for several learning schemes such as empirical risk minimization (ERM), least squares support vector machines (LS-SVMs) using given generic kernels, and SVMs using Gaussian kernels for both least squares and quantile regression. It turns out that for i.i.d. processes, our learning rates for ERM recover the optimal rates. On the other hand, for non-i.i.d. processes including geometrically α-mixing Markov processes, geometrically α-mixing processes with restricted decay, ϕ-mixing processes, and (time-reversed) geometrically $\mathcal{C}$-mixing processes, our learning rates for SVMs with Gaussian kernels match, up to some arbitrarily small extra term in the exponent, the optimal rates. For the remaining cases, our rates are at least close to the optimal rates. As a by-product, the assumed generalized Bernstein-type inequality also provides an interpretation of the so-called "effective number of observations" for various mixing processes.

1 Introduction and Motivation

In this paper, we study the supervised learning problem which aims at inferring a functional relation between explanatory variables and response variables [46]. In the literature of statistical learning theory, one of the main research topics is the generalization ability of different learning schemes which indicate their learnabilities on future observations. Nowadays, it has been well understood that the Bernstein-type inequalities play an important role in deriving fast learning rates. For example, the analysis of various algorithms from non-parametric statistics and machine learning crucially depends on these inequalities, see e.g. [12, 13, 18, 36]. Here, stronger results can typically be achieved since the Bernstein-type inequality allows for localization due to its specific dependence on the variance. In particular, most derivations of minimax optimal learning rates are based on it.

The classical Bernstein inequality assumes that the data are generated by an i.i.d. process. Unfortunately, however, this assumption is often violated in many real-world applications including financial prediction, signal processing, system identification and diagnosis, text and speech

recognition, and time series forecasting, among others. For this and other reasons, there has been some effort to establish Bernstein-type inequalities for non-i.i.d. processes. For instance, generalizations of Bernstein-type inequalities to the cases of α-mixing [31] and $\mathcal{C}$-mixing [21] processes have been found [7, 28, 27, 32] and [21, 19], respectively. These Bernstein-type inequalities have been applied to derive various convergence rates. For example, the Bernstein-type inequality established in [7] was employed in [54] to derive convergence rates for sieve estimates from strictly stationary α-mixing processes in the special case of neural networks. [20] applied the Bernstein-type inequality in [28] to derive an oracle inequality (see Page 220 in [36] for the meaning of the *oracle inequality*) for generic regularized empirical risk minimization algorithms with stationary α-mixing processes. By applying the Bernstein-type inequality in [27], [4] derived almost sure uniform convergence rates for the estimated Lévy density both in mixed-frequency and low-frequency setups and proved their optimality in the minimax sense. Particularly, concerning the least squares loss, [2] obtained the optimal learning rates for ϕ-mixing processes by applying the Bernstein-type inequality established in [32]. By developing a Bernstein-type inequality for $\mathcal{C}$-mixing processes that include ϕ-mixing processes and many discrete-time dynamical systems, [21] established an oracle inequality as well as fast learning rates for generic regularized empirical risk minimization algorithms with observations from $\mathcal{C}$-mixing processes.

The above-mentioned inequalities are termed as Bernstein-type since they rely on the variance of the random variables. However, we note that these inequalities are usually presented in similar but rather complicated forms which consequently are not easy to apply directly in analyzing the performance of statistical learning schemes and may be also lack of interpretability. On the other hand, existing studies on learning from mixing processes may diverse from one to another since they may be conducted under different assumptions and notations, which leads to barriers in comparing the learnability of these learning algorithms.

In this work, we first introduce a generalized Bernstein-type inequality and show that it can be instantiated to various stationary mixing processes. Based on the generalized Bernstein-type inequality, we establish an oracle inequality for a class of learning algorithms including ERM [36, Chapter 6] and SVMs. On the technical side, the oracle inequality is derived by refining and extending the analysis of [37]. To be more precise, the analysis in [37] partially ignored localization with respect to the regularization term, which in our study is addressed by a carefully arranged peeling approach inspired by [36]. This leads to a sharper stochastic error bound and consequently a sharper bound for the oracle inequality, comparing with that of [37]. Besides, based on the assumed generalized Bernstein-type inequality, we also provide an interpretation and comparison of the effective numbers of observations when learning from various mixing processes.

Our second main contribution made in the present study lies in that we present a unified treatment on analyzing learning schemes with various mixing processes. For example, we establish fast learning rates for α-mixing and (time-reversed) $\mathcal{C}$-mixing processes by tailoring the generalized oracle inequality. For ERM, our results match those in the i.i.d. case, if one replaces the number of observations with the effective number of observations. For LS-SVMs, as far as we know, the best learning rates for the case of geometrically α-mixing process are those derived in [51, 43, 17]. When applied to LS-SVMs, it turns out that our oracle inequality leads to faster learning rates that those reported in [51] and [17]. For sufficiently smooth kernels, our rates are also faster than those in [43]. For other mixing processes including geometrically α-mixing Markov chains, geometrically ϕ-mixing processes, and geometrically $\mathcal{C}$-mixing processes, our rates for LS-SVMs with Gaussian kernels match essentially the optimal learning rates, while for LS-SVMs with given generic kernel, we only obtain rates that are close to the optimal rates.

The rest of this work is organized as follows: In Section 2, we introduce some basics of statistical learning theory. Section 3 presents the key assumption of a generalized Bernstein-type inequality for stationary mixing processes, and present some concrete examples that satisfy this assumption. Based on the generalized Bernstein-type inequality, a sharp oracle inequality is developed in Section 4 while its proof is deferred to the Appendix. Section 5 provides some applications of the newly developed oracle inequality. The paper is ended in Section 6.

2 A Primer in Learning Theory

Let $(X, \mathcal{X})$ be a measurable space and $Y \subset \mathbb{R}$ be a closed subset. The goal of (supervised) statistical learning is to find a function $f : X \to \mathbb{R}$ such that for $(x, y) \in X \times Y$ the value $f(x)$ is a good prediction of y at x. The following definition will help us define what we mean by "good".

Definition 1. *Let $(X, \mathcal{X})$ be a measurable space and $Y \subset \mathbb{R}$ be a closed subset. Then a function $L : X \times Y \times \mathbb{R} \to [0, \infty)$ is called a loss function, or simply a loss, if it is measurable.*

In this study, we are interested in loss functions that in some sense can be restricted to domains of the form $X \times Y \times [-M, M]$ as defined below, which is typical in learning theory [36, Definition 2.22] and is in fact motivated by the boundedness of Y.

Definition 2. *We say that a loss $L : X \times Y \times \mathbb{R} \to [0, \infty)$ can be clipped at $M > 0$, if, for all $(x, y, t) \in X \times Y \times \mathbb{R}$, we have*

$$L(x, y, \widehat{t}\,) \leq L(x, y, t),$$

where $\widehat{t}$ denotes the clipped value of t at $\pm M$, that is

$$\widehat{t} := \begin{cases} -M & \text{if } t < -M, \\ t & \text{if } t \in [-M, M], \\ M & \text{if } t > M. \end{cases}$$

Throughout this work, we make the following assumptions on the loss function L:

Assumption 1. *The loss function $L : X \times Y \times \mathbb{R} \to [0, \infty)$ can be clipped at some $M > 0$. Moreover, it is both bounded in the sense of $L(x, y, t) \leq 1$ and locally Lipschitz continuous, that is,*

$$|L(x, y, t) - L(x, y, t')| \leq |t - t'|\,. \tag{1}$$

Here both inequalites are supposed to hold for all $(x, y) \in X \times Y$ and $t, t' \in [-M, M]$.

Note that the above assumption with Lipschitz constant equals to one can typically be enforced by scaling. To illustrate the generality of the above assumptions on L, let us first consider the case of binary classification, that is $Y := \{-1, 1\}$. For this learning problem one often uses a convex surrogate for the original discontinuous classification loss $\mathbf{1}_{(-\infty,0]}(y \operatorname{sign}(t))$, since the latter may lead to computationally infeasible approaches. Typical surrogates L belong to the class of margin-based losses, that is, L is of the form $L(y, t) = \varphi(yt)$, where $\varphi : \mathbb{R} \to [0, \infty)$ is

a suitable, convex function. Then L can be clipped, if and only if φ has a global minimum, see [36, Lemma 2.23]. In particular, the hinge loss, the least squares loss for classification, and the squared hinge loss can be clipped, but the logistic loss for classification and the AdaBoost loss cannot be clipped. On the other hand, [34] established a simple technique, which is similar to inserting a small amount of noise into the labeling process, to construct a clippable modification of an arbitrary convex, margin-based loss. Finally, both the Lipschitz continuity and the boundedness of L can be easily verified for these losses, where for the latter it may be necessary to suitably scale the loss.

Bounded regression is another class of learning problems, where the assumptions made on L are often satisfied. Indeed, if $Y := [-M, M]$ and L is a convex, distance-based loss represented by some $\psi : \mathbb{R} \to [0, \infty)$, that is $L(y, t) = \psi(y - t)$, then L can be clipped whenever $\psi(0) = 0$, see again [36, Lemma 2.23]. In particular, the least squares loss

$$L(y, t) = (y - t)^2 \tag{2}$$

and the τ-pinball loss

$$L_\tau(y, t) := \psi(y - t) = \begin{cases} -(1 - \tau)(y - t), & \text{if } y - t < 0, \\ \tau(y - t), & \text{if } y - t \geq 0, \end{cases} \tag{3}$$

used for quantile regression can be clipped. Again, for both losses, the Lipschitz continuity and the boundedness can be easily enforced by a suitable scaling of the loss.

Given a loss function L and an $f : X \to \mathbb{R}$, we often use the notation $L \circ f$ for the function $(x, y) \mapsto L(x, y, f(x))$. Our major goal is to have a small average loss for future unseen observations (x, y). This leads to the following definition.

Definition 3. *Let $L : X \times Y \times \mathbb{R} \to [0, \infty)$ be a loss function and P be a probability measure on $X \times Y$. Then, for a measurable function $f : X \to \mathbb{R}$, the L-risk is defined by*

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

Moreover, the minimal L-risk

$$\mathcal{R}^*_{L,P} := \inf\{\mathcal{R}_{L,P}(f) \,|\, f : X \to \mathbb{R} \text{ measurable}\}$$

is called the Bayes risk with respect to P and L. In addition, a measurable function $f^*_{L,P} : X \to \mathbb{R}$ satisfying $\mathcal{R}_{L,P}(f^*_{L,P}) = \mathcal{R}^*_{L,P}$ is called a Bayes decision function.

Let $(\Omega, \mathcal{A}, \mu)$ be a probability space, $\mathcal{Z} := (Z_i)_{i \geq 1}$ be an $X \times Y$-valued stochastic process on $(\Omega, \mathcal{A}, \mu)$, we write

$$D := \big((X_1, Y_1), \ldots, (X_n, Y_n)\big) := (Z_1, \ldots, Z_n) \in (X \times Y)^n$$

for a training set of length n that is distributed according to the first n components of $\mathcal{Z}$. Informally, the goal of learning from a training set D is to find a decision function f_D such that $\mathcal{R}_{L,P}(f_D)$ is close to the minimal risk $\mathcal{R}^*_{L,P}$. Our next goal is to formalize this idea. We begin with the following definition.

Definition 4. *Let X be a set and $Y \subset \mathbb{R}$ be a closed subset. A learning method $\mathcal{L}$ on $X \times Y$ maps every set $D \in (X \times Y)^n$, $n \geq 1$, to a function $f_D : X \to \mathbb{R}$.*

Now a natural question is whether the functions f_D produced by a specific learning method satisfy

$$\mathcal{R}_{L,P}(f_D) \to \mathcal{R}^*_{L,P}, \qquad n \to \infty\,.$$

If this convergence takes place for all P, then the learning method is called universally consistent. In the i.i.d. case many learning methods are known to be universally consistent, see e.g. [12] for classification methods, [18] for regression methods, and [36] for generic SVMs. For consistent methods, it is natural to ask how fast the convergence rate is. Unfortunately, in most situations uniform convergence rates are impossible, see [12, Theorem 7.2], and hence establishing learning rates require some assumptions on the underlying distribution P. Again, results in this direction can be found in the above-mentioned books. In the non-i.i.d. case, [30] showed that no uniform consistency is possible if one only assumes that the data generating process $\mathcal{Z}$ is stationary and ergodic. On the other hand, if some further assumptions of the dependence structure of $\mathcal{Z}$ are made, then consistency is possible, see e.g. [40].

Let us now describe the learning algorithms of particular interest to us. To this end, we assume that we have a hypothesis set $\mathcal{F}$ consisting of bounded measurable functions $f : X \to \mathbb{R}$, which is pre-compact with respect to the supremum norm $\|\cdot\|_\infty$. Since the cardinality of $\mathcal{F}$ can be infinite, we need to recall the following concept, which will enable us to approximate $\mathcal{F}$ by using finite subsets.

Definition 5. *Let (T, d) be a metric space and $\varepsilon > 0$. We call $S \subset T$ an ε-net of T if for all $t \in T$ there exists an $s \in S$ with $d(s, t) \le \varepsilon$. Moreover, the ε-covering number of T is defined by*

$$\mathcal{N}(T, d, \varepsilon) := \inf\left\{ n \ge 1 : \exists s_1, \ldots, s_n \in T \text{ such that } T \subset \bigcup_{i=1}^n B_d(s_i, \varepsilon) \right\},$$

where $\inf \emptyset := \infty$ and $B_d(s, \varepsilon) := \{t \in T : d(t, s) \le \varepsilon\}$ denotes the closed ball with center $s \in T$ and radius ε.

Note that our hypothesis set $\mathcal{F}$ is assumed to be pre-compact, and hence for all $\varepsilon > 0$, the covering number $\mathcal{N}(\mathcal{F}, \|\cdot\|_\infty, \varepsilon)$ is finite.

Denote $D_n := \frac{1}{n}\sum_{i=1}^n \delta_{(X_i,Y_i)}$, where $\delta_{(X_i,Y_i)}$ denotes the (random) Dirac measure at (X_i, Y_i). In other words, D_n is the empirical measure associated to the data set D. Then, the risk of a function $f : X \to \mathbb{R}$ with respect to this measure

$$\mathcal{R}_{L,D_n}(f) = \frac{1}{n}\sum_{i=1}^n L(X_i, Y_i, f(X_i))$$

is called the empirical L-risk.

With these preparations we can now introduce the class of learning methods of interest:

Definition 6. *Let $L : X \times Y \times \mathbb{R} \to [0, \infty)$ be a loss that can be clipped at some $M > 0$, $\mathcal{F}$ be a hypothesis set, that is, a set of measurable functions $f : X \to \mathbb{R}$, with $0 \in \mathcal{F}$, and Υ be a regularizer on $\mathcal{F}$, that is, $\Upsilon : \mathcal{F} \to [0, \infty)$ with $\Upsilon(0) = 0$. Then, for $\delta \ge 0$, a learning method whose decision functions $f_{D_n,\Upsilon} \in \mathcal{F}$ satisfy*

$$\Upsilon(f_{D_n,\Upsilon}) + \mathcal{R}_{L,D_n}(\widehat{f}_{D_n,\Upsilon}) \le \inf_{f\in\mathcal{F}} \left(\Upsilon(f) + \mathcal{R}_{L,D_n}(f)\right) + \delta \tag{4}$$

for all $n \ge 1$ and $D_n \in (X \times Y)^n$ is called δ-approximate clipped regularized empirical risk minimization (δ-CR-ERM) with respect to L, $\mathcal{F}$, and Υ.

In the case $\delta = 0$, we simply speak of clipped regularized empirical risk minimization (CR-ERM). In this case, $f_{D_n,\Upsilon}$ in fact can be also defined as follows:

$$f_{D_n,\Upsilon} = \arg\min_{f \in \mathcal{F}} \Upsilon(f) + \mathcal{R}_{L,D_n}(f).$$

Note that on the right-hand side of (4) the unclipped loss is considered, and hence CR-ERMs do not necessarily minimize the regularized clipped empirical risk $\Upsilon(\cdot) + \mathcal{R}_{L,D_n}(\cdot)$. Moreover, in general CR-ERMs do not minimize the regularized risk $\Upsilon(\cdot) + \mathcal{R}_{L,D_n}(\cdot)$ either, because on the left-hand side of (4) the clipped function is considered. However, if we have a minimizer of the unclipped regularized risk, then it automatically satisfies (4). In particular, ERM decision functions satisfy (4) for the regularizer $\Upsilon := 0$ and $\delta := 0$, and SVM decision functions satisfy (4) for the regularizer $\Upsilon := \lambda\|\cdot\|_H^2$ and $\delta := 0$. In other words, ERM and SVMs are CR-ERMs.

3 Mixing Processes and A Generalized Bernstein-type Inequality

In this section, we introduce a generalized Bernstein-type inequality. Here the inequality is said to be *generalized* in that it depends on the effective number of observations instead of the number of observations, which, as we shall see later, makes it applicable to various stationary stochastic processes. To this end, let us first introduce several mixing processes.

3.1 Several Stationary Mixing Processes

We begin with introducing some notations. Recall that $(X, \mathcal{X})$ is a measurable space and $Y \subset \mathbb{R}$ is closed. We further denote $(\Omega, \mathcal{A}, \mu)$ as a probability space, $\mathcal{Z} := (Z_i)_{i\geq 1}$ as an $X \times Y$-valued stochastic process on $(\Omega, \mathcal{A}, \mu)$, $\mathcal{A}_1^i$ and $\mathcal{A}_{i+n}^\infty$ as the σ-algebras generated by $(Z_1, \ldots, Z_i)$ and $(Z_{i+n}, Z_{i+n+1}, \ldots)$, respectively. Throughout, we assume that $\mathcal{Z}$ is stationary, that is, the $(X \times Y)^n$-valued random variables $(Z_{i_1}, \ldots, Z_{i_n})$ and $(Z_{i_1+i}, \ldots, Z_{i_n+i})$ have the same distribution for all n, i, $i_1, \ldots, i_n \geq 1$. Let $\chi : \Omega \to X$ be a measurable map. μ_χ is denoted as the χ-image measure of μ, which is defined as $\mu_\chi(B) := \mu(\chi^{-1}(B))$, $B \subset X$ measurable. We denote $L_p(\mu)$ as the space of (equivalence classes of) measurable functions $g : \Omega \to \mathbb{R}$ with finite L_p-norm $\|g\|_p$. Then $L_p(\mu)$ together with $\|g\|_p$ forms a Banach space. Moreover, if $\mathcal{A}' \subset \mathcal{A}$ is a sub-σ-algebra, then $L_p(\mathcal{A}', \mu)$ denotes the space of all $\mathcal{A}'$-measurable functions $g \in L_p(\mu)$. ℓ_2^d denotes the space of d-dimensional sequences with finite Euclidean norm. Finally, for a Banach space E, we write B_E for its closed unit ball.

In order to characterize the mixing property of a stationary stochastic process, various notions have been introduced in the literature [9]. Several frequently considered examples are α-mixing, β-mixing and ϕ-mixing, which are, respectively, defined as follows:

Definition 7 (α**-Mixing Process).** *A stochastic process $\mathcal{Z} = (Z_i)_{i\geq 1}$ is called α-**mixing** if there holds*

$$\lim_{n\to\infty} \alpha(\mathcal{Z}, n) = 0,$$

where $\alpha(\mathcal{Z}, n)$ is the α-mixing coefficient defined by

$$\alpha(\mathcal{Z}, n) = \sup_{A \in \mathcal{A}_1^i,\, B \in \mathcal{A}_{i+n}^\infty} |\mu(A \cap B) - \mu(A)\mu(B)|.$$

Moreover, a stochastic process $\mathcal{Z}$ is called ***geometrically α-mixing****, if*

$$\alpha(\mathcal{Z}, n) \leq c \exp(-bn^{\gamma}), \qquad n \geq 1, \tag{5}$$

for some constants $b > 0$, $c \geq 0$, and $\gamma > 0$.

Definition 8 (β**-Mixing Process).** *A stochastic process $\mathcal{Z} = (Z_i)_{i \geq 1}$ is called β-**mixing** if there holds*

$$\lim_{n \to \infty} \beta(\mathcal{Z}, n) = 0,$$

where $\beta(\mathcal{Z}, n)$ is the β-mixing coefficient defined by

$$\beta(\mathcal{Z}, n) := \mathbb{E} \sup_{B \in \mathcal{A}_{i+n}^{\infty}} |\mu(B) - \mu(B|\mathcal{A}_1^i)|.$$

Definition 9 (ϕ**-Mixing Process).** *A stochastic process $\mathcal{Z} = (Z_i)_{i \geq 1}$ is called ϕ-**mixing** if there holds*

$$\lim_{n \to \infty} \phi(\mathcal{Z}, n) = 0,$$

where $\phi(\mathcal{Z}, n)$ is the ϕ-mixing coefficient defined by

$$\phi(\mathcal{Z}, n) := \sup_{A \in \mathcal{A}_1^i, B \in \mathcal{A}_{i+n}^{\infty}} |\mu(B) - \mu(B|A)|.$$

The α-mixing concept was introduced by Rosenblatt [31] while the β-mixing coefficient was introduced by [49, 50], and was attributed there to Kolmogorov. Moreover, Ibragimov [22] introduced the ϕ-coefficient, see also [23]. An extensive and thorough account on mixing concepts including β- and ϕ-mixing is also provided by [10]. It is well-known that, see e.g. [20, Section 2], the β- and ϕ-mixing sequences are also α-mixing, see Figure 1. From the above definition, it is obvious that i.i.d. processes are also geometrically α-mixing processes since (5) is satisfied for $c = 0$ and all $b, \gamma > 0$. Moreover, several time series models such as ARMA and GARCH, which are often used to describe, e.g. financial data, satisfy (5) under natural conditions [16, Chapter 2.6.1], and the same is true for many Markov chains including some dynamical systems perturbed by dynamic noise, see e.g. [48, Chapter 3.5].

Another important class of mixing processes called (time-reversed) $\mathcal{C}$-mixing processes was originally introduced in [25] and recently investigated in [21]. As shown below, it is defined in association with a function class that takes into account of the smoothness of functions and therefore could be more general in the dynamical system context. As illustrated in [25] and [21], the $\mathcal{C}$-mixing process encounters a large family of dynamical systems. Given a semi-norm $\|\cdot\|$ on a vector space E of bounded measurable functions $f : Z \to \mathbb{R}$, we define the $\mathcal{C}$**-norm** by

$$\|f\|_{\mathcal{C}} := \|f\|_{\infty} + \|f\|, \tag{6}$$

and denote the space of all bounded $\mathcal{C}$-functions by $\mathcal{C}(Z) := \left\{ f : Z \to \mathbb{R} \,\middle|\, \|f\|_{\mathcal{C}} < \infty \right\}$.

Definition 10 ($\mathcal{C}$**-Mixing Process).** *Let $\mathcal{Z} = (Z_i)_{i \geq 1}$ be a stationary stochastic process. For $n \geq 1$, the $\mathcal{C}$-mixing coefficients are defined by*

$$\phi_{\mathcal{C}}(\mathcal{Z}, n) := \sup \left\{ \operatorname{cor}(\psi, h \circ Z_{k+n}) : k \geq 1,\, \psi \in B_{L_1(\mathcal{A}_1^k, \mu)}, h \in B_{\mathcal{C}(Z)} \right\},$$

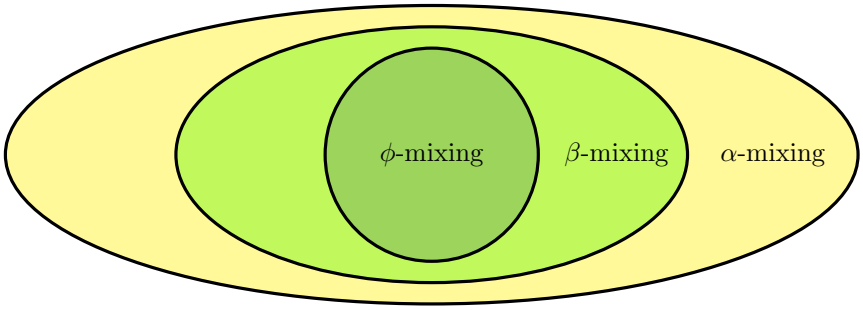

Figure 1: Relations among α-, β-, and ϕ-mixing processes

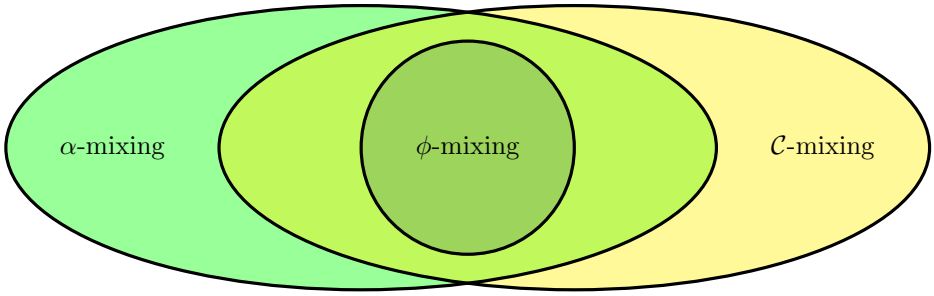

Figure 2: Relations among α-, ϕ-, and $\mathcal{C}$-mixing processes

and similarly, the time-reversed $\mathcal{C}$-mixing coefficients are defined by

$$\phi_{\mathcal{C},rev}(\mathcal{Z},n) := \sup\left\{\mathrm{cor}(h\circ Z_k,\varphi) : k\geq 1, h\in B_{\mathcal{C}(Z)}, \varphi\in B_{L_1(\mathcal{A}_{k+n}^{\infty},\mu)}\right\}.$$

*Let $(d_n)_{n\geq 1}$ be a strictly positive sequence converging to 0. Then we say that $\mathcal{Z}$ is **(time-reversed) $\mathcal{C}$-mixing** with rate $(d_n)_{n\geq 1}$, if we have $\phi_{\mathcal{C},(rev)}(\mathcal{Z},n)\leq d_n$ for all $n\geq 1$. Moreover, if $(d_n)_{n\geq 1}$ is of the form*

$$d_n := c\exp\bigl(-bn^{\gamma}\bigr), \qquad n\geq 1,$$

*for some constants $c>0$, $b>0$, and $\gamma>0$, then $\mathcal{X}$ is called **geometrically (time-reversed) $\mathcal{C}$-mixing**. If $(d_n)_{n\geq 1}$ is of the form*

$$d_n := c\cdot n^{-\gamma}, \qquad n\geq 1,$$

*for some constants $c>0$, and $\gamma>0$, then $\mathcal{X}$ is called **polynomial (time-reversed) $\mathcal{C}$-mixing**.*

Figure 2 illustrates the relations among α-mixing processes, ϕ-mixing processes, and $\mathcal{C}$-mixing processes. Clearly, ϕ-mixing processes are $\mathcal{C}$-mixing [21]. Furthermore, various discrete-time dynamical systems including Lasota-Yorke maps, uni-modal maps, and piecewise expanding maps in higher dimension are $\mathcal{C}$-mixing, see [25]. Moreover, smooth expanding maps on manifolds, piecewise expanding maps, uniformly hyperbolic attractors, and non-uniformly hyperbolic uni-modal maps are time-reversed geometrically $\mathcal{C}$-mixing, see [47, Proposition 2.7, Proposition 3.8, Corollary 4.11 and Theorem 5.15], respectively.

3.2 A Generalized Bernstein-type Inequality

As discussed in the introduction, the Bernstein-type inequality plays an important role in many areas of probability and statistics. In the statistical learning theory literature, it is also crucial in conducting concentrated estimation for learning schemes. As mentioned previously, these inequalities are usually presented in rather complicated forms under different assumptions, which therefore limit their portability to other contexts. However, what is common behind these inequalities is their relying on the boundedness assumption of the variance. Given the above discussions, in this subsection, we introduce the following generalized Bernstein-type inequality, with the hope of making it as an off-the-shelf tool for various mixing processes.

Assumption 2. *Let $\mathcal{Z} := (Z_i)_{i\geq 1}$ be an $X\times Y$-valued, stationary stochastic process on $(\Omega, \mathcal{A}, \mu)$ and $P := \mu_{Z_1}$. Furthermore, let $h : X \times Y \to \mathbb{R}$ be a bounded measurable function for which there exist constants $B > 0$ and $\sigma \geq 0$ such that $\mathbb{E}_P h = 0$, $\mathbb{E}_P h^2 \leq \sigma^2$, and $\|h\|_\infty \leq B$. Assume that, for all $\varepsilon > 0$, there exist constants $n_0 \geq 1$ independent of ε and $n_{\mathit{eff}} \geq 1$ such that for all $n \geq n_0$, we have*

$$P\left(\frac{1}{n}\sum_{i=1}^{n} h(Z_i) \geq \varepsilon\right) \leq C\exp\left(-\frac{\varepsilon^2 n_{\mathit{eff}}}{c_\sigma\sigma^2 + c_B\varepsilon B}\right), \tag{7}$$

where $n_{\mathit{eff}} \leq n$ is the effective number of observations, C is a constant independent of n, and c_σ, c_B are positive constants.

Note that in Assumption 2, the generalized Bernstein-type inequality (7) is assumed with respect to n_{eff} instead of n, which is a function of n and is termed as the *effective number of observations*. The terminology, *effective number of observations*, "provides a heuristic understanding of the fact that the statistical properties of autocorrelated data are similar to a suitably defined number of independent observations" [55]. We will continue our discussion on the effective number of observations n_{eff} in Subsection 3.4 below.

3.3 Instantiation to Various Mixing Processes

We now show that the generalized Bernstein-type inequality in Assumption 2 can be instantiated to various mixing processes, e.g., i.i.d processes, geometrically α-mixing processes, restricted geometrically α-mixing processes, geometrically α-mixing Markov chains, ϕ-mixing processes, geometrically $\mathcal{C}$-mixing processes, polynomially $\mathcal{C}$-mixing processes, among others.

3.3.1 I.I.D Processes

Clearly, the classical Bernstein inequality [5] satisfies (7) with $n_0 = 1$, $C = 1$, $c_\sigma = 2$, $c_B = 2/3$, and $n_{\mathit{eff}} = n$.

3.3.2 Geometrically α-Mixing Processes

For stationary geometrically α-mixing processes $\mathcal{Z}$, [28, Theorem 4.3] bounds the left-hand side of (7) by

$$(1 + 4e^{-2}c)\exp\left(-\frac{3\varepsilon^2 n^{(\gamma)}}{6\sigma^2 + 2\varepsilon B}\right)$$

for any $n \geq 1$, and $\varepsilon > 0$, where

$$n^{(\gamma)} := \Big\lfloor n \Big\lceil (8n/b)^{1/(\gamma+1)} \Big\rceil^{-1} \Big\rfloor,$$

where $\lfloor t \rfloor$ is the largest integer less than or equal to t and $\lceil t \rceil$ is the smallest integer greater than or equal to t for $t \in \mathbb{R}$. Observe that $\lceil t \rceil \leq 2t$ for all $t \geq 1$ and $\lfloor t \rfloor \geq t/2$ for all $t \geq 2$. From this it is easy to conclude that, for all $n \geq n_0$ with

$$n_0 := \max\{b/8, 2^{2+5/\gamma} b^{-1/\gamma}\}, \tag{8}$$

we have $n^{(\gamma)} \geq 2^{-\frac{2\gamma+5}{\gamma+1}} b^{\frac{1}{\gamma+1}} n^{\frac{\gamma}{\gamma+1}}$. Hence, the right-hand side of (7) takes the form

$$(1 + 4e^{-2}c) \exp\left(-\frac{\varepsilon^2 n^{\gamma/(\gamma+1)}}{(8^{2+\gamma}/b)^{1/(1+\gamma)}(\sigma^2 + \varepsilon B/3)} \right).$$

It is easily seen that this bound is of the generalized form (7) with n_0 given in (8), $C = 1+4e^{-2}c$, $c_\sigma = (8^{2+\gamma}/b)^{1/(1+\gamma)}$, $c_B = (8^{2+\gamma}/b)^{1/(1+\gamma)}/3$, and $n_{\mathit{eff}} = n^{\gamma/(\gamma+1)}$.

3.3.3 Restricted Geometrically α-Mixing Processes

A restricted geometrically α-mixing process is referred to as a geometrically α-mixing process (see Definition 10) with $\gamma \geq 1$. For this kind of α-mixing processes, [27, Theorem 2] established a bound for the right-hand side of (7) that takes the following form

$$c_c \exp\left(-\frac{c_b \varepsilon^2 n}{v^2 + B^2/n + \varepsilon B (\log n)^2} \right), \tag{9}$$

for all $\varepsilon > 0$ and $n \geq 2$, where c_b is some constant depending only on b, c_c is some constant depending only on c, and v^2 is defined by

$$v^2 := \sigma^2 + 2 \sum_{2 \leq i \leq n} \left| \mathrm{cov}(h(X_1), h(X_i)) \right|. \tag{10}$$

In fact, for any $\epsilon > 0$, by using Davydov's covariance inequality [11, Corollary to Lemma 2.1] with $p = q = 2 + \epsilon$ and $r = (2+\epsilon)/\epsilon$, we obtain for $i \geq 2$,

$$\begin{aligned}
\mathrm{cov}(h(Z_1), h(Z_i)) &\leq 8 \|h(Z_1)\|_{2+\epsilon} \|h(Z_i)\|_{2+\epsilon} \alpha(\mathcal{Z}, i-1)^{\epsilon/(2+\epsilon)} \\
&\leq 8 \big(\mathbb{E}_P h^{2+\epsilon} \big)^{2/(2+\epsilon)} \big(c e^{-b(i-1)} \big)^{\epsilon/(2+\epsilon)} \\
&\leq 8 c^{\epsilon/(2+\epsilon)} B^{2\epsilon/(2+\epsilon)} \sigma^{2 \cdot 2/(2+\epsilon)} \exp\big(-b\epsilon(i-1)/(2+\epsilon) \big).
\end{aligned}$$

Consequently, we have

$$\begin{aligned}
v^2 &\leq \sigma^2 + 16 c^{\epsilon/(2+\epsilon)} B^{2\epsilon/(2+\epsilon)} \sigma^{2 \cdot 2/(2+\epsilon)} \sum_{2 \leq i \leq n} \exp\big(-b\epsilon(i-1)/(2+\epsilon) \big) \\
&\leq \sigma^2 + 16 c^{\epsilon/(2+\epsilon)} B^{2\epsilon/(2+\epsilon)} \sigma^{2 \cdot 2/(2+\epsilon)} \sum_{i \geq 1} \exp(-b\epsilon i/(2+\epsilon)).
\end{aligned}$$

Setting

$$c_\epsilon := 16 c^{\epsilon/(2+\epsilon)} B^{2\epsilon/(2+\epsilon)} \sum_{i \geq 1} \exp(-b\epsilon i/(2+\epsilon)), \tag{11}$$

then the probability bound (9) can be reformulated as

$$c_c \exp\left(-\frac{c_b \varepsilon^2 n}{\sigma^2 + c_\epsilon \sigma^{4/(2+\epsilon)} + B^2/n + \varepsilon B (\log n)^2}\right).$$

When $n \geq n_0$ with

$$n_0 := \max\left\{3, \exp\left(\sqrt{c_\epsilon}\sigma^{-\epsilon/(2+\epsilon)}\right), B^2/\sigma^2\right\},$$

it can be further upper bounded by

$$c_c \exp\left(-\frac{c_b \varepsilon^2 (n/(\log n)^2)}{3\sigma^2 + \varepsilon B}\right).$$

Therefore, the Bernstein-type inequality for the restricted $\mathcal{C}$-mixing process is also of the generalized form (7) where $C = c_c$, $c_\sigma = 3/c_b$, $c_B = 1/c_b$, and $n_{\mathit{eff}} = n/(\log n)^2$.

3.3.4 Geometrically α-Mixing Markov Chains

For the stationary geometrically α-mixing Markov chain with centered and bounded random variables, [1] bounds the left-hand side of (7) by

$$\exp\left(-\frac{n\varepsilon^2}{\tilde{\sigma}^2 + \varepsilon B \log n}\right),$$

where $\tilde{\sigma}^2 = \lim_{n\to\infty} \frac{1}{n} \cdot \mathrm{Var} \sum_{i=1}^n h(X_i)$.

Following the similar arguments as in the restricted geometrically $\mathcal{C}$-mixing case, we know that for an arbitrary $\epsilon > 0$, there holds

$$\begin{aligned}
\mathrm{Var} \sum_{i=1}^n h(X_i) &= n\sigma^2 + 2 \sum_{1 \leq i < j \leq n} |\mathrm{cov}(h(X_i), h(X_j))| \\
&\leq n\sigma^2 + 2n \sum_{2 \leq i \leq n} |\mathrm{cov}(h(X_1), h(X_i))| \\
&= n \cdot v^2 \leq n\left(\sigma^2 + c_\epsilon \sigma^{4/(2+\epsilon)}\right),
\end{aligned}$$

where v^2 is defined in (10) and c_ϵ is given in (11). Consequently the following inequality holds

$$\exp\left(-\frac{n\varepsilon^2}{\sigma^2 + c_\epsilon \sigma^{4/(2+\epsilon)} + \varepsilon B \log n}\right) \leq \exp\left(-\frac{\varepsilon^2 (n/\log n)}{2\sigma^2 + \varepsilon B}\right),$$

for $n \geq n_0$ with

$$n_0 := \max\left\{3, \exp\left(c_\epsilon \sigma^{-2\epsilon/(2+\epsilon)}\right)\right\}.$$

That is, when $n \geq n_0$ the Bernstein-type inequality for the geometrically α-mixing Markov chain can be also formulated as the generalized form (7) with $C = 1$, $c_\sigma = 2$, $c_B = 1$, and $n_{\mathit{eff}} = n/\log n$.

3.3.5 ϕ-Mixing Processes

For a ϕ-mixing process $\mathcal{Z}$, [32] provides the following bound for the left-hand side of (7)

$$\exp\left(-\frac{\varepsilon^2 n}{8c_\phi(4\sigma^2+\varepsilon B)}\right),$$

where $c_\phi := \sum_{k=1}^{\infty}\sqrt{\phi(\mathcal{Z},k)}$. Obviously, it is of the general form (7) with $n_0 = 1$, $C = 1$, $c_\sigma = 32c_\phi$, $c_B = 8c_\phi$, and $n_{\mathit{eff}} = n$.

3.3.6 Geometrically $\mathcal{C}$-Mixing Processes

For the geometrically $\mathcal{C}$-mixing process in Definition 10, [21] recently developed a Bernstein-type inequality. To state the inequality, the following assumption on the semi-norm $\|\cdot\|$ in (6) is needed

$$\left\|e^f\right\| \le \left\|e^f\right\|_\infty \|f\|, \quad f \in \mathcal{C}(Z).$$

Under the above restriction on the semi-norm $\|\cdot\|$ in (6), and the assumptions that $\|h\| \le A$, $\|h\|_\infty \le B$, and $\mathbb{E}_P h^2 \le \sigma^2$, [21] states that when $n \ge n_0$ with

$$n_0 := \max\left\{\min\left\{m \ge 3 : m^2 \ge 808c(3A+B)/B \text{ and } m/(\log m)^{2/\gamma} \ge 4\right\}, e^{3/b}\right\}. \tag{12}$$

the right-hand side of (7) takes the form

$$2\exp\left(-\frac{\varepsilon^2 n/(\log n)^{2/\gamma}}{8(\sigma^2+\varepsilon B/3)}\right). \tag{13}$$

It is easy to see that (13) is also of the generalized form (7) with n_0 given in (12), $C = 2$, $c_\sigma = 8$, $c_B = 8/3$, and $n_{\mathit{eff}} = n/(\log n)^{2/\gamma}$.

3.3.7 Polynomially $\mathcal{C}$-Mixing Processes

For the polynomially $\mathcal{C}$-mixing processes, a Bernstein-type inequality was established recently in [19]. Under the same restriction on the semi-norm $\|\cdot\|$ and assumption on h as in the geometrically $\mathcal{C}$-mixing case, it states that when $n \ge n_0$ with

$$n_0 := \max\left\{\left(808c(3A+B)/B\right)^{1/2}, 4^{(\gamma+1)/(\gamma-2)}\right\}, \ \gamma > 2, \tag{14}$$

the right-hand side of (7) takes the form

$$2\exp\left(-\frac{\varepsilon^2 n^{(\gamma-2)/(\gamma+1)}}{8(\sigma^2+\varepsilon B/3)}\right).$$

An easy computation shows that it is also of the generalized form (7) with n_0 given in (14), $C = 2$, $c_\sigma = 8$, $c_B = 8/3$, and $n_{\mathit{eff}} = n^{(\gamma-2)/(\gamma+1)}$ with $\gamma > 2$.

3.4 From Observations to Effective Observations

The generalized Bernstein-type inequality in Assumption 2 is assumed with respect to the *effective number of observations* n_{eff}. As verified above, the assumed generalized Bernstein-type inequality indeed holds for many mixing processes whereas n_{eff} may take different values in different circumstances. Supposing that we have n observations drawn from a certain mixing process discussed above, Table 1 reports its effective number of observations. As mentioned above, it can be roughly treated as the number of independent observations when inferring the statistical properties of correlated data. In this subsection, we make some effort in presenting an intuitive understanding towards the meaning of the effective number of observations.

Table 1: Effective Number of Observations for Different Mixing Processes

examples	*effective number of observations*
i.i.d processes	n
geometrically α-mixing processes	$n^{\gamma/(\gamma+1)}$
restricted geometrically α-mixing processes	$n/(\log n)^2$
geometrically α-mixing Markov chains	$n/\log n$
ϕ-mixing processes	n
geometrically $\mathcal{C}$-mixing processes	$n/(\log n)^{2/\gamma}$
polynomially $\mathcal{C}$-mixing processes	$n^{(\gamma-2)/(\gamma+1)}$ with $\gamma > 2$

The terminology - *effective observations*, which may be also referred as the *effective number of observations* depending on the context, appeared probably first in [3] when studying the autocorrelated time series data. In fact, many similar concepts can be found in the literature of statistical learning from mixing processes, see e.g., [24, 52, 28, 33, 55]. For stochastic processes, *mixing* indicates the *asymptotic independence*. In some sense, the effective observations can be taken as the independent observations that can contribute when learning from a certain mixing process.



Figure 3: An illustration of the effective number of observations when inferring the statistical properties of the data drawn from mixing processes. The data of size n are split into n_{eff} blocks, each of size n/n_{eff}.

In fact, when inferring statistical properties with data drawn from mixing processes, a frequently employed technique is to split the data of size n into k blocks, each of size ℓ [53, 28, 8, 29, 21]. Each block may be constructed either by choosing consecutive points in the original observation set or by a jump selection [28, 21]. With the constructed blocks, one can then introduce a new sequence of blocks that are independent between the blocks by using the coupling technique. Due to the mixing assumption, the difference between the two sequences of blocks can be measured with respect to a certain metric. Therefore, one can deal with the independent blocks instead of dependent blocks now. On the other hand, for observations in each

originally constructed block, one can again apply the coupling technique [8, 14] to tackle, e.g., introducing ℓ new i.i.d observations and bounding the difference between the newly introduced observations and the original observations with respect to a certain metric. During this process, one tries to ensure that the number of blocks k is as large as possible, for which n_{eff} turns out to be the choice. An intuitive illustration of this procedure is shown in Fig. 3.

4 A Generalized Sharp Oracle Inequality

In this section we present one of our main results: an oracle inequality for learning from mixing processes satisfying the generalized Bernstein-type inequality (7). We first introduce a few more notations. Let $\mathcal{F}$ be a hypothesis set in the sense of Definition 6. For

$$r^* := \inf_{f\in\mathcal{F}} \Upsilon(f) + \mathcal{R}_{L,P}(\widehat{f}\,) - \mathcal{R}^*_{L,P}, \tag{15}$$

and $r > r^*$, we write

$$\mathcal{F}_r := \left\{ f \in \mathcal{F} : \Upsilon(f) + \mathcal{R}_{L,P}(\widehat{f}\,) - \mathcal{R}^*_{L,P} \le r \right\}. \tag{16}$$

Since $L(x,y,0) \le 1$, $0 \in \mathcal{F}$, and $\Upsilon(0) = 0$, then we have $r^* \le 1$, Furthermore, we assume that there exists a function $\varphi : (0,\infty) \to (0,\infty)$ that satisfies

$$\ln \mathcal{N}(\mathcal{F}_r, \|\cdot\|_\infty, \varepsilon) \le \varphi(\varepsilon) r^p \tag{17}$$

for all $\varepsilon > 0$, $r > 0$ and a suitable constant $p \in (0,1]$. Note that there are actually many hypothesis sets satisfying Assumption (17), see Section 5 for some examples.

Now, we present the oracle inequality as follows:

Theorem 1. *Let $\mathcal{Z}$ be a stochastic process satisfying Assumption 2 with constants $n_0 \ge 1$, $C > 0$, $c_\sigma > 0$, and $c_B > 0$. Furthermore, let L be a loss satisfying Assumption 1. Moreover, assume that there exists a Bayes decision function $f^*_{L,P}$ and constants $\vartheta \in [0,1]$ and $V \ge 1$ such that*

$$\mathbb{E}_P(L\circ\widehat{f} - L\circ f^*_{L,P})^2 \le V \cdot \left(\mathbb{E}_P(L\circ\widehat{f} - L\circ f^*_{L,P})\right)^\vartheta, \qquad f \in \mathcal{F}, \tag{18}$$

where $\mathcal{F}$ is a hypothesis set with $0 \in \mathcal{F}$. We define r^* and $\mathcal{F}_r$ by (15) and (16), respectively and assume that (17) is satisfied. Finally, let $\Upsilon : \mathcal{F} \to [0,\infty)$ be a regularizer with $\Upsilon(0) = 0$, $f_0 \in \mathcal{F}$ be a fixed function, and $B_0 \ge 1$ be a constant such that $\|L\circ f_0\|_\infty \le B_0$. Then, for all fixed $\varepsilon > 0$, $\delta \ge 0$, $\tau \ge 1$, $n \ge n_0$, and $r \in (0,1]$ satisfying

$$r \ge \max\left\{ \left(\frac{c_V(\tau + \varphi(\varepsilon/2)2^p r^p)}{n_{\mathit{eff}}} \right)^{\frac{1}{2-\vartheta}}, \frac{8c_B B_0 \tau}{n_{\mathit{eff}}}, r^* \right\} \tag{19}$$

with $c_V := 64(4c_\sigma V + c_B)$, every learning method defined by (4) satisfies with probability μ not less than $1 - 8Ce^{-\tau}$:

$$\Upsilon(f_{D_n,\Upsilon}) + \mathcal{R}_{L,P}(\widehat{f}_{D_n,\Upsilon}) - \mathcal{R}^*_{L,P} < 2\Upsilon(f_0) + 4\mathcal{R}_{L,P}(f_0) - 4\mathcal{R}^*_{L,P} + 4r + 5\varepsilon + 2\delta. \tag{20}$$

The proof of Theorem 1 will be provided in the Appendix. Before we illustrate this oracle inequality in the next section with various examples, let us briefly discuss the variance bound (18). For example, if $Y = [-M, M]$ and L is the least squares loss, then it is well-known that (18) is satisfied for $V := 16M^2$ and $\vartheta = 1$, see e.g. [36, Example 7.3]. Moreover, under some assumptions on the distribution P, [38] established a variance bound of the form (18) for the so-called pinball loss used for quantile regression. In addition, for the hinge loss, (18) is satisfied for $\vartheta := q/(q+1)$, if Tsybakov's noise assumption [45, Proposition 1] holds for q, see [36, Theorem 8.24]. Finally, based on [6], [34] established a variance bound with $\vartheta = 1$ for the earlier mentioned clippable modifications of strictly convex, twice continuously differentiable margin-based loss functions.

We remark that in Theorem 1 the constant B_0 is necessary since the assumed boundedness of L only guarantees $\|L \circ \widehat{f}\|_\infty \leq 1$, while B_0 bounds the function $L \circ f_0$ for an *unclipped* $f_0 \in \mathcal{F}$. We do not assume that all $f \in \mathcal{F}$ satisfy $\widehat{f} = f$, therefore in general B_0 is necessary. We refer to Examples 2, 3 and 4 for situations, where B_0 is significantly larger than 1.

5 Applications to Statistical Learning

To illustrate the oracle inequality developed in Section 4, we now apply it to establish learning rates for some algorithms including ERM over finite sets and SVMs using either a given generic kernel or a Gaussian kernel with varying widths. In the ERM case, our results match those in the i.i.d. case, if one replaces the number of observations n with the effective number of observations n_{eff} while, for LS-SVMs with given generic kernels, our rates are slightly worse than the recently obtained optimal rates [41] for i.i.d. observations. The latter difference is not surprising when considering the fact that [41] used heavy machinery from empirical process theory such as Talagrand's inequality and localized Rademacher averages while our results only use a light-weight argument based on the generalized Bernstein-type inequality and the peeling method. However, when using Gaussian kernels, we indeed recover the optimal rates for LS-SVMs and SVMs for quantile regression with i.i.d. observations.

Let us now present the first example, that is, the empirical risk minimization scheme over a finite hypothesis set.

Example 1 (ERM)**.** *Let $\mathcal{Z}$ be a stochastic process satisfying Assumption 2 and the hypothesis set $\mathcal{F}$ be finite with $0 \in \mathcal{F}$ and $\Upsilon(f) = 0$ for all $f \in \mathcal{F}$. Moreover, assume that $\|f\|_\infty \leq M$ for all $f \in \mathcal{F}$. Then, for accuracy $\delta := 0$, the learning method described by (4) is ERM, and Theorem 1 shows by some simple estimates that*

$$\mathcal{R}_{L,P}(f_{D_n,\Upsilon}) - \mathcal{R}^*_{L,P} \leq 4 \inf_{f \in \mathcal{F}} \left(\mathcal{R}_{L,P}(f) - \mathcal{R}^*_{L,P} \right) + 4 \left(\frac{c_V(\tau + \ln|\mathcal{F}|)}{n_{\mathit{eff}}} \right)^{1/(2-\vartheta)} + \frac{32 c_B \tau}{n_{\mathit{eff}}}$$

hold with probability μ not less than $1 - 8Ce^{-\tau}$.

Recalling that for the i.i.d. case we have $n_{\mathit{eff}} = n$, therefore, in Example 1 the oracle inequality (20) is thus an exact analogue to standard oracle inequality for ERM learning from i.i.d. processes (see e.g. [36, Theorem 7.2]), albeit with different constants.

For further examples let us begin by briefly recalling SVMs [36]. To this end, let X be a measurable space, $Y := [-1, 1]$ and k be a measurable (reproducing) kernel on X with reproducing kernel Hilbert space (RKHS) H. Given a regularization parameter $\lambda > 0$ and a convex

loss L, SVMs find the unique solution

$$f_{D_n,\lambda} = \operatorname*{arg\,min}_{f\in H} \left(\lambda\|f\|_H^2 + \mathcal{R}_{L,D_n}(f)\right).$$

In particular, SVMs using the least squares loss (2) are called least squares SVMs (LS-SVMs) [44] where a primal-dual characterization is given, and also termed as kernel ridge regression in the case of zero bias term as studied in [36]. SVMs using the τ-pinball loss (3) are called SVMs for quantile regression. To describe the approximation properties of H, we further need the approximation error function

$$A(\lambda) := \inf_{f\in H} \left(\lambda\|f\|_H^2 + \mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^*\right), \qquad \lambda > 0, \tag{21}$$

and denote $f_{P,\lambda}$ as the population version of $f_{D_n,\lambda}$, which is given by

$$f_{P,\lambda} := \operatorname*{arg\,min}_{f\in H} \left(\lambda\|f\|_H^2 + \mathcal{R}_{L,P}(f) - \mathcal{R}_{L,P}^*\right). \tag{22}$$

The next example discusses learning rates for LS-SVMs using a given generic kernel.

Example 2 (Generic Kernels)**.** *Let $(X,\mathcal{X})$ be a measurable space, $Y = [-1,1]$, and $\mathcal{Z}$ be a stochastic process satisfying Assumption 2. Furthermore, let L be the least squares loss and H be an RKHS over X such that the closed unit ball B_H of H satisfies*

$$\ln \mathcal{N}(B_H, \|\cdot\|_\infty, \varepsilon) \le a\varepsilon^{-2p}, \qquad \varepsilon > 0,$$

for some constants $p \in (0,1]$ and $a > 0$. In addition, assume that the approximation error function satisfies $A(\lambda) \le c\lambda^\beta$ for some $c > 0$, $\beta \in (0,1]$, and all $\lambda > 0$.

Recall that for SVMs we always have $f_{D_n,\lambda} \in \lambda^{-1/2}B_H$, see [36, (5.4)]. Consequently we only need to consider the hypothesis set $\mathcal{F} = \lambda^{-1/2}B_H$. Then, (16) implies that $\mathcal{F}_r \subset r^{1/2}\lambda^{-1/2}B_H$ and consequently we find

$$\ln \mathcal{N}(\mathcal{F}_r, \|\cdot\|_\infty, \varepsilon) \le a\lambda^{-p}\varepsilon^{-2p}r^p. \tag{23}$$

Thus, we can define the function φ in (17) as $\varphi(\varepsilon) := a\lambda^{-p}\varepsilon^{-2p}$. For the least squares loss, the variance bound (18) is valid with $\vartheta = 1$, hence the condition (19) is satisfied if

$$r \ge \max\left\{\left(c_V 2^{1+3p} a\right)^{\frac{1}{1-p}} \lambda^{-\frac{p}{1-p}} n_{\mathit{eff}}^{-\frac{1}{1-p}} \varepsilon^{-\frac{2p}{1-p}}, \frac{2c_V\tau}{n_{\mathit{eff}}}, \frac{8c_B B_0 \tau}{n_{\mathit{eff}}}, r^*\right\}. \tag{24}$$

Therefore, let r be the sum of the terms on the right-hand side. Since for large n the first and next-to-last term in (24) dominate, the oracle inequality (20) becomes

$$\begin{aligned}\lambda\|f_{D_n,\lambda}\|_H^2 + \mathcal{R}_{L,P}(\widehat{f}_{D_n,\lambda}) - \mathcal{R}_{L,P}^* &\le 4\lambda\|f_{P,\lambda}\|_H^2 + 4\mathcal{R}_{L,P}(f_{P,\lambda}) - 4\mathcal{R}_{L,P}^* + 4r + 5\varepsilon \\ &\le C\left(\lambda^\beta + \lambda^{-\frac{p}{1-p}} n_{\mathit{eff}}^{-\frac{1}{1-p}} \varepsilon^{-\frac{2p}{1-p}} + \lambda^{\beta-1} n_{\mathit{eff}}^{-1}\tau + \varepsilon\right),\end{aligned}$$

where $f_{P,\lambda}$ is defined in (22) and C is a constant independent of n, λ, τ, or ε. Now optimizing over ε, we then see by [36, Lemma A.1.7] that the LS-SVM using $\lambda_n := n_{\mathit{eff}}^{-\rho/\beta}$ learns with the rate $n_{\mathit{eff}}^{-\rho}$, where

$$\rho := \min\left\{\beta, \frac{\beta}{\beta + p\beta + p}\right\}. \tag{25}$$

In particular, for geometrically α-mixing processes, we obtain the learning rate $n^{-\alpha\rho}$, where $\alpha := \frac{\gamma}{\gamma+1}$ and ρ as in (25). Let us compare this rate with the ones previously established for LS-SVMs in the literature. For example, [37] proved a rate of the form

$$n^{-\alpha \min\{\beta, \frac{\beta}{\beta+2p\beta+p}\}}$$

under exactly the same assumptions. Since $\beta > 0$ and $p > 0$, our rate is always better than that of [37]. In addition, [17] generalized the rates of [37] to regularization terms of the form $\lambda\|\cdot\|_H^q$ with $q \in (0, 2]$. The resulting rates are again always slower than the ones established in this work. For the standard regularization term, that is $q = 2$, [51] established the rate

$$n^{-\frac{\alpha\beta}{2p+1}},$$

which is always slower than ours, too. Finally, in the case $p = 1$, [42] established the rate

$$n^{-\frac{2\alpha\beta}{\beta+3}},$$

which was subsequently improved to

$$n^{-\frac{3\alpha\beta}{2\beta+4}}$$

in [43]. The latter rate is worse than ours, if and only if $(1 + \beta)(1 + 3p) \leq 5$. In particular, for $p \in (0, 1/2]$ we always get better rates. Furthermore, the analysis of [42, 43] is restricted to LS-SVMs, while our results hold for rather generic learning algorithms.

Example 3 (Smooth Kernels)**.** *Let $X \subset \mathbb{R}^d$ be a compact subset, $Y = [-1, 1]$, and $\mathcal{Z}$ be a stochastic process satisfying Assumption 2. Furthermore, let L be the least squares loss and $H = W^m(X)$ be a Sobolev space with smoothness $m > d/2$. Then it is well-known, see e.g. [41] or [36, Theorem 6.26], that*

$$\ln \mathcal{N}(B_H, \|\cdot\|_\infty, \varepsilon) \leq a\varepsilon^{-2p}, \qquad \varepsilon > 0,$$

where $p := \frac{d}{2m}$ and $a > 0$ is some constant. Let us additionally assume that the marginal distribution P_X is absolutely continuous with respect to the uniform distribution, where the corresponding density is bounded away from 0 and ∞. Then there exists a constant $C_p > 0$ such that

$$\|f\|_\infty \leq C_p \|f\|_H^p \|f\|_{L_2(P_X)}^{1-p}, \qquad f \in H$$

for the same p, see [26] and [41, Corollary 3]. Consequently, we can bound $B_0 \leq \lambda^{(\beta-1)p}$ as in [41]. Moreover, the assumption on the approximation error function is satisfied for $\beta := s/m$, whenever $f^*_{L,P} \in W^s(X)$ and $s \in (0, m]$. Therefore, the resulting learning rate is

$$n_{\mathit{eff}}^{-\frac{2s}{2s+d+ds/m}}. \tag{26}$$

Note that in the i.i.d. case, where $n_{\mathit{eff}} = n$, this rate is worse than the optimal rate $n^{-\frac{2s}{2s+d}}$, where the discrepancy is the term ds/m in the denominator. However, this difference can be made arbitrarily small by picking a sufficiently large m, that is, a sufficiently smooth kernel k. Moreover, in this case, for geometrically α-mixing processes, the rate (26) becomes

$$n^{-\frac{2s\alpha}{2s+d+ds/m}},$$

where $\alpha := \frac{\gamma}{\gamma+1}$. Comparing this rate with the one from [43], it turns out that their rate is worse than ours, if $m \geq \frac{1}{16}(2s+3d+\sqrt{4s^2+108sd+9d^2})$. Note that by the constraint $s \leq m$, the latter is always satisfied for $m \geq d$.

In the following, we are mainly interested in the commonly used Gaussian kernels $k_\sigma : X \times X \to \mathbb{R}$ defined by

$$k_\sigma(x,x') := \exp\bigl(-\|x-x'\|_2^2/\sigma^2\bigr), \qquad x,x' \in X,$$

where $X \subset \mathbb{R}^d$ is a nonempty subset and $\sigma > 0$ is a free parameter called the width. We write H_σ for the corresponding RKHSs, which are described in some detail in [39]. The entropy numbers for Gaussian kernels [36, Theorem 6.27] and the equivalence of covering and entropy numbers [36, Lemma 6.21] yield that

$$\ln \mathcal{N}(B_{H_\sigma}, \|\cdot\|_\infty, \varepsilon) \leq a\sigma^{-d}\varepsilon^{-2p}, \qquad \varepsilon > 0, \tag{27}$$

for some constants $a > 0$ and $p \in (0,1)$. Then (16) implies $\mathcal{F}_r \subset r^{1/2}\lambda^{-1/2}B_{H_\sigma}$ and consequently

$$\ln \mathcal{N}(\mathcal{F}_r, \|\cdot\|_\infty, \varepsilon) \leq a\sigma^{-d}\lambda^{-p}\varepsilon^{-2p}r^p.$$

Therefore, we can define the function φ in Theorem 1 as

$$\varphi(\varepsilon) := a\sigma^{-d}\lambda^{-p}\varepsilon^{-2p}. \tag{28}$$

Moreover, [15, Section 2] shows that there exists a constant $c > 0$ such that for all $\lambda > 0$ and all $\sigma \in (0,1]$, there is an $f_0 \in H_\sigma$ with $\|f_0\|_\infty \leq c$ and

$$A(\lambda) \leq \lambda\|f_0\|_{H_\sigma}^2 + \mathcal{R}_{L,P}(f_0) - \mathcal{R}_{L,P}^* \leq c\lambda\sigma^{-d} + c\sigma^{2t}.$$

Example 4 (Gaussian Kernels). *Let $Y := [-M,M]$ for some $M > 0$, $Z := \mathbb{R}^d \times Y$, $\mathcal{Z}$ be a stochastic process satisfying Assumption 2, and P be a distribution on Z whose marginal distribution on $\mathbb{R}^d$ is concentrated on $X \subset B_{\ell_2^d}$ and absolutely continuous w.r.t. the Lebesgue measure μ on $\mathbb{R}^d$. We denote the corresponding density $g : \mathbb{R}^d \to [0,\infty)$ and assume $\mu(\partial X) = 0$ and $g \in L_q(\mu)$ for some $q \geq 1$. Moreover, assume that the Bayes decision function $f_{L,P}^* = \mathbb{E}_P(Y|x)$ satisfies $f_{L,P}^* \in L_2(\mu) \cap L_\infty(\mu)$ as well as $f_{L,P}^* \in B_{2s,\infty}^t$ for some $t \geq 1$ and $s \geq 1$ with $\frac{1}{q} + \frac{1}{s} = 1$. Here, $B_{2s,\infty}^t$ denotes the Besov space with the smoothness parameter t, see also [15, Section 2]. Recall that, for the least squares loss, the variance bound (18) is valid with $\vartheta = 1$. Consequently, Condition (19) is satisfied if*

$$r \geq \max\left\{\left(c_V 2^{1+3p}a\right)^{\frac{1}{1-p}} \sigma^{-\frac{d}{1-p}} \lambda^{-\frac{p}{1-p}} n_{\mathit{eff}}^{-\frac{1}{1-p}} \varepsilon^{-\frac{2p}{1-p}}, \frac{2c_V\tau}{n_{\mathit{eff}}}, \frac{8c_B B_0 \tau}{n_{\mathit{eff}}}, r^*\right\}. \tag{29}$$

Note that in the right-hand side of (29), the first term dominates when n goes large. In this context, the oracle inequality (20) becomes

$$\lambda\|f_{D_n,\lambda}\|_{H_\sigma}^2 + \mathcal{R}_{L,P}(\widehat{f}_{D_n,\lambda}) - \mathcal{R}_{L,P}^* \leq C\Bigl(\lambda\sigma^{-d} + \sigma^{2t} + \sigma^{-\frac{d}{1-p}}\lambda^{-\frac{p}{1-p}} n_{\mathit{eff}}^{-\frac{1}{1-p}} \varepsilon^{-\frac{2p}{1-p}}\tau + \varepsilon\Bigr).$$

Here C is a constant independent of n, λ, σ, τ, or ε. Again, optimizing over ε together with some standard techniques, see [36, Lemma A.1.7], we then see that for all $\xi > 0$, the LS-SVM using Gaussian RKHS H_σ and

$$\lambda_n = n_{\mathit{eff}}^{-1} \qquad \text{and} \qquad \sigma_n = n_{\mathit{eff}}^{-\frac{1}{2t+d}},$$

learns with the rate

$$n_{\mathit{eff}}^{-\frac{2t}{2t+d}+\xi}. \tag{30}$$

In the i.i.d. case we have $n_{\mathrm{eff}} = n$, and hence the learning rate (30) becomes

$$n^{-\frac{2t}{2t+d}+\xi}\,. \tag{31}$$

Recall that modulo the arbitrarily small $\xi > 0$ these learning rates are essentially optimal, see e.g. [41, Theorem 13] or [18, Theorem 3.2]. Moreover, for geometrically α-mixing processes, the rate (30) becomes

$$n^{-\frac{2t}{2t+d}\alpha+\xi}\,,$$

where $\alpha := \frac{\gamma}{\gamma+1}$. This rate is optimal up to the factor α and the additional ξ in the exponent. Particularly, for restricted geometrically α-mixing processes, geometrically α-mixing Markov chains, ϕ-mixing processes, we obtain the essentially optimal learning rates (31). Moreover, the same essentially optimal learning rates can be achieved for (time-reversed) geometrically $\mathcal{C}$-mixing processes, if we additionally assume $f^*_{L,P} \in \mathrm{Lip}(\mathbb{R}^d)$, see also [21, Example 4.7].

In the last example, we will briefly discuss learning rates for SVMs for quantile regression. For more information on such SVMs we refer to [15, Section 4].

Example 5 (Quantile Regression with Gaussian Kernels)**.** *Let $Y := [-1,1]$, $Z := \mathbb{R}^d \times Y$, $\mathcal{Z}$ be a stochastic process satisfying Assumption 2, P be a distribution on Z, and Q be the marginal distribution of P on $\mathbb{R}^d$. Assume that $X := \operatorname{supp} Q \subset B_{\ell_2^d}$ and that for Q-almost all $x \in X$, the conditional probability $P(\cdot|x)$ is absolutely continuous w.r.t. the Lebesgue measure on Y and the conditional densities $h(\cdot, x)$ of $P(\cdot|x)$ are uniformly bounded away from 0 and ∞, see also [15, Example 4.5]. Moreover, assume that Q is absolutely continuous w.r.t. the Lebesgue measure on X with associated density $g \in L_u(X)$ for some $u \geq 1$. For $\tau \in (0,1)$, let $f^*_{\tau,P} : \mathbb{R}^d \to \mathbb{R}$ be a conditional τ-quantile function that satisfies $f^*_{\tau,P} \in L_2(\mu) \cap L_\infty(\mu)$. In addition, we assume that $f^*_{\tau,P} \in B^t_{2s,\infty}$ for some $t \geq 1$ and $s \geq 1$ such that $\frac{1}{s} + \frac{1}{u} = 1$. Then [38, Theorem 2.8] yields a variance bound of the form*

$$\mathbb{E}_P(L_\tau \circ \widehat{f} - L_\tau \circ f^*_{\tau,P})^2 \leq V \cdot \mathbb{E}_P(L_\tau \circ \widehat{f} - L_\tau \circ f^*_{\tau,P})\,, \tag{32}$$

for all $f : X \to \mathbb{R}$, where V is a suitable constant and L_τ is the τ-pinball loss. Then, following similar arguments with those in Example 4, with the same choices of λ_n and σ_n, the same rates can be obtained as in Example 4.

Here, we give two remarks on Example 5. First, it is noted that the Bernstein condition (32) holds when the distribution P is of a τ-quantile of p-average type q in the sense of Definition 2.6 in [38]. Two distributions of this type can be found in Examples 2.3 and 2.4 in [35]. On the other hand, the rates obtained in Example 5 are in fact for the excess L_τ-risk. However, since [38, Theorem 2.7] shows

$$\|\widehat{f} - f^*_{\tau,P}\|^2_{L_2(P_X)} \leq c\big(\mathcal{R}_{L_\tau,P}(\widehat{f}\,) - \mathcal{R}^*_{L_\tau,P}\big)$$

for some constant $c > 0$ and all $f : X \to \mathbb{R}$, we also obtain the same rates for $\|\widehat{f} - f^*_{\tau,P}\|^2_{L_2(P_X)}$. Last but not least, optimality for various mixing processes can be discussed along the lines of LS-SVMs.

6 Conclusions

In the present paper, we proposed a unified learning theory approach to studying learning schemes sampling from various commonly investigated stationary mixing processes that include geometrically α-mixing processes, geometrically α-mixing Markov chains, ϕ-mixing processes, and geometrically $\mathcal{C}$-mixing processes. The proposed approach is considered to be unified in the following sense: First, in our study, the empirical processes of the above-mentioned mixing processes were assumed to satisfy a generalized Bernstein-type inequality, which includes many commonly considered cases; Second, by instantiating the generalized Bernstein-type inequality to different scenarios, we illustrated the effective number of observations for different mixing processes; Third, based on the above generalized Bernstein-type concentration assumption, a generalized sharp oracle inequality was established within the statistical learning theory framework. Finally, faster or at least comparable learning rates can be obtained by applying the established oracle inequality to various learning schemes with different mixing processes.

Acknowledgement

The authors would like to thank the editor and the reviewers for their insightful comments and helpful suggestions that improved the quality of this paper. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC AdG A-DATADRIVE-B (290923). This paper reflects only the authors' views, the Union is not liable for any use that may be made of the contained information; Research Council KUL: GOA/10/09 MaNet, CoE PFV/10/002 (OPTEC), BIL12/11T; PhD/Postdoc grants; Flemish Government: FWO: PhD/Postdoc grants, projects: G.0377.12 (Structured systems), G.088114N (Tensor based data similarity); IWT: PhD/Postdoc grants, projects: SBO POM (100031); iMinds Medical Information Technologies SBO 2014; Belgian Federal Science Policy Office: IUAP P7/19 (DYSCO, Dynamical systems, control and optimization, 2012-2017).

Appendix

Proof of Theorem 1 in Section 4

Since the proof of Theorem 1 is rather complicated, we first describe its main steps briefly: First we decompose the regularized excess risk into an approximation error term and two stochastic error terms. The approximation error and the first stochastic error term can be estimated by standard techniques. Similarly, the first step in the estimation of the second error term is a rather standard quotient approach, see e.g. [36, Theorem 7.20], which allows for localization with respect to both the variance and the regularization. Due to the absence of tools from empirical process theory, however, the remaining estimation steps become more involved. To be more precise, we split the "unit ball" of the hypothesis space $\mathcal{F}$ into disjoint "spheres". For each sphere, we then use localized covering numbers and the generalized Bernstein-type inequality from Assumption 2, and the resulting estimates are then combined using the peeling method. This yields a quasi-geometric series with rate smaller than 1 if the radius of the innermost ball is sufficiently large. As a result, the estimated error probability on the whole "unit ball" nearly equals the estimated error probability of the innermost "ball", which unsurprisingly leads to a significant improvement compared to [37].

Before we prove Theorem 1, we need to reformulate (7). Setting $\tau := \frac{\varepsilon^2 n_{\mathit{eff}}}{c_\sigma \sigma^2 + \varepsilon c_B B}$, with some simple

transformations we obtain

$$\mu\left(\left\{\omega\in\Omega : \frac{1}{n}\sum_{i=1}^{n} h(Z_i(\omega)) \geq \sqrt{\frac{\tau c_\sigma \sigma^2}{n_{\mathit{eff}}}} + \frac{c_B B \tau}{n_{\mathit{eff}}}\right\}\right) \leq Ce^{-\tau} \tag{33}$$

for all $\tau > 0$ and $n \geq n_0$.

Proof of Theorem 1. **Main Decomposition.** For $f : X \to \mathbb{R}$ we define $h_f := L \circ f - L \circ f^*_{L,P}$. By the definition of $f_{D_n,\Upsilon}$, we then have

$$\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}} \leq \Upsilon(f_0) + \mathbb{E}_{D_n} h_{f_0} + \delta,$$

and consequently we obtain

$$\begin{aligned}
\Upsilon(f_{D_n,\Upsilon}) + \mathcal{R}_{L,P}(\widehat{f}_{D_n,\Upsilon}) - \mathcal{R}^*_{L,P} &= \Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \\
&\leq \Upsilon(f_0) + \mathbb{E}_{D_n} h_{f_0} - \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}} + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} + \delta \\
&= (\Upsilon(f_0) + \mathbb{E}_P h_{f_0}) + (\mathbb{E}_{D_n} h_{f_0} - \mathbb{E}_P h_{f_0}) + (\mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} - \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}}) + \delta.
\end{aligned} \tag{34}$$

Estimating the First Stochastic Term. Let us first bound the term $\mathbb{E}_{D_n} h_{f_0} - \mathbb{E}_P h_{f_0}$. To this end, we further split this difference into

$$\mathbb{E}_{D_n} h_{f_0} - \mathbb{E}_P h_{f_0} = \left(\mathbb{E}_{D_n}(h_{f_0} - h_{\widehat{f}_0}) - \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})\right) + (\mathbb{E}_{D_n} h_{\widehat{f}_0} - \mathbb{E}_P h_{\widehat{f}_0}). \tag{35}$$

Now $L \circ f_0 - L \circ \widehat{f}_0 \geq 0$ implies $h_{f_0} - h_{\widehat{f}_0} = L \circ f_0 - L \circ \widehat{f}_0 \in [0, B_0]$, and hence we obtain

$$\mathbb{E}_P\left((h_{f_0} - h_{\widehat{f}_0}) - \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})\right)^2 \leq \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})^2 \leq B_0 \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0}).$$

Inequality (33) applied to $h := (h_{f_0} - h_{\widehat{f}_0}) - \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})$ thus shows that

$$\mathbb{E}_{D_n}(h_{f_0} - h_{\widehat{f}_0}) - \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0}) \leq \sqrt{\frac{\tau c_\sigma B_0 \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})}{n_{\mathit{eff}}}} + \frac{c_B B_0 \tau}{n_{\mathit{eff}}}$$

holds with probability μ not less than $1 - Ce^{-\tau}$. Moreover, using $\sqrt{ab} \leq \frac{a}{2} + \frac{b}{2}$, we find

$$\sqrt{n_{\mathit{eff}}^{-1} \tau c_\sigma B_0 \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0})} \leq \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0}) + n_{\mathit{eff}}^{-1} c_\sigma B_0 \tau/4,$$

and consequently we have with probability μ not less than $1 - Ce^{-\tau}$ that

$$\mathbb{E}_{D_n}(h_{f_0} - h_{\widehat{f}_0}) - \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0}) \leq \mathbb{E}_P(h_{f_0} - h_{\widehat{f}_0}) + \frac{7 c_B B_0 \tau}{4 n_{\mathit{eff}}}. \tag{36}$$

In order to bound the remaining term in (35), that is $\mathbb{E}_{D_n} h_{\widehat{f}_0} - \mathbb{E}_P h_{\widehat{f}_0}$, we first observe that (1) implies $\|h_{\widehat{f}_0}\|_\infty \leq 1$, and hence we have $\|h_{\widehat{f}_0} - \mathbb{E}_P h_{\widehat{f}_0}\|_\infty \leq 2$. Moreover, (18) yields

$$\mathbb{E}_P(h_{\widehat{f}_0} - \mathbb{E}_P h_{\widehat{f}_0})^2 \leq \mathbb{E}_P h_{\widehat{f}_0}^2 \leq V(\mathbb{E}_P h_{\widehat{f}_0})^\vartheta.$$

In addition, if $\vartheta \in (0, 1]$, the first inequality in [36, Lemma 7.1] implies for $q := \frac{2}{2-\vartheta}$, $q' := \frac{2}{\vartheta}$, $a := (n_{\mathit{eff}}^{-1} c_\sigma 2^{-\vartheta} \vartheta^\vartheta V \tau)^{1/2}$, and $b := (2\vartheta^{-1} \mathbb{E}_P h_{\widehat{f}_0})^{\vartheta/2}$, that

$$\sqrt{\frac{c_\sigma V \tau (\mathbb{E}_P h_{\widehat{f}_0})^\vartheta}{n_{\mathit{eff}}}} \leq \left(1 - \frac{\vartheta}{2}\right)\left(\frac{c_\sigma 2^{-\vartheta} \vartheta^\vartheta V \tau}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \mathbb{E}_P h_{\widehat{f}_0} \leq \left(\frac{c_\sigma V \tau}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \mathbb{E}_P h_{\widehat{f}_0}.$$

Since $\mathbb{E}_P h_{\widehat{f}_0} \geq 0$, this inequality also holds for $\vartheta = 0$, and hence (33) shows that we have

$$\mathbb{E}_{D_n} h_{\widehat{f}_0} - \mathbb{E}_P h_{\widehat{f}_0} \leq \mathbb{E}_P h_{\widehat{f}_0} + \left(\frac{c_\sigma V \tau}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \frac{2c_B \tau}{n_{\mathit{eff}}}$$

with probability μ not less than $1 - Ce^{-\tau}$. By combining this estimate with (36) and (35), we now obtain that with probability μ not less than $1 - 2Ce^{-\tau}$ we have

$$\mathbb{E}_{D_n} h_{f_0} - \mathbb{E}_P h_{f_0} \leq \mathbb{E}_P h_{f_0} + \left(\frac{c_\sigma V \tau}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \frac{2c_B \tau}{n_{\mathit{eff}}} + \frac{7c_B B_0 \tau}{4n_{\mathit{eff}}}, \tag{37}$$

since $1 \leq B_0$, i.e., we have established a bound on the second term in (34).

Estimating the Second Stochastic Term. For the third term in (34) let us first consider the case $n_{\mathit{eff}} < c_V(\tau + \varphi(\varepsilon/2)2^p r^p)$. Combining (37) with (34) and using $1 \leq B_0$, $1 \leq V$, $c_\sigma V \leq c_V$, $2 \leq 4^{1/(2-\vartheta)}$, and $\mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} - \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}} \leq 2$, then we find

$$\begin{aligned}
\Upsilon(f_{D_n,\Upsilon}) + \mathcal{R}_{L,P}(\widehat{f}_{D_n,\Upsilon}) - \mathcal{R}^*_{L,P} &\leq \Upsilon(f_0) + 2\mathbb{E}_P h_{f_0} + \left(\frac{c_\sigma V \tau}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \frac{2c_B \tau}{n_{\mathit{eff}}} + \frac{7c_B B_0 \tau}{4n_{\mathit{eff}}} \\
&\quad + (\mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} - \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}}) + \delta \\
&\leq \Upsilon(f_0) + 2\mathbb{E}_P h_{f_0} + \left(\frac{c_\sigma V(\tau + \varphi(\varepsilon/2)2^p r^p)}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \frac{4c_B B_0 \tau}{n_{\mathit{eff}}} \\
&\quad + 2\left(\frac{c_V(\tau + \varphi(\varepsilon/2)2^p r^p)}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \delta \\
&\leq 2\Upsilon(f_0) + 4\mathbb{E}_P h_{f_0} + 3\left(\frac{c_V(\tau + \varphi(\varepsilon/2)2^p r^p)}{n_{\mathit{eff}}}\right)^{\frac{1}{2-\vartheta}} + \frac{8c_B B_0 \tau}{n_{\mathit{eff}}} + 2\delta
\end{aligned}$$

with probability μ not less than $1 - 2Ce^{-\tau}$. It thus remains to consider the case $n_{\mathit{eff}} \geq c_V(\tau + \varphi(\varepsilon/2)2^p r^p)$.

Introduction of the Quotients. To establish a non-trivial bound on the term $\mathbb{E}_P h_{\widehat{f}_D} - \mathbb{E}_{D_n} h_{\widehat{f}_D}$ in (34), we define functions

$$g_{f,r} := \frac{\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}}{\Upsilon(f) + \mathbb{E}_P h_{\widehat{f}} + r}, \qquad f \in \mathcal{F},\ r > r^*.$$

For $f \in \mathcal{F}$, we have $\|\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}\|_\infty \leq 2$. Moreover, for $f \in \mathcal{F}_r$, the variance bound (18) implies

$$\mathbb{E}_P(h_{\widehat{f}} - \mathbb{E}_P h_{\widehat{f}})^2 \leq \mathbb{E}_P h^2_{\widehat{f}} \leq V(\mathbb{E}_P h_{\widehat{f}})^\vartheta \leq V r^\vartheta. \tag{38}$$

Peeling. For a fixed $r \in (r^*, 1]$, let K be the largest integer satisfying $2^K r \leq 1$. Then we can get the following disjoint partition of the function set $\mathcal{F}_1$:

$$\mathcal{F}_1 \subset \mathcal{F}_r \cup \bigcup_{k=1}^{K+1} (\mathcal{F}_{2^k r} \setminus \mathcal{F}_{2^{k-1} r}).$$

We further write $\overline{C}_{\varepsilon,r,0}$ for a minimal ε-net of $\mathcal{F}_r$ and $\overline{C}_{\varepsilon,r,k}$ for minimal ε-nets of $\mathcal{F}_{2^k r} \setminus \mathcal{F}_{2^{k-1} r}$, $1 \leq k \leq K+1$, respectively. Then the union of these nets $\bigcup_{k=0}^{K+1} \overline{C}_{\varepsilon,r,k} =: \overline{C}_{\varepsilon,1}$ is an ε-net of the set $\mathcal{F}_1$. Moreover, we define

$$\widetilde{\mathcal{C}}_{\varepsilon,r,k} := \bigcup_{l=0}^{k} \overline{C}_{\varepsilon,r,l}, \qquad 0 \leq k \leq K+1,$$

which are ε-nets of $\mathcal{F}_{2^k r}$ with $\widetilde{\mathcal{C}}_{\varepsilon,r,k} \subset \widetilde{\mathcal{C}}_{\varepsilon,r,k+1}$ for all $0 \le k \le K$, and the net $\widetilde{\mathcal{C}}_{\varepsilon,r,K+1}$ coincide with $\overline{C}_{\varepsilon,1}$. For $A \subset B$ an elementary calculation shows that

$$\mathcal{N}(A, \|\cdot\|_\infty, \varepsilon) \le \mathcal{N}(B, \|\cdot\|_\infty, \varepsilon/2). \tag{39}$$

By using (39) for $\mathcal{F}_{2^k r} \backslash \mathcal{F}_{2^{k-1} r} \subset \mathcal{F}_{2^k r}$ we can estimate the cardinality of $\widetilde{\mathcal{C}}_{\varepsilon,r,k}$ by

$$\begin{aligned} |\widetilde{\mathcal{C}}_{\varepsilon,r,k}| &= \left| \bigcup_{l=0}^{k} \overline{C}_{\varepsilon,r,l} \right| \le \sum_{l=0}^{k} |\overline{C}_{\varepsilon,r,l}| = \sum_{l=0}^{k} \mathcal{N}(\mathcal{F}_{2^k r} \backslash \mathcal{F}_{2^{k-1} r}, \|\cdot\|_\infty, \varepsilon) \\ &\le \sum_{l=0}^{k} \mathcal{N}(\mathcal{F}_{2^k r}, \|\cdot\|_\infty, \varepsilon/2) \le \sum_{l=0}^{k} \exp\left(\varphi(\varepsilon/2)(2^l r)^p\right) \\ &\le (k+1) \exp\left(\varphi(\varepsilon/2) 2^{kp} r^p\right), \qquad 0 \le k \le K+1. \end{aligned} \tag{40}$$

Using the peeling technique in [20, Theorem 5.2] with $\mathcal{Z}_f := \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}})$, $\Gamma(f) := \Upsilon(f) + \mathbb{E}_P h_{\widehat{f}}$,

$$m_k := \begin{cases} r^* & \text{for } k = 0, \\ 2^{k-1} r & \text{for } 1 \le k \le K, \\ 1 & \text{for } k = K+1, \end{cases}$$

and choosing $\epsilon = 1/4$, we get

$$\begin{aligned} \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,1}} \mathbb{E}_{D_n} g_{f,r} > \frac{1}{4} \Big) &= \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,1}} \frac{\mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}})}{\Upsilon(f) + \mathbb{E}_P h_{\widehat{f}} + r} > \frac{1}{4} \Big) \\ &\le \sum_{k=1}^{K+2} \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,r,k}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > (2^{k-1} r + r)/4 \Big) \\ &\le \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,r,0}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > (r^* + r)/4 \Big) \\ &\quad + \sum_{k=1}^{K+1} \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,r,k}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > (2^{k-1} r + r)/4 \Big) \\ &\le \mu\Big(\sup_{f \in \widetilde{\mathcal{C}}_{\varepsilon,r,1}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > r/4 \Big) \\ &\quad + \sum_{k=1}^{K+1} \mu\Big(\sup_{f \in \widetilde{\mathcal{C}}_{\varepsilon,r,k}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > 2^{k-1} r/4 \Big) \\ &\le 2 \sum_{k=1}^{K+1} \mu\Big(\sup_{f \in \widetilde{\mathcal{C}}_{\varepsilon,r,k}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > 2^{k-3} r \Big). \end{aligned}$$

Estimating the Error Probabilities on the "Spheres". Our next goal is to estimate all the error probabilities by using (7), (38) and the union bound. From $\vartheta \in [0,1]$ and the estimations of the covering numbers (40) follows that

$$\begin{aligned} &\mu\Big(\sup_{f \in \widetilde{\mathcal{C}}_{\varepsilon,r,k}} \mathbb{E}_{D_n}(\mathbb{E}_P h_{\widehat{f}} - h_{\widehat{f}}) > 2^{k-3} r \Big) \\ &\le C |\widetilde{\mathcal{C}}_{\varepsilon,r,k}| \exp\Big(- \frac{(2^{k-3} r)^2 n_{\mathit{eff}}}{c_\sigma V (2^k r)^\vartheta + 2 c_B (2^{k-3} r)} \Big) \\ &\le C \cdot (k+1) \exp\left(\varphi(\varepsilon/2) 2^{kp} r^p\right) \cdot \exp\Big(- \frac{(2^{k-1} r)^2 n_{\mathit{eff}}}{32 c_\sigma V (2^{k-1} r)^\vartheta + 8 c_B (2^{k-1} r)} \Big). \end{aligned}$$

For $k \geq 1$, we denote the right-hand side of this estimate by $p_k(r)$, that is

$$p_k(r) := C \cdot (k+1) \exp\left(\varphi(\varepsilon/2) 2^{kp} r^p\right) \cdot \exp\Big(-\frac{(2^{k-1}r)^2 n_{\mathit{eff}}}{32 c_\sigma V (2^{k-1}r)^\vartheta + 8 c_B (2^{k-1}r)}\Big).$$

Then we have

$$\begin{aligned} q_k(r) := \frac{p_{k+1}(r)}{p_k(r)} &\leq \frac{k+2}{k+1} \cdot \exp\left(\varphi(\varepsilon/2)(2^{k+1}r)^p - \varphi(\varepsilon/2)(2^k r)^p\right) \\ &\quad \cdot \exp\Big(-\frac{2^2(2^{k-1}r)^2 n_{\mathit{eff}}}{32 c_\sigma V \cdot 2(2^{k-1}r)^\vartheta + 8 c_B \cdot 2(2^{k-1}r)} + \frac{(2^{k-1}r)^2 n_{\mathit{eff}}}{32 c_\sigma V (2^{k-1}r)^\vartheta + 8 c_B (2^{k-1}r)}\Big) \\ &\leq 2 \exp\left(\varphi(\varepsilon/2) 2^{kp+1} r^p\right) \cdot \exp\Big(-\frac{(2^{k-1}r)^2 n_{\mathit{eff}}}{32 c_\sigma V (2^{k-1}r)^\vartheta + 8 c_B (2^{k-1}r)}\Big), \end{aligned}$$

and our assumption $2^k r \leq 1$, $0 \leq k \leq K$ implies

$$\begin{aligned} q_k(r) &\leq 2 \exp\left(\varphi(\varepsilon/2) 2^{kp+1} r^p\right) \cdot \exp\Big(-\frac{(2^{k-1}r)^2 n_{\mathit{eff}}}{32 c_\sigma V (2^{k-1}r)^\vartheta + 8 c_B (2^{k-1}r)}\Big) \\ &\leq 2 \exp\Big(2^{(k-1)p} \cdot 4 r^p \varphi(\varepsilon/2) - 2^{(k-1)(2-\vartheta)} \cdot \frac{r^{2-\vartheta} n_{\mathit{eff}}}{32 c_\sigma V + 8 c_B}\Big). \end{aligned}$$

Since $p \in (0,1]$, $k \geq 1$ and $\vartheta, \in [0,1]$, we have

$$2^{(k-1)p} \leq 2^{(k-1)(2-\vartheta)}.$$

The first assumption in (19) implies that $r \geq \big(64(4c_\sigma V + c_B)\varphi(\varepsilon/2) r^p / n_{\mathit{eff}}\big)^{1/(2-\vartheta)}$ or equivalently that

$$4 r^p \varphi(\varepsilon/2) \leq \frac{1}{2} \cdot \frac{r^{2-\vartheta} n_{\mathit{eff}}}{32 c_\sigma V + 8 c_B},$$

thus, using $2^{(k-1)(2-\vartheta)} \geq 1$, we find

$$q_k(r) \leq 2 \exp\Big(-\frac{1}{2} \cdot \frac{r^{2-\vartheta} n_{\mathit{eff}}}{32 c_\sigma V + 8 c_B}\Big).$$

Moreover, since $\tau \geq 1$, the first assumption in (19) implies also $r \geq \big(64(4c_\sigma V + c_B)/n_{\mathit{eff}}\big)^{1/(2-\vartheta)}$ or equivalently that

$$\frac{1}{2} \cdot \frac{r^{2-\vartheta} n_{\mathit{eff}}}{32 c_\sigma V + 8 c_B} \geq 4,$$

and hence $q_k(r) \leq 2e^{-4}$, that is, $p_{k+1}(r) \leq 2e^{-4} p_k(r)$ for all $k \geq 1$.

Summing all the Error Probabilities. From the above discussion we have

$$\begin{aligned} \mu\Big(\sup_{f \in \overline{\mathcal{C}}_{\varepsilon,1}} \mathbb{E}_{D_n} g_{f,r} > \frac{1}{4}\Big) &\leq 2 \sum_{k=1}^{K+1} p_k(r) \leq 2 \cdot p_1(r) \cdot \sum_{k=0}^{K} (2e^{-4})^k \leq 3 p_1(r) \\ &= 6\, C \exp\left(\varphi(\varepsilon/2) 2^p r^p\right) \cdot \exp\Big(-\frac{r^2 n_{\mathit{eff}}}{32 c_\sigma V r^\vartheta + 8 c_B r}\Big) \\ &\leq 6\, C \exp\left(\varphi(\varepsilon/2) 2^p r^p\right) \cdot \exp\Big(-\frac{r^2 n_{\mathit{eff}}}{32 c_\sigma V r^\vartheta + 8 c_B r^\vartheta}\Big) \\ &\leq 6\, C \exp\left(\varphi(\varepsilon/2) 2^p r^p\right) \cdot \exp\left(-\frac{r^{2-\vartheta} n_{\mathit{eff}}}{32 c_\sigma V + 8 c_B}\right). \end{aligned}$$

Then once again the first assumption in (19) gives

$$r \geq \left(\frac{(32c_\sigma V + 8c_B)(\tau + \varphi(\varepsilon/2)2^p r^p)}{n_{\mathit{eff}}} \right)^{\frac{1}{2-\vartheta}}$$

and a simple transformation thus yields

$$\mu \left(D_n \in (X \times Y)^n : \sup_{f \in \overline{\mathcal{C}}_{\varepsilon,1}} \mathbb{E}_{D_n} g_{f,r} \leq \frac{1}{4} \right) \geq 1 - 6Ce^{-\tau}.$$

Consequently we see that with probability μ not less than $1 - 6Ce^{-\tau}$ we have

$$\mathbb{E}_P h_{\widehat{f}} - \mathbb{E}_{D_n} h_{\widehat{f}} \leq \frac{1}{4} \left(\Upsilon(f) + \mathbb{E}_P h_{\widehat{f}} + r \right) \tag{41}$$

for all $f \in \overline{\mathcal{C}}_{\varepsilon,1}$. Since $r \in (0,1]$, we have $f_{D_n,\Upsilon} \in \mathcal{F}_1$, i.e. either $f_{D_n,\Upsilon} \in \mathcal{F}_r$, or there exists an integer $k \leq K+1$ such that $f_{D_n,\Upsilon} \in \mathcal{F}_{2^k r} \backslash \mathcal{F}_{2^{k-1} r}$. Thus there exists an $f_{D_n} \in \overline{\mathcal{C}}_{\varepsilon,r,0} \subset \mathcal{F}_r$ or $f_{D_n} \in \overline{\mathcal{C}}_{\varepsilon,r,k} \subset \mathcal{F}_{2^k r} \backslash \mathcal{F}_{2^{k-1} r}$ with $\|f_{D_n,\Upsilon} - f_{D_n}\|_\infty \leq \varepsilon$. By the assumed Lipschitz continuity of the clipped L the latter implies

$$|h_{\widehat{f}_{D_n}}(x,y) - h_{\widehat{f}_{D_n,\Upsilon}}(x,y)| \leq |\widehat{f}_{D_n}(x) - \widehat{f}_{D_n,\Upsilon}(x)| \leq |f_{D_n}(x) - f_{D_n,\Upsilon}(x)| \leq \varepsilon \tag{42}$$

for all $(x,y) \in X \times Y$. For $f_{D_n,\Upsilon}, f_{D_n} \in \mathcal{F}_r$ we obviously have

$$\Upsilon(f_{D_n}) + \mathbb{E}_P h_{\widehat{f}_{D_n}} \leq r$$

and for the other cases $f_{D_n,\Upsilon}, f_{D_n} \in \mathcal{F}_{2^k r} \backslash \mathcal{F}_{2^{k-1} r}$ we obtain

$$\Upsilon(f_{D_n}) + \mathbb{E}_P h_{\widehat{f}_{D_n}} \leq 2^k r = 2 \cdot 2^{k-1} r \leq 2 \left(\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \right),$$

consequently, we always have

$$\Upsilon(f_{D_n}) + \mathbb{E}_P h_{\widehat{f}_{D_n}} \leq 2 \left(\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \right) + r. \tag{43}$$

Combining (42) with (41) and (43), we obtain

$$\mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} - \mathbb{E}_{D_n} h_{\widehat{f}_{D_n,\Upsilon}} \leq \frac{1}{2} \left(\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} + \varepsilon + r \right) + 2\varepsilon$$

with probability μ not less than $1 - 6Ce^{-\tau}$. By combining this estimate with (34) and (37), we then obtain that

$$\begin{aligned} \Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \leq \Upsilon(f_0) + 2\mathbb{E}_P h_{f_0} + \left(\frac{c_\sigma V \tau}{n_{\mathit{eff}}} \right)^{\frac{1}{2-\vartheta}} + \frac{2c_B \tau}{n_{\mathit{eff}}} + \frac{7c_B B_0 \tau}{4n_{\mathit{eff}}} + \delta \\ + \frac{\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}}}{2} + \frac{5}{2}\varepsilon + \frac{1}{2} r \end{aligned}$$

holds with probability μ not less than $1 - 8Ce^{-\tau}$. From the assumptions in (19) follows that

$$\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \leq \Upsilon(f_0) + 2\mathbb{E}_P h_{f_0} + 2r + \delta + \frac{\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}}}{2} + \frac{5\varepsilon}{2}$$

holds with probability μ not less than $1 - 8Ce^{-\tau}$. Consequently, we have

$$\Upsilon(f_{D_n,\Upsilon}) + \mathbb{E}_P h_{\widehat{f}_{D_n,\Upsilon}} \leq 2\Upsilon(f_0) + 4\mathbb{E}_P h_{f_0} + 4r + 5\varepsilon + 2\delta.$$

Therefore, we have proved the assertion. □

References

[1] R. Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electron. J. Probab.*, 13:1000–1034, 2008.

[2] P. Alquier and O. Wintenberger. Fast rates in learning with dependent observations. *arXiv preprint arXiv:1202.4283*, 2012.

[3] G. V. Bayley and J. M. Hammersley. The "effective" number of independent observations in an autocorrelated time series. *Supplement to J. R. Stat. Soc.*, 8(2):184–197, 1946.

[4] D. Belomestny. Spectral estimation of the Lévy density in partially observed affine models. *Stochastic Process. Appl.*, 121(6):1217–1244, 2011.

[5] S. Bernstein. *The Theory of Probabilities*. Gastehizdat Publishing House, Moscow, 1946.

[6] G. Blanchard, G. Lugosi, and N. Vayatis. On the rate of convergence of regularized boosting classifiers. *J. Mach. Learn. Res.*, 4(5):861–894, 2004.

[7] D. Bosq. Bernstein-type large deviations inequalities for partial sums of strong mixing processes. *Statistics*, 24(1):59–70, 1993.

[8] D. Bosq. *Nonparametric Statistics for Stochastic Processes: Estimation and Prediction*, volume 110. Springer Science & Business Media, 2012.

[9] R. C. Bradley. Basic properties of strong mixing conditions. A survey and some open questions. *Probab. Surv.*, 2:107–144, 2005. Update of, and a supplement to, the 1986 original.

[10] R. C. Bradley. *Introduction to Strong Mixing Conditions. Vol. 1-3*. Kendrick Press, Heber City, UT, 2007.

[11] Y. A. Davydov. Convergence of distributions generated by stationary stochastic processes. *Theor. Probab. Appl.*, 13(4):691–696, 1968.

[12] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1996.

[13] L. Devroye and G. Lugosi. *Combinatorial Methods in Density Estimation*. Springer, New York, 2001.

[14] J. C. Duchi, A. Agarwal, M. Johansson, and M. I. Jordan. Ergodic mirror descent. *SIAM J. Optimiz.*, 22(4):1549–1578, 2012.

[15] M. Eberts and I. Steinwart. Optimal regression rates for SVMs using Gaussian kernels. *Electron. J. Stat.*, 7:1–42, 2013.

[16] J. Fan and Q. Yao. *Nonlinear Time Series: Nonparametric and Parametric Methods*. Springer-Verlag, New York, 2003.

[17] Y. Feng. Least-squares regularized regression with dependent samples and q-penalty. *Appl. Anal.*, 91(5):979–991, 2012.

[18] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-free Theory of Nonparametric Regression*. Springer-Verlag, New York, 2002.

[19] H. Hang, Y. Feng, and J. A.K. Suykens. Learning from $\mathcal{C}$-mixing processes. *Internal Report, ESAT, KU Leuven, Leuven, Belgium*, 2016.

[20] H. Hang and I. Steinwart. Fast learning from α-mixing observations. *J. Multivariate Anal.*, 127:184–199, 2014.

[21] H. Hang and I. Steinwart. A Bernstein-type inequality for some mixing processes and dynamical systems with an application to learning. *Ann. Stat., in press*, 2016. Available at http://www.e-publications.org/ims/submission/AOS/user/submissionFile/22219?confirm=ec9efb84.

[22] I. A. Ibragimov. Some limit theorems for stationary processes. *Theor. Probab. Appl.*, 7(4):349–382, 1962.

[23] I. A. Ibragimov and Y. A. Rozanov. *Gaussian Random Processes*. Springer-Verlag, New York-Berlin, 1978. Translated from the Russian by A. B. Aries.

[24] D. Lubman. Spatial averaging in a diffuse sound field. *J. Acoust. Soc. Am.*, 46(3A):532–534, 1969.

[25] V. Maume-Deschamps. Exponential inequalities and functional estimations for weak dependent data; applications to dynamical systems. *Stoch. Dyn.*, 6(4):535–560, 2006.

[26] S. Mendelson and J. Neeman. Regularization in kernel learning. *Ann. Stat.*, 38(1):526–565, 2010.

[27] F. Merlevède, M. Peligrad, and E. Rio. Bernstein inequality and moderate deviations under strong mixing conditions. In *High Dimensional Probability V: the Luminy Volume*, volume 5 of *Inst. Math. Stat. Collect.*, pages 273–292. Inst. Math. Statist., Beachwood, OH, 2009.

[28] D. S. Modha and E. Masry. Minimum complexity regression estimation with weakly dependent observations. *IEEE Trans. Inform. Theory*, 42(6, part 2):2133–2145, 1996.

[29] M. Mohri and A. Rostamizadeh. Stability bounds for stationary ϕ-mixing and β-mixing processes. *J. Mach. Learn. Res.*, 4:1–26, 2009.

[30] A. B. Nobel. Limits to classification and regression estimation from ergodic processes. *Ann. Stat.*, 27(1):262–273, 1999.

[31] M. Rosenblatt. A central limit theorem and a strong mixing condition. *Proc. Nat. Acad. Sci. USA.*, 42:43–47, 1956.

[32] P.-M. Samson. Concentration of measure inequalities for Markov chains and Φ-mixing processes. *Ann. Probab.*, 28(1):416–461, 2000.

[33] Z. Şen. Small sample estimation of the variance of time-averages in climatic time series. *Int. J. Climatol.*, 18(15):1725–1732, 1998.

[34] I. Steinwart. Two oracle inequalities for regularized boosting classifiers. *Stat. Interface*, 2(3):271–284, 2009.

[35] I. Steinwart and A. Christmann. How SVMs can estimate quantiles and the median. In *Advances in Neural Information Processing Systems*, pages 305–312, 2007.

[36] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.

[37] I. Steinwart and A. Christmann. Fast learning from non-i.i.d. observations. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22, pages 1768–1776. MIT Press, Cambridge, MA, 2009.

[38] I. Steinwart and A. Christmann. Estimating conditional quantiles with the help of the pinball loss. *Bernoulli*, 17(1):211–225, 2011.

[39] I. Steinwart, D. Hush, and C. Scovel. An explicit description of the reproducing kernel Hilbert spaces of Gaussian RBF kernels. *IEEE Trans. Inform. Theory*, 52(10):4635–4643, 2006.

[40] I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *J. Multivariate Anal.*, 100:175–194, 2009.

[41] I. Steinwart, D. Hush, and C. Scovel. Optimal rates for regularized least squares regression. In S. Dasgupta and A. Klivans, editors, *Proceedings of the 22nd Annual Conference on Learning Theory*, pages 79–93. 2009.

[42] H. Sun and Q. Wu. A note on application of integral operator in learning theory. *Appl. Comput. Harmon. Anal.*, 26(3):416–421, 2009.

[43] H. Sun and Q. Wu. Regularized least square regression with dependent samples. *Adv. Comput. Math.*, 32:175–189, 2010.

[44] J. A. K. Suykens, T. Van Gestel, J. De Brabanter, B. De Moor, and J. Vandewalle. *Least Squares Support Vector Machines*. World Scientific, Singapore, 2002.

[45] A. B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *Ann. Stat.*, 32(1):135–166, 2004.

[46] V. N. Vapnik. *Statistical Learning Theory*. John Wiley & Sons, Inc., New York, 1998.

[47] M. Viana. *Stochastic Dynamics of Deterministic Systems*, volume 21. IMPA Brazil, 1997.

[48] M. Vidyasagar. *A Theory of Learning and Generalization: With Applications to Neural Networks and Control Systems*. Springer, London, second edition, 2003.

[49] V. A. Volkonskii and Y. A. Rozanov. Some limit theorems for random functions. I. *Theor. Probab. Appl.*, 4:178–197, 1959.

[50] V. A. Volkonskii and Y. A. Rozanov. Some limit theorems for random functions. II. *Teor. Verojatnost. i Primenen.*, 6:202–215, 1961.

[51] Y.-L. Xu and D.-R. Chen. Learning rates of regularized regression for exponentially strongly mixing sequence. *J. Statist. Plann. Inference*, 138:2180–2189, 2008.

[52] A. M. Yaglom. *Correlation Theory of Stationary and Related Random Functions*. Springer-Verlag, 1987.

[53] B. Yu. Rates of convergence for empirical processes of stationary mixing sequences. *Ann. Probab.*, 22:94–116, 1994.

[54] T. Zhang. Statistical behavior and consistency of classification methods based on convex risk minimization. *Ann. Stat.*, 32(1):56–85, 2004.

[55] A. Zieba. Effective number of observations and unbiased estimators of variance for autocorrelated data-an overview. *Metrol. Meas. Syst.*, 17(1):3–16, 2010.