
A Case of Exponential Convergence Rates for SVM

Vivien Cabannes
Meta AI

Stefano Vigogna
University of Rome Tor Vergata

Abstract

Optimizing the misclassification risk is in general NP-hard. Tractable solvers can be obtained by considering a surrogate regression problem. While convergence to the regression function is typically sublinear, the corresponding classification error can decay much faster. Fast and super fast rates (up to exponential) have been established for general smooth losses on problems where a hard margin is present between classes. This leaves out models based on non-smooth losses such as support vector machines, and problems where there is no hard margin, begging several questions. Are such models incapable of fast convergence? Are they therefore structurally inferior? Is the hard margin condition really necessary to obtain exponential convergence? Developing a new strategy, we provide an answer to these questions. In particular, we show not only that support vector machines can indeed converge exponentially fast, but also that they can do so even without hard margin.

1 INTRODUCTION

To solve a problem with computer calculations, classical computer science consists in handcrafting a set of rules. In contrast, machine learning is based on the collection of a vast amount of solved instances of this problem, and on the automatic tuning of an algorithm that maps inputs defining the problem to the desired outputs. Denote the inputs by $x \in \mathcal{X}$, the outputs by $y \in \mathcal{Y}$, and the input/output mappings by $f : \mathcal{X} \rightarrow \mathcal{Y}$. To learn a mapping f^* , it is customary to introduce an explicit metric of error, and search for the function that minimizes it. Define this metric through a loss $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ that quantifies how bad a prediction $f(x)$ is when y is observed. Assuming the existence of a distribution ρ over $\mathcal{X} \times \mathcal{Y}$, that generates the instances of

the problem meant to be solved, one aims to minimize the average loss value

$$\mathcal{R}(f) = \mathbb{E}_{(X,Y) \sim \rho} [\ell(f(X), Y)]. \quad (1)$$

In practice, this “risk” \mathcal{R} can be evaluated approximately with samples $\mathcal{D}_n = (X_i, Y_i)_{i \leq n}$, collected by the machine learning scientist and assumed to have been drawn independently accordingly to ρ .

This work focuses on the binary classification problem where $\mathcal{Y} = \{-1, 1\}$, and ℓ is the zero-one loss $\ell(y, z) = \mathbf{1}_{y \neq z}$. In this setting, the risk $\mathcal{R}(f)$ captures the probability of mistakes of a classifier f , and its minimizer is characterized by

$$f^* = \arg \min_{f: \mathcal{X} \rightarrow \mathcal{Y}} \mathcal{R}(f) = \text{sign } \eta, \quad (2)$$

$$\text{where } \eta(x) = \mathbb{E}[Y | X = x]. \quad (3)$$

Ideally, leveraging the dataset \mathcal{D}_n , one would like to find a mapping $f_{\mathcal{D}_n} : \mathcal{X} \rightarrow \mathcal{Y}$ that is close to be optimal, in the sense that the excess of risk $\mathcal{E}(f_{\mathcal{D}_n}) = \mathcal{R}(f_{\mathcal{D}_n}) - \mathcal{R}(f^*)$ is as small as it could be. Since this quantity is actually random, inheriting from the randomness of the samples, statisticians focus on controlling its average. While classification is often the first problem described in introductory machine learning classes, several recent works have shown that, when the model is well-specified, as the number of samples grows, it is possible to show that this average decays much faster than what usual statistical learning theory suggests. This section provides a brief historical review of related literature before précising our contributions.

1.1 Statistical Learning Theory

The classical approach to minimize (1) without the knowledge of ρ but with the sole access to samples $\mathcal{D}_n \sim \rho^{\otimes n}$ is to restrict the search over functions in a class \mathcal{F} , and look for an empirical risk minimizer

$$f_{\mathcal{D}_n}^* \in \arg \min_{f \in \mathcal{F}} \mathcal{R}_{\mathcal{D}_n}, \quad (4)$$

$$\text{where } \mathcal{R}_{\mathcal{D}_n}(f) = \frac{1}{n} \sum_{i=1}^n \ell(f(X_i), Y_i).$$

If we denote by $f_{\mathcal{F}}^*$ the minimizer of \mathcal{R} in \mathcal{F} , using the fact that $\mathcal{R}_{\mathcal{D}_n}(f_{\mathcal{F}}^*) \geq \mathcal{R}_{\mathcal{D}_n}(f_{\mathcal{D}_n}^*)$, the excess of risk can be

bounded as

$$\begin{aligned} \mathcal{R}(f_{\mathcal{D}_n}^*) - \mathcal{R}(f^*) &\leq 2 \sup_{f \in \mathcal{F}} |\mathcal{R}(f) - \mathcal{R}_{\mathcal{D}_n}(f)| \quad (\text{estimation error}) \\ &\quad + \mathcal{R}(f_{\mathcal{F}}^*) - \mathcal{R}(f^*). \quad (\text{approximation error}) \end{aligned} \quad (5)$$

This bound can be seen as highly suboptimal because it bounds the deviation of a random function with the worst deviation in the function class. However, for any class \mathcal{F} , there exists an ‘‘adversarial’’ distribution ρ for which convergence rates (of the excess of risk toward zero as a function of the number of samples $n \in \mathbb{N}$) derived through this bound can not be improved beside lowering some multiplicative constants (Vapnik, 1995). On the one hand, the estimation error can be controlled with general tools to bound the supremum of a random process (e.g., Dudley, 1967), and will decrease as $O(n^{-1/2})$ with a multiplicative constant that depends on the size of the class \mathcal{F} . On the other hand, the approximation error depends on assumptions of the problem, and the bigger the size of the class \mathcal{F} , the less restrictive it will be to assume that f^* is not too different from $f_{\mathcal{F}}^*$. Hence, there is a clear trade-off between controlling both errors, which should be balanced in order to optimize a bound on the full excess of risk.

1.2 Surrogate Methods

In practice, due to the combinatorial nature of discrete-valued functions, finding the empirical risk minimizer (4) is often an intractable problem (e.g., Höffgen and Simon, 1992; Arora et al., 1997). Therefore, people have approached the original problem with other perspectives. A straightforward approach is given by *plug-in classifiers*, i.e., classifiers of the form $\text{sign } \hat{\eta}$, for $\hat{\eta}$ some estimator of η . For example, such an estimator can be constructed as $\hat{\eta}(x) = \sum_{i=1}^n \alpha_i(x) Y_i$, for $\alpha_i(x)$ some weights that specify how much the observation Y_i made at the point X_i should diffuse to the point x (see Friedman, 1994, for an example). Another popular approach to solve classification problems is provided by *support vector machines* (SVM), which were introduced from geometric considerations to maximize the margin between the classes $\{x \in \mathcal{X} \mid f^*(x) = y\}$ for $y \in \{-1, 1\}$ (Cortes and Vapnik, 1995).

These two approaches can be conjointly understood as introducing a surrogate loss $L : \mathbb{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ and looking for a continuous-valued function $g : \mathcal{X} \rightarrow \mathbb{R}$ that solves the surrogate problem defined by L before retrieving a solution $f : \mathcal{X} \rightarrow \mathcal{Y}$ of the original problem as

$$f = \text{sign } g, \quad g^* \in \arg \min_{g: \mathcal{X} \rightarrow \mathbb{R}} \mathcal{R}_S(g), \quad (6)$$

$$\text{where } \mathcal{R}_S(g) = \mathbb{E}_{(X,Y) \sim \rho} [L(g(X), Y)],$$

where the notation S stands for ‘‘surrogate’’. To an estimate $g : \mathcal{X} \rightarrow \mathbb{R}$ of g^* we associate an estimate $f : \mathcal{X} \rightarrow \mathcal{Y}$ of f^* through the decoding step $f = \text{sign } g$. In particular,

using the variational characterization of the mean, η can be estimated through $L(z, y) = |z - y|^2$. Meanwhile, SVM are related to the *hinge loss*, see Appendix B,

$$L(z, y) = (1 - zy)_+ = \max(0, 1 - zy), \quad (7)$$

Surrogate methods benefit from their relative easiness to optimize and the quality of their practical results. Arguably, they define the current state of the art in classification, softmax regression being particularly popular to train neural networks on classification tasks.

Surrogate methods were studied in depth by Zhang (2004); Bartlett et al. (2006), who proposed a generic framework to relate the excess of the original risk to the excess of surrogate risk through calibration inequalities of the type

$$\mathcal{R}(f) - \mathcal{R}(f^*) \leq \psi(\mathcal{R}_S(g) - \mathcal{R}_S(g^*)), \quad (8)$$

where $f = \text{sign } g$ and ψ is a concave function, uniquely defined from L and verifying $\psi(0) = 0$. The use of a concave function is motivated by Jensen inequality, allowing to integrate an inequality derived pointwise (conditionally on an input x).

When $\psi(x) \propto x$ and $\mathbb{E}_{\mathcal{D}_n} \mathcal{R}_S(g_{\mathcal{D}_n}) - \mathcal{R}_S(g^*)$ is controlled as $O(n^{-1/2})$ for $g_{\mathcal{D}_n}$ the minimizer of the empirical surrogate risk, one can control $\mathbb{E}_{\mathcal{D}_n} \mathcal{R}(\text{sign } g_{\mathcal{D}_n}) - \mathcal{R}(f^*)$ as $O(n^{-1/2})$. Because of minimax optimality of VC theory, those rates in $n^{-1/2}$ are wildly regarded as optimal. Yet, a closer look reveals that examples to show minimax optimality are based on degenerated data distributions. Those worst cases are unlikely to appear in practice. Indeed, some works have shown that under relatively mild low-noise conditions, much faster rates can be derived.

1.3 Exponential Convergence Rates

On the one hand, calibration inequalities (8) are appealing, as they allow casting directly rates derived on the surrogate problem to rates on the original problem. On the other hand, because ψ has to be concave, rates in $O(n^{-r})$ on the surrogate problem can not be cast as better rates on the original problem, corresponding to the optimal inequalities where $\psi(x) = cx$ for some $c > 0$. Yet, one can find cases where the sign of η can be estimated much faster than η itself, even when this sign is estimated with surrogate methods.

Most works that prove faster convergence rates are based on a specific derivation, which consists in relating the classification excess of risk with some power of the supremum norm. More specifically, Mammen and Tsybakov (1999) (see also Massart and Nédélec, 2006) introduced the following condition.

Assumption 1 (Hard low-noise condition). *The binary classification problem defined through the distribution ρ is said to verify the hard low-noise (or hard Tsybakov margin¹)*

¹Note that the wording margin here refers to a form of separation

condition if the conditional mean η is bounded away from zero, i.e.,

$$\exists \eta_0 > 0; \quad |\eta(X)| > \eta_0 \quad a.s., \quad (9)$$

where the notation *a.s.* stands for almost surely. Equivalently, $|\eta|^{-1} \in L^\infty(\rho_{\mathcal{X}})$.

Under Assumption 1, we know that $|g(x) - \eta(x)| < |\eta(x)|$ implies $\text{sign } g(x) = \text{sign } \eta(x) = f^*(x)$, thus $\mathcal{R}(\text{sign } g) - \mathcal{R}(f^*) \leq \mathbb{P}_X(\text{sign } g(X) \neq f^*(X)) \leq \mathbb{1}_{\|g - \eta\|_\infty \geq \eta_0}$. Hence, we get for any estimate $g_{\mathcal{D}_n} : \mathcal{X} \rightarrow \mathbb{R}$ computed from the dataset \mathcal{D}_n ,

$$\mathbb{E}_{\mathcal{D}_n}[\mathcal{R}(\text{sign } g_{\mathcal{D}_n})] - \mathcal{R}(f^*) \leq \mathbb{P}_{\mathcal{D}_n}(\|g_{\mathcal{D}_n} - \eta\|_{L^\infty} > \eta_0).$$

As a consequence, an exponential concentration inequality on the L^∞ distance between $g_{\mathcal{D}_n}$ and η directly translates to exponential convergence rates on the average excess of risk. In particular, estimation methods for η based on Hölder classes of functions, such as local polynomials, are known to be well-behaved with respect to the L^∞ norm (see, e.g., the construction of covering number by Kolmogorov and Tikhomirov, 1959). This was leveraged by Audibert and Tsybakov (2007) in a seminal paper that shows how better rates can be achieved on the classification problem under Assumption 1 and a variety of weaker conditions (described later in Assumption 2).

Surprisingly, such an approach has remained somehow less popular than approaches based on calibration inequalities, and we are missing a framework to fully apprehend fast rates phenomena. Some results for logistic regression were achieved by Koltchinskii and Beznosova (2005). Recently, Cabannes et al. (2021b) showed that this result generalizes to any discrete output learning problem, and that approaches that naturally lead to concentration in L^2 could be turned into fast rates based on interpolation inequalities that relate the L^2 norm with the L^∞ one (notably reusing the work of Fischer and Steinwart (2020) on interpolation spaces). Exploiting the work of Marteau-Ferey et al. (2019), this can be generalized to any self-concordant loss (using self-concordance to reduce the problem to a least-squares problem); and, through the work of Lin et al. (2020), to any spectral filtering technique (beyond Tikhonov regularization), such as stochastic gradient descent, which was actually shown earlier for binary classification by Pillaud-Vivien et al. (2018) and Nitanda and Suzuki (2019). In the same stream of research, Vigogna et al. (2022) proposed a general framework to study exponential rates for smooth losses in multiclass classification beyond least-squares. However, we believe that there is a bigger picture to be uncovered.

in the output space, and not to the usual margin of SVM that describes separation between classes in the input space.

1.4 Contribution

The proofs of exponential convergence in the works quoted above are all based on the basic mechanism outlined in Audibert and Tsybakov (2007), which, in substance, consists in relating the excess of classification risk with concentration on the supremum norm as, with notations of Assumption 2,

$$\mathcal{R}(\text{sign } g) - \mathcal{R}(\text{sign } g^*) \leq c \|g - g^*\|_\infty^{p+1},$$

Unfortunately, such a mechanism relies on L^∞ concentration, and does not easily extend to the hinge loss. Does this mean that support vector machines do not exhibit superfast rates, and thus they are inferior to other surrogate methods? The practice seems to answer negatively. In this paper, we give a firm theoretical answer to this question. In particular, we show not only that support vector machines do achieve exponential rates, but also that they can do so even without assuming the hard low-noise condition. Our main contribution is to introduce a general framework to prove exponential convergence rates, and show how this framework can be applied to the hinge loss while only considering classical assumptions. Stated otherwise, our goal is to show fast rates for a non-smooth loss, which we will do for the hinge loss formalization of kernelized SVM.

Outline. Our general strategy is illustrated on Figure 1 and consists in first finding a relation

$$\mathcal{R}_S(g_\theta) - \mathcal{R}_S(g_{\theta^*}) \geq \|\theta - \theta^*\|,$$

for some natural parameter θ in a Banach space Θ parametrizing a class of functions $g_\theta \in \mathcal{F}$, and then show that $\text{sign } g_\theta = \text{sign } g_{\theta^*}$ when $\|\theta - \theta^*\|$ is small enough, that is,

$$\exists \varepsilon > 0; \quad \|\theta - \theta^*\| \leq \varepsilon \quad \Rightarrow \quad \text{sign } g_\theta = \text{sign } g_{\theta^*}.$$

Assuming that $\text{sign } g_{\theta^*} = f^*$, we deduce that

$$\begin{aligned} & \mathbb{E}_{\mathcal{D}_n}[\mathcal{R}(\text{sign } g_{\theta_n})] - \mathcal{R}(f^*) \\ &= \mathbb{E}_{\mathcal{D}_n}[\mathcal{R}(\text{sign } g_{\theta_n}) - \mathcal{R}(\text{sign } g_{\theta^*})] \\ &\leq \mathbb{E}_{\mathcal{D}_n}[\mathbb{1}_{\text{sign } g_{\theta_n} \neq (\text{sign } g_{\theta^*})}] \\ &\leq \mathbb{E}_{\mathcal{D}_n}[\mathbb{1}_{\|\theta_n - \theta^*\| \geq \varepsilon}] \\ &\leq \mathbb{P}_{\mathcal{D}_n}(\|\theta_n - \theta^*\| \geq \varepsilon) \\ &\leq \mathbb{P}_{\mathcal{D}_n}(\mathcal{R}_S(g_{\theta_n}) - \mathcal{R}_S(g_{\theta^*}) \geq \varepsilon). \end{aligned}$$

where g_{θ_n} is an estimate of g_θ based on the samples \mathcal{D}_n . Finally, we conclude with an exponential concentration inequality that controls the deviation of the excess of risk based on classical statistical learning theory.

2 EXPONENTIAL CONVERGENCE OF SVM

This section is devoted to the proof of exponential convergence rates for the hinge loss. We shall fix the notation \mathcal{R}_S

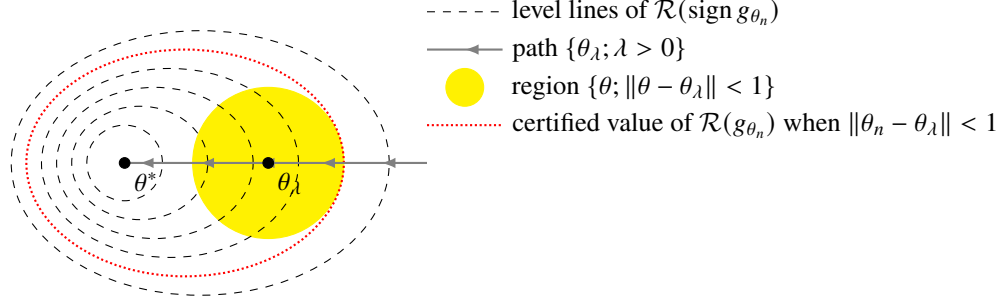


Figure 1: Our convergence analysis consists in relating natural concentration given by surrogate methods to the original excess of risk without passing by the surrogate excess of risk. We denote by θ_λ the regularized surrogate risk minimizer $\arg \min \mathcal{R}_S(g_\theta) + \lambda \|\theta\|^2$. As the drawing shows, concentration in parameter space Θ can be cast as deviation on the original excess of risk. Yet, such a casting relation depends on the geometry of this picture, which itself depends on what surrogate is used, what is the function to learn, how a regularized estimator approached it, and how our empirical estimate concentrates around the regularized estimator. Note that this figure illustrates an abstract mechanism that generalizes the simpler mechanism we use to derive exponential convergence rates. It contrasts with usual statistical learning theory that combines approximation and estimation error in an additive fashion.

as the surrogate risk associated with (7). All the proofs are collected in Appendix A.

2.1 Refined Calibration for the Hinge Loss

We start by introducing the classical weak low-noise condition (Mammen and Tsybakov, 1999).

Assumption 2 (Weak low-noise condition). *The binary classification problem defined through the distribution ρ is said to verify the p -low-noise condition, with $p \in (0, \infty)$, if there exists a constant $c > 0$ such that*

$$\mathbb{P}_{\rho_{\mathcal{X}}}(0 < |\eta(X)| < t) \leq ct^p, \quad (10)$$

where the notation $\rho_{\mathcal{X}}$ denotes the marginal of ρ over \mathcal{X} .

Assumption 2 is equivalent to asking for the inverse of the conditional mean $|\eta|^{-1}$ (with the convention $0^{-1} = 0$) to belong to the Lorentz space $L^{p,\infty}(\rho_{\mathcal{X}})$ (also known as weak- L^p space), which is the Banach space endowed with the norm (quasi-norm and quasi-Banach if $p < 1$)

$$\|f\|_{p,\infty} = \sup_{t>0} t \mathbb{P}_{\rho_{\mathcal{X}}}(f(X) > t)^{\frac{1}{p}}, \quad (11)$$

where the $\rho_{\mathcal{X}}$ denotes the marginal of ρ with respect to \mathcal{X} . This definition can be extended to the case $p = \infty$ by setting $L^{p,\infty}(\rho_{\mathcal{X}}) = L^\infty(\rho_{\mathcal{X}})$, which characterizes the hard low-noise condition in Assumption 1. We will also use $\|\cdot\|_p$, for $p \in [1, \infty]$, to denote the L^p -norm on \mathcal{X} endowed with $\rho_{\mathcal{X}}$.

We now relate the excess of risk on the hinge loss to the deviation in these spaces.

Lemma 1 (Weak- L^q concentration due to the hinge loss). *For any functions $g_1, g_2 : \mathcal{X} \rightarrow [-1, 1]$,*

$$\mathcal{R}_S(g_2) - \mathcal{R}_S(g_1) = \mathbb{E}_{\rho_{\mathcal{X}}}[-\eta(X)(g_2(X) - g_1(X))]. \quad (12)$$

In particular, under Assumption 1, for any $g : \mathcal{X} \rightarrow \mathbb{R}$,

$$\mathcal{R}_S(g) - \mathcal{R}_S(g^*) \geq \|\eta\|^{-1} \|\pi(g) - g^*\|_1. \quad (13)$$

where $g^* = \text{sign } \eta$ is a minimizer of \mathcal{R}_S and π is the projection of \mathbb{R} on $[-1, 1]$, defined as mapping $t \in \mathbb{R}$ to $\pi(t) = \text{sign}(t) \min\{|t|, 1\}$. Similarly, under Assumption 2, with $q = p/p + 1$,

$$\mathcal{R}_S(g) - \mathcal{R}_S(g^*) \geq 2^{-1} \|\eta\|^{-1} \|\pi(g) - g^*\|_{q,\infty}. \quad (14)$$

Lemma 1 shows that we can set the minimizer $g^* = f^* \in \{-1, 1\}^{\mathcal{X}}$. This is a useful fact as it implies that the excess of the original risk is zero as soon as $\|g - g^*\|_\infty < 1$. In essence, the only piece missing in order to prove fast convergence rates is an interpolation inequality between $L^{q,\infty}$ and L^∞ . In the following, we will leverage Lemma 1 more subtly by considering a class of functions \mathcal{G} and assumptions on the distribution $\rho_{\mathcal{X}}$ such that, if an estimate $g \in \mathcal{G}$ has not the same sign almost everywhere as the estimand g^* , then $\|g - g^*\|_{q,\infty}$ is bounded away from zero. By contraposition, if $g \in \mathcal{G}$ presents a small excess of surrogate risk, then $\text{sign } g = \text{sign } g^*$. When \mathcal{X} is a metric space, one way to proceed is to assume that g is Lipschitz-continuous, together with some minimal mass assumptions. Let us begin with the minimal mass assumption. We first need the following definition.

Definition 1 (Well-behaved sets). *A set $U \subset \mathcal{X}$ is said to be well-behaved with respect to ρ if there exist constants $c, r > 0$ and an exponent $d > 0$ such that, for any $x \in U$,*

$$\forall \varepsilon \in [0, r]; \quad \rho_{\mathcal{X}}(U \cap \mathcal{B}(x, \varepsilon)) \geq c\varepsilon^d, \quad (15)$$

and $\mathcal{B}(x, \varepsilon)$ the ball in \mathcal{X} of center x and radius ε .

The following examples show that the coefficient d that appears in (15) results from the dimension of the ambient space, the regularity of singularities of the border of the set, and the decay of the density when approaching the frontier of the set.

Example 1. *The set $[0, 1]^p$ is well-behaved with coefficients $r = 1$, $d = p$ and $c = 2^{-d} \text{vol}(\mathbb{S}^{d-1})$ with respect to the Lebesgue measure in \mathbb{R}^d .*

Example 2. The set $\{(x, y) \in \mathbb{R}^2 \mid x \in [0, 1], y \in [0, x^{n-1}]\}$ is well-behaved with coefficient $r = 1$, $d = n$ and $c = n^{-1}$ with respect to the Lebesgue measure. Reciprocally, the set $[0, 1]$ is well-behaved with coefficient $r = 1$, $d = n$ and $c = n^{-1}$ with respect to the measure whose density equals $\rho(x) = x^{n-1}$.

Assumption 3 (Minimal mass assumption). The classification problem is said to verify the d -minimal mass assumption if the decision regions $\mathcal{X}_y = \{x \in \text{supp } \rho_{\mathcal{X}} \mid y\eta(x) > 0\}$ for $y \in \{-1, 1\}$ are both well-behaved with exponent d .

Assumption 3 is a weakening of an assumption that is commonly found in the statistical learning literature. More precisely, it is often assumed that ρ is absolutely continuous according to the Lebesgue measure λ on \mathcal{X} (assumed to be a Euclidean space), that its density is bounded away from zero on its support, and that its support has smooth boundary, so that $\lambda(\text{supp } \rho_{\mathcal{X}} \cap \mathcal{B}(x, \varepsilon)) > c'\lambda(\mathcal{B}(x, \varepsilon))$ (see the strong density assumption in Audibert and Tsybakov, 2007).

The minimal mass requirement allows relating misclassification events to $L^{q,\infty}$ deviation.

Lemma 2. Under Assumption 3, there exists a constant c_0 such that if g is G -Lipschitz-continuous for $G > r^{-1}$, for any $q \in (0, 1]$

$$\begin{aligned} & \exists x \in \text{supp } \rho_{\mathcal{X}}, \quad |g(x) - g^*(x)| \geq 1 \\ \Rightarrow & \|g - g^*\|_{q,\infty} \geq c_0 G^{-\frac{d}{q}}. \end{aligned} \quad (16)$$

Putting together Lemmas 1 and 2, we obtain the following refined calibration.

Proposition 1. Under Assumptions 2 and 3, if g is G -Lipschitz-continuous with $G > r^{-1}$, we have

$$\begin{aligned} & \mathcal{R}_S(g) - \mathcal{R}_S(g^*) \leq 2^{-1} \|\eta\|_{p,\infty}^{-1} c_0 G^{-\frac{d(p+1)}{p}} \\ \Rightarrow & \mathcal{R}(\text{sign } g) = \mathcal{R}(f^*). \end{aligned} \quad (17)$$

2.2 Trade-off between Estimation and Approximation Errors

We are now left with the research of $g_{\mathcal{D}_n}$ inside a class of Lipschitz-continuous functions such that $\mathcal{R}_S(g_{\mathcal{D}_n}) - \mathcal{R}_S(g^*)$ is sub-Gaussian (its randomness being inherited from the dataset \mathcal{D}_n from which $g_{\mathcal{D}_n}$ is built). To do so, let us consider a linear (a.k.a. kernelized) class of functions

$$\mathcal{G}_{M,\sigma} = \{x \mapsto \langle \theta, \varphi(\sigma^{-1}x) \rangle \mid \theta \in \mathcal{H}, \|\theta\|_{\mathcal{H}} \leq M\}, \quad (18)$$

where \mathcal{H} is a separable Hilbert space, $\varphi : \mathcal{X} \rightarrow \mathcal{H}$ is a G_φ -Lipschitz-continuous mapping, and $\sigma > 0$ is a scaling (or bandwidth) parameter. Such a class of functions can be entirely described from the kernel $k(x, x') = \langle \varphi(x), \varphi(x') \rangle$ (see Scholkopf and Smola, 2001, for a primer on kernel methods). An example for \mathcal{G} is given by the Gaussian kernel, a.k.a. radial basis function, $k(x, x') = \exp(-d(x, x')^2)$.

Using Cauchy-Schwarz, it is easy to show that any function in $\mathcal{G}_{M,\sigma}$ is $MG_\varphi\sigma^{-1}$ -Lipschitz-continuous.

In order to find a function $g_{\mathcal{D}_n}$ that is likely to minimize \mathcal{R}_S without accessing the distribution ρ , but only i.i.d. samples $\mathcal{D}_n = (X_i, Y_i)_{i \leq n} \sim \rho^{\otimes n}$, it is classical to consider the empirical risk minimizer

$$g_{\mathcal{D}_n} \in \arg \min_{g \in \mathcal{G}_{M,\sigma}} \frac{1}{n} \sum_{i=1}^n L(g(X_i), Y_i). \quad (19)$$

This problem is convex with respect to θ parametrizing $g \in \mathcal{G}_{M,\sigma}$, and is easily optimized with duality. We refer the curious reader to the extensive literature on kernelized SVM (see Cristianini and Shawe-Taylor, 2000; Scholkopf and Smola, 2001; Steinwart and Christmann, 2008, for books on the matter).

In order to show that $\mathcal{R}_S(g_{\mathcal{D}_n})$ is close to $\mathcal{R}_S(g^*)$, one can apply classical results from statistical learning theory, and in particular (5). The estimation error can be bounded using the extensive literature on Rademacher complexity for linear classes of functions on Lipschitz-continuous losses (Bartlett and Mendelson, 2002). To bound the approximation error, one needs to make additional assumptions on the problem. We refer to Steinwart and Scovel (2007); Blaschzyk and Steinwart (2018) for advanced considerations on the matter. In view of our calibration result (17), the following additional assumption suffices to prove exponential convergence of SVM.

Assumption 4 ((c_0, p, d) -Source condition). There exist M, σ and a function $g \in \mathcal{G}_{M,\sigma}$ such that $\mathcal{R}_S(g) - \mathcal{R}_S(g^*) \leq 4^{-1} \|\eta\|_{p,\infty}^{-1} c_0 M^{-r} G_\varphi^{-r} \sigma^r$ with $r = d(p+1)/p$.

It should be noted that, because of Proposition 1, the function g in Assumption 4 is a perfect classifier. This implies that the decision frontier $\overline{\mathcal{X}_{-1}} \cap \overline{\mathcal{X}_1}$ (the bar notation corresponding to space closure) inherits from the regularity of g , since it is included in the set $\{x \in \mathcal{X} \mid g(x) = 0\}$. In particular, if $\mathcal{G}_{M,\sigma}$ is included in \mathcal{C}^m , this frontier would be in \mathcal{C}^m . Hence, for Assumptions 4 to hold, the boundary frontier should match the regularity implicitly defined by $\mathcal{G}_{M,\sigma}$.

We are finally ready to state our main result, establishing exponential convergence rates for SVM.

Theorem 1 (Exponential convergence rates for SVM). Under Assumptions 2, 3 and 4, there exists a constant $c > 0$ such that the empirical minimizer $g_{\mathcal{D}_n}$ defined by (19) verifies

$$\mathbb{E}_{\mathcal{D}_n} \mathcal{R}(\text{sign } g_{\mathcal{D}_n}) - \mathcal{R}(f^*) \leq 2 \exp(-cn). \quad (20)$$

While this result is achieved for constrained SVM, the same result can be achieved with regularized SVM, which is used in practice.

Corollary 1. Under Assumptions 2, 3 and 4, there exist $\lambda \geq 0$ and a constant $c > 0$ such that $g_{\mathcal{D}_n,\lambda}$ verifies

$$\mathbb{E}_{\mathcal{D}_n} \mathcal{R}(\text{sign } g_{\mathcal{D}_n,\lambda}) - \mathcal{R}(f^*) \leq 2 \exp(-cn).$$

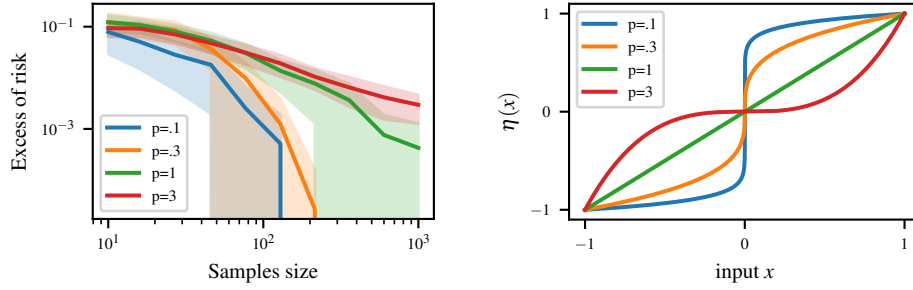


Figure 2: SVM generalization error as a function of the number of samples (left) for a problem where X is uniform on $[-1, . - 1] \cup [.1, 1]$ and $\eta(x) = \text{sign}(x) |x|^p$ (right). We observe exponential convergence rates on the blue and orange curve. The green and red curves showcase regimes where exponential convergence rates have not been reached yet, one would need more samples to observe them.

where

$$g_{\mathcal{D}_n, \lambda} = \arg \min_{g; \theta \in \mathcal{H}} \sum_{i=1}^n L(g_\theta(X_i), Y_i) + \lambda \|\theta\|_{\mathcal{H}}^2.$$

2.3 Relaxing Assumptions

Exponential convergence rates rely on strong assumptions in order to set the approximation error to zero. In particular, it is customary to assume that the surrogate function to learn lies in the model we have chosen, that is, in our notation, $g^* \in \mathcal{G}_{M, \sigma}$. In our case this would be a strong assumption, since g^* is piecewise constant while $\mathcal{G}_{M, \sigma}$ is a smooth space of functions. It turns out that the assumption $g^* \in \mathcal{G}_{M, \sigma}$ is not necessary, and what we actually need is the ability to reach a sufficiently small risk within the class $\mathcal{G}_{M, \sigma}$. How small is enough is quantified by the statement of Proposition 1 and Assumption 4. In particular, this assumption is verified when the function class $\mathcal{G}_{M, \sigma}$ is rich enough and classes are separated by a margin in the input space as specified by the following assumption.

Assumption 5 (Cluster Assumption). *The classes $\mathcal{X}_1 = (f^*)^{-1}(1) := \{x \in \text{supp } \rho_{\mathcal{X}} \mid \eta(x) > 0\}$ and $\mathcal{X}_{-1} = (f^*)^{-1}(-1)$ are separated by a margin, in the sense that the distance between any two points in each set is bounded away from zero. Formally,*

$$\exists \delta_0 > 0; \quad \forall (x, x') \in \mathcal{X}_1 \times \mathcal{X}_{-1}, \quad d(x, x') \geq \delta_0. \quad (21)$$

Proposition 2 (Source condition example). *Assumption 5 implies Assumption 4 for any (c_0, p, d) , as long as \mathcal{X} is a Euclidean space and k is taken as the exponential kernel $k(x, y) = \exp(-\sigma^{-1} \|x - y\|)$ for any $\sigma > 0$.*

In terms of practical applications, the cluster assumption says that no one can continuously modify an input to go from a region of the space linked with one class to a region linked with another class without going through inputs that will never exist. This is typically true for well-curated image datasets such as CIFAR10: one can not continuously

transform an image of a truck into an image of a horse without going through images that will never appear in the CIFAR10 dataset (Krizhevsky, 2009). As stated in the seminal work of Seeger (2001), “the ‘cluster assumption’ is a very general and weak assumption, therefore applicable as prior assumption to many unsupervised tasks”. It has been popular in unsupervised, weakly-supervised and semi-supervised learning (see Rigollet, 2007; Cabannes et al., 2021a, for exponential convergence rates in those settings).

To deepen the study of the approximation error, one could leverage the following geometrical characterization of the risk of misclassification. For $f : \mathcal{X} \rightarrow \{-1, 1\}$, we have

$$\begin{aligned} \mathcal{R}(f) - \mathcal{R}(f^*) &= \mathbb{E}[|\eta(X)| \mathbf{1}_{f(X) \neq f^*(X)}] \\ &\leq \mathbb{P}(f(X) \neq f^*(X)) = \rho_{\mathcal{X}}(f^{-1}(\{1\}) \Delta \mathcal{X}_1), \end{aligned} \quad (22)$$

where Δ denotes the symmetric difference of sets, *i.e.* $A \Delta B = (A \cup B) \setminus (A \cap B)$. In particular, under Assumption 5 the minimizer $g_{\mathcal{G}_{M, \sigma}}$ of the surrogate risk in $\mathcal{G}_{M, \sigma}$ verifies

$$\rho_{\mathcal{X}} \left((\text{sign } g_{\mathcal{G}_{M, \sigma}})^{-1}(\{1\}) \Delta \mathcal{X}_1 \right) \leq \psi(M, \sigma), \quad (23)$$

for ψ a function that vanishes for sufficiently large M and small σ .

On the one hand, one could control the approximation error by assuming or deriving inequalities akin to (22) and (23), with different profiles of ψ . We conjecture that this can be done by assuming low-noise conditions that are well adapted to the geometric nature of SVM, such as the one proposed by Steinwart and Scovel (2007) (see also Gentile and Warmuth, 1999; Cristianini and Shawe-Taylor, 2000). On the other hand, the estimation error can be controlled by extending the ideas presented in this paper to study the worst value of the estimation error $\mathcal{R}(\text{sign } g_{\mathcal{D}_n}) - \mathcal{R}(\text{sign } g_{\mathcal{G}_{M, \sigma}})$ under the knowledge of $\mathcal{R}_S(g_{\mathcal{D}_n}) - \mathcal{R}_S(g_{\mathcal{G}_{M, \sigma}})$. Fitting M and σ to trade estimation and approximation errors, such derivations would open the way to fast polynomial rates under less restrictive assumptions.

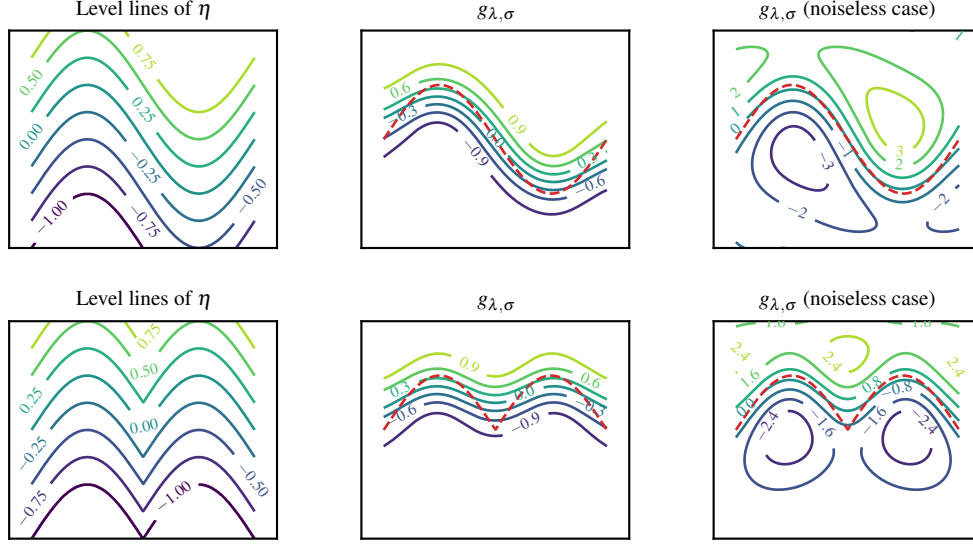


Figure 3: Study of the level lines of $g_{\lambda,\sigma}$ when $\eta^{-1}(0) \in C^\infty$ (top) and $\eta^{-1}(0) \in C^0 \setminus C^1$ (bottom). The function g^* takes values -1 below the optimal decision frontier plotted in red and $+1$ above, independently of the noise. We observe that the bias error $\mathcal{R}(\text{sign } g_{\lambda,\sigma}) - \mathcal{R}(f^*)$, which is bounded by the volume between the level lines $\{x \in \mathcal{X} \mid g_{\lambda,\sigma}(x) = 0\}$ and $\{x \in \mathcal{X} \mid \eta(x) = 0\}$ (plotted in red), depends on both the regularity of the latter, and on the noise level. Here, σ is taken to be of the order of 15% of the diameter of the domain, which explains the regularity of the observed level lines. The noiseless cases on the right correspond to the situations where $\mathbb{E}[Y|X] = \text{sign } \eta(X)$ for η plotted on the left.

2.4 Prior Assumptions

In this section, we discuss the mildness of our assumptions in comparison to existing proofs of exponential convergence rates. Those proofs are based on the combination of the hard low-noise condition together with the existence of a regular function close to η , which implies Assumption 5. In contrast, *the cluster assumption is both stronger than our assumptions and weaker than the existing assumptions to prove exponential convergence rates.*

To our knowledge, Audibert and Tsybakov (2007) first proves exponential convergence rates. They did so by assuming the regression function η to be Hölder-continuous, together with the hard Tsybakov condition (see Proposition 3.7 and the class \mathcal{P}). Under the hard Tsybakov condition, $\mathcal{X}_1 = \eta^{-1}([\eta_0, +\infty))$, hence, for $x \in \mathcal{X}_1$ and $x' \in \mathcal{X}_{-1}$, if η is (L, α) -Hölder

$$2\eta_0 \leq \|\eta(x) - \eta(x')\| \leq L \|x - x'\|^\alpha.$$

It follows that

$$d(\mathcal{X}_1, \mathcal{X}_{-1}) \geq \left(\frac{2\eta_0}{L}\right)^{1/\alpha}.$$

In other terms, those two assumptions implies the cluster assumption 5. Similarly, Rigollet (2007) proved exponential convergence rates under the cluster assumption for semi-supervised learning.

More recently, a renewed interest was triggered by results for SGD achieved by Pillaud-Vivien et al. (2018). Once

again, the authors assumed both the hard Tsybakov margin condition (A1) and the existence of a perfect classifier that is in the linear space of functions considered (implied by A4). They discuss the fact that their assumptions are met under the cluster assumption (A5) plus some regularity of η (see Proposition 3). Indeed, the cluster assumptions is a necessary condition for (A1+A4) to hold when the considered space of functions is included in the class of Hölder functions (which is true for all classical reproducing kernel Hilbert space). This can be proven with the same derivations as the one above (replacing η by g_λ and η_0 by $\delta/2$ with their notations). In contrast, we do not need regularity of η , which we show in practice on Figure 4.

3 NUMERICAL ANALYSIS

In this section, we provide experiments to illustrate and validate our theoretical findings. In order to be inline with the current practice of machine learning, instead of considering the hard constraint $\|\theta\| \leq M$ when minimizing a risk functional, we add a penalty $\lambda \|\theta\|^2$ to the risk to be minimized. Going from a constrained to a penalized framework does not change the nature of the statistical analysis, and one might loosely think of λ as $1/M$ (see, for example, Bach, 2023). All experiments are made with the Gaussian kernel. Precise details of the different settings are provided in Appendix C.

First, we observe that the regime described in this paper kicks in when the error is already pretty small. On many

real-world problems, we do not expect the generalization error as a function of the number of data used for training to exhibit a clear exponential behavior until an unusually big number of samples is used. This fact is illustrated on Figure 2, where for hard problems, the exponential behaviors still do not kick in after a thousand of samples.

Second, this paper shows that, in order to get exponential convergence rates for SVM, one needs the minimizer $g_{M,\sigma}$ of the surrogate risk over the selected class of functions to be a perfect classifier, *i.e.* its sign equals the sign of g^* . While this is not constraining under the cluster assumption, we inspect divergences from this condition on Figure 3. We observe that, even if g^* does not depend on the noise, $g_{M,\sigma}$ does. We also observe that the regularity of the decision boundary $\{x \in \mathcal{X} \mid \eta(x) = 0\}$ should match the regularity defined implicitly by the kernel k and the scale parameter σ .

Experimental comparisons of different classification approaches have been done by many people, and our goal is not to showcase the superiority of the SVM over least-squares, which might be considered as general wisdom that led to the golden age of SVM in the pre-deep-learning area (see Joachims, 1998, for example). In comparison with previous works based on calibration inequalities (Rosasco et al., 2004; Steinwart, 2007), our analysis proves the robustness of SVM to noise far away from the decision boundary, in the sense that one does not need η to be bounded away from zero. This is a distinctive aspect of SVM compared to smooth surrogate methods (Nowak-Vila et al., 2020), such as softmax regression, that implicitly estimate conditional probabilities and whose performance depends on the regularity of η . We illustrate this fact graphically on Figure 4.

4 LIMITATIONS

Are Surrogate Methods Only a Proxy for Classification?

From a theoretical perspective, if we are only interested in the optimal mapping $f : \mathcal{X} \rightarrow \mathcal{Y}$, learning surrogate quantities can be seen as a waste of resources. In essence, this waste of resources is similar to the one occurring when we learn the full probability function $(p(y))_{y \in \mathcal{Y}}$ for some probability distribution p on \mathcal{Y} , while we only care about its mode. Yet, in practice, what we call a “surrogate” problem might actually be a problem of prime interest when we do not only want to predict $f^*(x)$, but we would also like to know how much we can confidently discard other potential outputs for an input x . Furthermore, assuming that a problem is exactly defined through an “original” loss that defines a clear and unique measure of error can be questioned when some practitioners evaluate methods with several metrics of performance (*e.g.* Chowdhery et al., 2022).

Do PAC-Mounds Provide Confidence Levels? Since the parameters in Assumptions 2 and 3 are hard to estimate in practice, it would be difficult to directly plug our bounds

into a practical problem to derive confidence levels on how much error one might expect when deploying a model in production. Less ambitiously, we see theorems akin to Theorem 1 as providing theoretical indications that a learning method or a set of hyperparameters is sound. This is a generic downfall of probably approximately correct (PAC) generalization bounds (Valiant, 2013), which might explain why practitioners often prefer to derive error indications from test samples (see, *e.g.*, Géron, 2017). Along this line, research on conformal prediction provides interesting considerations to obtain useful confidence information from test samples (Vovk and Shafer, 2008). Finally, all these statistical methods to get confidence intervals assume representative (if not i.i.d.) data, an assumption sometimes hard to meet in practice, which is a problem that has found echoes in the civil society (*e.g.* Benjamin, 2019).

Societal Impact The theoretical nature of the present work prevents us to discuss its potential negative social impact without questioning the impact of the whole field of machine learning, which is out of scope of the current paper.

5 CONCLUSIONS

In this work, we were keen to illustrate a simple mechanism to get exponential convergence rates for support vector machines. Our proof relates the misclassification to the surrogate risk through a sort of $L^p - L^q$ interpolation inequality. Thanks to this new strategy, we were able to deal with a non-smooth loss such as the hinge loss, which is quite popular and whose understanding can not be easily reduced to previous work. Remarkably, our assumptions are strictly weaker than all the existing assumptions we are aware of used to prove exponential convergence rates. In particular, we showed that the hard low-noise condition is not crucial in order to derive exponential convergence rates for the SVM.

This provides a crucial step to better understand convergence rates on classification problems. An extension to generic discrete output problems could be made by considering polyhedral losses, and deriving variants of Lemma 1 (see Frongillo and Waggoner, 2021, for calibration inequalities for such losses). An important follow-up would be to provide a more global picture of fast polynomial rates for SVM under relaxations of Assumptions 3 and 4.

Finally, Chizat and Bach (2020) have made a link between two-layer wide neural networks in the interpolation regime (which implies Assumption 1 with $\eta_0 = 1$) and max-margin classifiers over specific linear classes of functions. As a consequence, we could directly plug in our analysis to prove exponential convergence rates for those small neural networks in this noiseless setting. Studying rates, constants and hyperparameter tuning in this setting would be of particular interest if it was to provide practical guidelines to deep learning practitioners in the spirit of Yang et al. (2021).

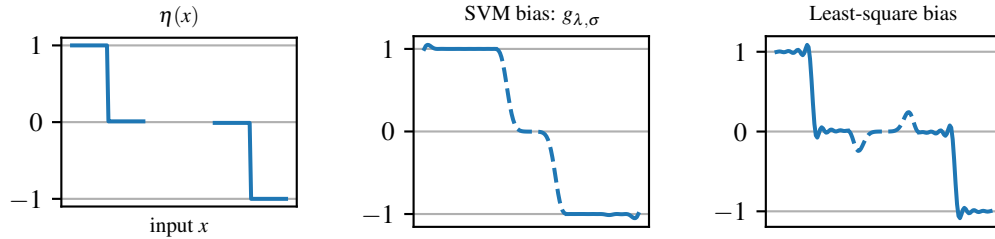


Figure 4: Comparison of the regularized risk (i.e. $\mathcal{R}_S(g_\theta) + \lambda \|\theta\|^2$) minimizer for the hinge loss surrogate (middle) and the least-squares surrogate (right), when η is not regular (left). In this setting, the hinge loss is minimized for $g = \text{sign}(\eta)$, which can be chosen regular, while the least-squares loss is minimized for $g = \eta$, which can not be chosen regular. The reconstruction is made with σ about 3% of the domain diameter, and λ relatively small. We assume no density in the middle of the domain, explaining the absence of definition of η and the dashed lines of the right figures. The oscillation on the later figures is related to the Gibbs phenomenon (Wilbraham, 1848). This phenomenon prevents the regularized least-squares solution from being a perfect classifier.

Acknowledgements

VC would like to thank Alex Nowak-Vila for sharing insights that led to this work, as well as Francis Bach for useful comments. SV is partially supported by the MIUR Excellence Department Project MatMod@TOV awarded to the Department of Mathematics, University of Rome Tor Vergata. The authors are also deeply grateful for the highly valuable feedback from the anonymous reviewers.

References

- Arora, S., Babai, L., Stern, J., and Sweedyk, E. (1997). The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*.
- Audibert, J.-Y. and Tsybakov, A. (2007). Fast learning rates for plug-in classifiers. *The Annals of Statistics*.
- Bach, F. (2023). *Learning Theory from First Principles*. To appear at MIT press.
- Bartlett, P., Jordan, M., and McAuliffe, J. (2006). Convexity, classification, and risk bounds. *Journal of the American Statistical Association*.
- Bartlett, P. and Mendelson, S. (2002). Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*.
- Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity.
- Blaschzyk, I. and Steinwart, I. (2018). Improved classification rates under refined margin conditions. *Electronic Journal of Statistics*.
- Cabannes, V., Rudi, A., and Bach, F. (2021a). Disambiguation of weak supervision with exponential convergence rates. In *International Conference on Machine Learning*.
- Cabannes, V., Rudi, A., and Bach, F. (2021b). Fast rates in structured prediction. In *Conference on Learning Theory*.
- Castillo, R. E. and Rafeiro, H. (2016). *An Introductory Course in Lebesgue Spaces*. Springer.
- Chang, C. and Lin, C. (2011). LIBSVM: A library for support vector machines. *ACM TIST*.
- Chizat, L. and Bach, F. (2020). Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*.
- Chowdhery, A. et al. (2022). PaLM: Scaling language modeling with pathways. Technical Report 2204.02311, arXiv.
- Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Journal of Machine Learning*.
- Cristianini, N. and Shawe-Taylor, J. (2000). *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge university press.
- Dudley, R. (1967). The sizes of compact subsets of hilbert space and continuity of Gaussian processes. *Journal of Functional Analysis*.
- Fischer, S. and Steinwart, I. (2020). Sobolev norm learning rates for regularized least-squares algorithms. *Journal of Machine Learning Research*.
- Friedman, J. (1994). Flexible metric nearest neighbor classification. Technical report, Department of Statistics, Stanford University.
- Frongillo, R. and Waggoner, B. (2021). Surrogate regret bounds for polyhedral losses. In *Advances in Neural Information Processing Systems*.
- Gentile, C. and Warmuth, M. (1999). Linear hinge loss and average margin. In *Advances in Neural Information Processing Systems*.
- Géron, A. (2017). *Hands-On Machine Learning with Scikit-Learn & TensorFlow*. OReilly.
- Harris, C. et al. (2020). Array programming with NumPy. *Nature*.

-
- Höfgen, K.-U. and Simon, H. U. (1992). Robust trainability of single neurons. In *Computational Learning Theory*.
- Hunter, J. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*.
- Joachims, T. (1998). Text categorization with support vector machines: Learning with many relevant features. In *European Conference on Machine Learning*.
- Kolmogorov, A. and Tikhomirov, V. (1959). ε -entropy and ε -capacity of sets in functional spaces. *Uspekhi Matematicheskikh Nauk*.
- Koltchinskii, V. and Beznosova, O. (2005). Exponential convergence rates in classification. In *International Conference on Computational Learning Theory*.
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Technical report, Canadian Institute for Advanced Research.
- Lin, J., Rudi, A., Rosasco, L., and Cevher, V. (2020). Optimal rates for spectral algorithms with least-squares regression over Hilbert spaces. *Applied and Computational Harmonic Analysis*.
- Mammen, E. and Tsybakov, A. (1999). Smooth discrimination analysis. *The Annals of Statistics*.
- Marteau-Ferey, U., Ostrovskii, D., Bach, F., and Rudi, A. (2019). Beyond least-squares: Fast rates for regularized empirical risk minimization through self-concordance. In *Conference on Learning Theory*.
- Massart, P. and Nédélec, É. (2006). Risk bounds for statistical learning. *The Annals of Statistics*.
- Maurer, A. (2016). A vector-contraction inequality for rademacher complexities. In *Algorithmic Learning Theory*.
- Nitanda, A. and Suzuki, T. (2019). Stochastic gradient descent with exponential convergence rates of expected classification errors. In *International Conference on Artificial Intelligence and Statistics*.
- Nowak-Vila, A., Bach, F., and Rudi, A. (2020). A general theory for structured prediction with smooth convex surrogates. Technical Report 1902.01958, ArXiv.
- Pedregosa, F. et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*.
- Pillaud-Vivien, L., Rudi, A., and Bach, F. (2018). Exponential convergence of testing error for stochastic gradient methods. In *Conference on Learning Theory*.
- Rigollet, P. (2007). Generalization error bounds in semi-supervised classification under the cluster assumption. *Journal of Machine Learning Research*.
- Rosasco, L., Vito, E. D., Caponnetto, A., Piana, M., and Verri, A. (2004). Are loss functions all the same? *Neural Computation*.
- Scholkopf, B. and Smola, A. (2001). *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT press.
- Seeger, M. (2001). Learning with labeled and unlabeled data. Technical report, Institute for Adaptive and Neural Computation.
- Steinwart, I. (2007). How to compare different loss functions and their risks. *Constructive Approximation*.
- Steinwart, I. and Christmann, A. (2008). *Support Vector Machines*. Springer Science & Business Media.
- Steinwart, I. and Scovel, C. (2007). Fast rates for support vector machines using Gaussian kernels. *The Annals of Statistics*.
- Sun, H.-W. and Zhou, D.-X. (2008). Reproducing kernel hilbert spaces associated with analytic translation-invariant mercer kernels. *Journal of Fourier Analysis and Applications*.
- Valiant, L. (2013). *Probably Approximately Correct*. Basic Books.
- Vapnik, V. (1995). *The Nature of Statistical Learning Theory*. Springer-Verlag.
- Vigogna, S., Meanti, G., De Vito, E., and Rosasco, L. (2022). Multiclass learning with margin: exponential rates with no bias-variance trade-off. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pages 22260–22269. PMLR.
- Vovk, V. and Shafer, G. (2008). A tutorial on conformal prediction. *Journal of Machine Learning Research*.
- Wilbraham, H. (1848). On a certain periodic function. *The Cambridge and Dublin Mathematical Journal*.
- Yang, G. et al. (2021). Tensor programs V: Tuning large neural networks via zero-shot hyperparameter transfer. In *Advances in Neural Information Processing Systems*.
- Zhang, T. (2004). Statistical behavior and consistency of classification methods based on convex risk minimization. *Annals of Statistics*.

A PROOFS

In this section we provide the proofs of Lemma 1, Lemma 2, Proposition 1, Theorem 1 and Proposition 2.

A.1 Proof of Lemma 1

The first part follows by integration of a pointwise result. Consider the function $h_p : \mathbb{R} \rightarrow \mathbb{R}; q \mapsto p(1-q)_+ + (1-p)(1+q)_+$, where $p \in (0, 1)$ represents $\mathbb{P}(Y = 1|X)$ and q represents $g(x)$. The function h_p has a slope equal to $-p$ for $q < -1$, then slope $1 - 2p$ for $q \in (-1, 1)$, and $1 - p$ for $q > 1$. Therefore, when $q_1, q_2 \in (-1, 1)$, we have

$$h_p(q_2) - h_p(q_1) = (1 - 2p)(q_2 - q_1).$$

Taking $p = \mathbb{P}(Y = 1|X)$, $q_2 = g_2(X)$ and $q_1 = g_1(X)$, we get $1 - 2p = -\mathbb{E}[Y|X] = -\eta(X)$. By integration, we obtain the claim. From the previous slope considerations, it also follows that h_p is minimized by $q = \text{sign}(2p - 1)$, meaning that one can take $g^*(X) = \text{sign } \eta(X)$.

More exactly, the slopes reasoning shows that: for $x \notin \text{supp } \rho_{\mathcal{X}}$, $g^*(x)$ and $f^*(x)$ can be arbitrarily chosen; for $x \in \text{supp } \rho_{\mathcal{X}}$ such that $\eta(x) = 0$, $g^*(x)$ and $f^*(x)$ can be arbitrarily chosen in $[-1, 1]$ and $\{-1, 1\}$; for $x \in \text{supp } \rho_{\mathcal{X}}$ with $\eta(x) = 1$, $f^*(x) = 1$ and $g^*(x)$ can be arbitrarily chosen in $[1, +\infty)$; finally, for $x \in \text{supp } \rho_{\mathcal{X}}$ and $\eta(x) \in (0, 1)$, $g^*(x) = f^*(x) = 1$.

The second part follows from the fact that projecting on $[-1, 1]$ can only reduce the value of the hinge loss, that $\eta(x)(\pi(g)(x) - g^*(x))$ is always negative, and the reverse Hölder inequality:

$$\mathcal{R}_S(g) - \mathcal{R}(g^*) \geq \mathcal{R}_S(\pi(g)) - \mathcal{R}(g^*) = \|\eta(\pi(g) - g^*)\|_1 \geq \|\pi(g) - g^*\|_q \|\eta\|_p^{-1}.$$

A Hölder inequality also holds for weak Lebesgue spaces (see Castillo and Rafeiro, 2016, Theorem 5.23), whence

$$\mathcal{R}_S(g) - \mathcal{R}(g^*) \geq \|\eta(\pi(g) - g^*)\|_1 \geq \|\eta(\pi(g) - g^*)\|_{1,\infty} \geq \frac{1}{2} \|\eta\|_{p,\infty}^{-1} \|\pi(g) - g^*\|_{\frac{p}{p+1},\infty}.$$

This completes the proof.

A.2 Proof of Lemma 2

Assume without restrictions that there exists $x \in \mathcal{X}_1$ such that $|g^*(x) - g(x)| \geq 1$. For any event $A = A(X)$, by the law of total probability we have

$$\mathbb{P}(A) = \rho_{\mathcal{X}}(\mathcal{X}_1) \mathbb{P}(A | X \in \mathcal{X}_1) + \rho_{\mathcal{X}}(\mathcal{X}_{-1}) \mathbb{P}(A | X \in \mathcal{X}_{-1}) \geq \rho_{\mathcal{X}}(\mathcal{X}_1) \mathbb{P}(A | X \in \mathcal{X}_1).$$

Hence, since $g^*(\mathcal{X}_1) = \{1\}$,

$$\|g - g^*\|_{q,\infty}^q = \sup_{t>0} t^q \mathbb{P}(|g(X) - g^*(X)| > t) \geq \sup_{t>0} t^q \mathbb{P}(|g(X) - 1| > t | X \in \mathcal{X}_1) \rho_{\mathcal{X}}(\mathcal{X}_1).$$

Using the triangular inequality, the G -Lipschitz continuity of g , and the definition of x , we have that, for any $x' \in \mathcal{X}$,

$$|g(x') - 1| \geq |g(x) - 1| - |g(x') - g(x)| \geq 1 - Gd(x, x').$$

As a consequence,

$$\mathbb{P}(|g(X) - 1| > t | X \in \mathcal{X}_1) \geq \mathbb{P}\left(X \in \mathcal{B}\left(x, \frac{1-t}{G}\right) | X \in \mathcal{X}_1\right) = \frac{\rho_{\mathcal{X}}\left(\mathcal{X}_1 \cap \mathcal{B}\left(x, \frac{1-t}{G}\right)\right)}{\rho_{\mathcal{X}}(\mathcal{X}_1)}.$$

Combined with the previous facts, we get

$$\|g - g^*\|_{q,\infty}^q \geq \sup_{t>0} t^q \rho_{\mathcal{X}}\left(\mathcal{X}_1 \cap \mathcal{B}\left(x, \frac{1-t}{G}\right)\right).$$

Thanks to Assumption 3, there exists (c, r, d) such that (15) holds for \mathcal{X}_1 . Hence, when $G^{-1} < r$, we get the following lower bound:

$$\|g - g^*\|_{q,\infty}^q \geq cG^{-d} \sup_{t \in [0,1]} t^q (1-t)^d = cG^{-d} \frac{q^q d^d}{(d+q)^{d+q}}.$$

This proves the statement in the lemma.

A.3 Proof of Proposition 1

Suppose $\mathcal{R}(\text{sign } g) > \mathcal{R}(f^*)$. Then, observing that $\text{sign}(\pi(t)) = \text{sign}(t)$ for all $t \in \mathbb{R}$, and taking $g^* = f^*$, we know there must be $x \in \text{supp } \rho_{\mathcal{X}}$ such that $|\pi(g(x)) - g^*(x)| \geq 1$. Hence, by Lemma 2, we get $\|\pi(g) - g^*\|_{q,\infty} \geq c_0 G^{-\frac{d}{q}}$, and therefore, by Lemma 1, $\mathcal{R}_S(g) - \mathcal{R}_S(g^*) \geq 2^{-1} c_0 G^{-\frac{d}{q}} \|\eta\|^{-1}_{p,\infty}$. Thus, the proposition is proved.

A.4 Proof of Theorem 1

From Proposition 1 and Assumption 4, we get, with $\tilde{L} = \max\{MG_\varphi\sigma^{-1}, r^{-1}\}$ and $q = p/p+1$, and $g_{M,\sigma}$ the minimizer of \mathcal{R}_S insider $G_{M,\sigma}$,

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_n}[\mathcal{R}(\text{sign } g_{\mathcal{D}_n})] - \mathcal{R}(f^*) &\leq \mathbb{P}_{\mathcal{D}_n}\left(\mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_S(g^*) \geq 2^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}}\right) \\ &\leq \mathbb{P}_{\mathcal{D}_n}\left(\mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_S(g_{M,\sigma}) \geq 4^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}}\right). \end{aligned}$$

To deal with this last quantity, we proceed by using the fact that

$$\mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g_{\mathcal{D}_n}) \leq \mathcal{R}_{S,\mathcal{D}_n}(g_{\mathcal{D}_n}) \leq \mathcal{R}_{S,\mathcal{D}_n}(g_{M,\sigma}),$$

where $\mathcal{R}_{S,\mathcal{D}_n}$ denotes the empirical surrogate risk, to deduce that

$$\mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_S(g_{M,\sigma}) \leq \mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g_{\mathcal{D}_n}) + \mathcal{R}_{S,\mathcal{D}_n}(g_{M,\sigma}) + \mathcal{R}_S(g_{M,\sigma}).$$

Hence, we get the following union bound

$$\begin{aligned} \mathbb{E}_{\mathcal{D}_n}[\mathcal{R}(\text{sign } g_{\mathcal{D}_n})] - \mathcal{R}(f^*) &\leq \mathbb{P}_{\mathcal{D}_n}\left(\mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g_{\mathcal{D}_n}) \geq 8^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}}\right) \\ &\quad + \mathbb{P}_{\mathcal{D}_n}\left(\mathcal{R}_S(g_{M,\sigma}) - \mathcal{R}_{S,\mathcal{D}_n}(g_{M,\sigma}) \geq 8^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}}\right). \end{aligned}$$

Regarding the first term, we can reuse the literature on Rademacher complexity for linear models on convex risks (Bartlett and Mendelson, 2002), as well as the contraction principle in order to add π which is 1-Lipschitz Maurer (2016), which ensures that

$$\mathbb{E}_{\mathcal{D}_n} \left[\sup_{g \in \mathcal{G}_{M,\sigma}} |\mathcal{R}_S(\pi \circ g) - \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g)| \right] \leq M \|\varphi\|_\infty n^{-1/2}.$$

Note that Assumption 3 implies that $\text{supp } \rho_{\mathcal{X}}$ is compact, hence, if φ is Lipschitz-continuous, it is bounded on $\text{supp } \rho_{\mathcal{X}}$. This allows us to use McDiarmid inequality to get the same type of bound on the deviation of $\mathcal{R}_S(g_{\mathcal{D}_n})$ around its mean. Let $H(\mathcal{D}_n) = \sup_{g \in \mathcal{G}_{M,\sigma}} \mathcal{R}_S(\pi \circ g) - \mathcal{R}_{\mathcal{D}_n}(\pi \circ g)$. Let us decompose $\mathcal{D}_n = ((x_1, y_1), \dots, (x_n, y_n))$. We would like to show that if \mathcal{D}'_n is equal to \mathcal{D}_n for each datapoint but for (x_i, y_i) that becomes (x'_i, y'_i) then $H(\mathcal{D}_n) - H(\mathcal{D}'_n)$ is bounded. We have

$$\begin{aligned} H(\mathcal{D}_n) - H(\mathcal{D}'_n) &= \sup_{g \in \mathcal{G}_{M,\sigma}} \mathcal{R}_S(\pi \circ g) - \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g) - \sup_{g' \in \mathcal{G}_{M,\sigma}} \mathcal{R}_S(\pi \circ g') - \mathcal{R}_{S,\mathcal{D}'_n}(\pi \circ g') \\ &\leq \sup_{g \in \mathcal{G}_{M,\sigma}} \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g) - \mathcal{R}_{S,\mathcal{D}'_n}(\pi \circ g) \\ &= n^{-1} \sup_{g \in \mathcal{G}_{M,\sigma}} L(\pi \circ g(x'_i), y'_i) - L(\pi \circ g(x_i), y_i) \leq n^{-1} \end{aligned}$$

Using McDiarmid's inequality, we get

$$\mathbb{P}(H(\mathcal{D}_n) - \mathbb{E}[H(\mathcal{D}_n)] \geq t) \leq \exp(-2nt^2).$$

In other terms, when adding the control we have on the expectation, we get

$$\mathbb{P}_{\mathcal{D}_n} \left(\sup_{g \in \mathcal{G}_{M,\sigma}} \mathcal{R}_S(\pi \circ g) - \mathcal{R}_{S,\mathcal{D}_n}(\pi \circ g) > t + M \|\varphi\|_\infty n^{-1/2} \right) \leq \exp(-2nt^2). \quad (24)$$

When $8^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}} \geq \|\varphi\|_\infty M n^{-1/2}$, this leads to

$$\mathbb{P}_{\mathcal{D}_n} \left(\mathcal{R}_S(\pi \circ g_{\mathcal{D}_n}) - \mathcal{R}_{S,\mathcal{D}_n}(g_{\mathcal{D}_n}) \geq 8^{-1} \|\eta\|^{-1}_{p,\infty} c_0 \tilde{L}^{-\frac{d}{q}} \right)$$

$$\begin{aligned}
&\leq \exp\left(-\frac{n}{8}\left(8^{-1}\|\eta\|^{-1}\|_{p,\infty}^{-1}c_0\tilde{L}^{-\frac{d}{q}}-M\|\varphi\|_\infty n^{-1/2}\right)^2\right) \\
&\leq \exp\left(-\frac{c_0^2\sigma^{\frac{2d(p+1)}{p}}}{512\|\eta\|^{-1}\|_{p,\infty}^2(MG_\varphi)^{\frac{2d(p+1)}{p}}}\cdot n + \frac{c_0\|\varphi\|_\infty}{32\|\eta\|^{-1}\|_{p,\infty}M^{\frac{d(p+1)}{p}-1}G_\varphi^{\frac{d(p+1)}{p}}}\cdot n^{-1/2} - \frac{M^2\|\varphi\|_\infty^2}{8}\right).
\end{aligned}$$

Regarding the second term, we can use the classical concentration of $\mathcal{R}_{\mathcal{D}_n}(g_{M,\sigma})$ around its mean. For example, using the fact that Assumption 3 implies that $\rho_{\mathcal{X}}$ is compact, and using the fact that L and $g_{\sigma,M}$ are Lipschitz, we deduce that $L(g_{\sigma,M}, Y)$ is bounded, hence one can apply Hoeffding's inequality to get the same type of exponential control on this term.

The result follows from those concentration inequalities and the fact that \mathcal{R} is bounded by one, since any function $h : \mathbb{N}^* \rightarrow \mathbb{R}; n \rightarrow \min(1, a \exp(-bn))$ given two constants $a, b > 0$ can be bounded by $h' : \mathbb{N}^* \rightarrow \mathbb{R}; n \rightarrow 2 \exp(-cn)$ for a constant $c > 0$.

A.5 Proof of Proposition 2

Since the Hinge loss is 1-Lipschitz from the proof of Lemma 1, we have that, for any function $g : \mathcal{X} \rightarrow \mathbb{R}$,

$$\mathcal{R}_S(g) - \mathcal{R}_S(g^*) \leq \|g - g^*\|_{L^1(\rho_{\mathcal{X}})} \leq \|g - g^*\|_{L^2(\rho_{\mathcal{X}})}.$$

To verify Assumption 4, it is sufficient to prove that there exists a $\sigma > 0$ such that for any $c, r > 0$, there exists a $M > 0$ and a function $g \in \mathcal{G}_{M,\sigma}$ such that $\|g - g^*\|_1 \leq cM^{-r}$.

Under Assumption 5, g^* can be taken as a non-analytic smooth function, *e.g.*

$$g^*(x) = \frac{\exp(-d(x, \mathcal{X}_{-1})) - \exp(-d(x, \mathcal{X}_1))}{\exp(-d(x, \mathcal{X}_{-1})) + \exp(-d(x, \mathcal{X}_1))}.$$

In particular, it can be taken as belonging to any Sobolev space. On the one hand, any kernel that can be written as $k(x, y) = \varphi(\|x - y\|^2)$ for some function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ only contains analytic functions in the resulting function space $\mathcal{G} = \cup_{M,\sigma} \mathcal{G}_{M,\sigma}$ (Sun and Zhou, 2008), hence some extra work is needed to show that g^* can be approximated well enough from those spaces: more precisely, we need to show that the approximation error within $\mathcal{G}_{M,\sigma}$ decays faster than M^{-r} . On the other hand, as soon as $m \geq (d+1)/2$, the Sobolev space $H^m(\mathbb{R}^d) = W^{m,2}(\mathbb{R}^d)$ is a reproducing kernel Hilbert space. In particular, the case $m = (d+1)/2$ is associated with the exponential kernel $k(x, y) = \exp(-\|x - y\|)$ (see section 7.3.3 in Bach, 2023, for example), hence one can consider $g^* \in \mathcal{G}_{M,\sigma}$ for M big enough, which trivially implies Assumption 4.

A.6 Proof of Corollary 1

This theorem is a direct application of strong duality. Because $\theta \rightarrow \mathcal{R}(g_\theta)$ is convex and $\|\theta\| \leq M$ has a non-empty relative interior, strong duality holds and

$$\begin{aligned}
\inf_{\theta; \|\theta\| \leq M} \mathcal{R}_S(g_\theta) &= \sup_{\lambda \geq 0} \inf_{\theta} \mathcal{R}_S(g_\theta) + \lambda(\|\theta\|^2 - M^2) \\
&= \mathcal{R}_S(g_{\theta^*}) + \lambda^*(\|\theta^*\|^2 - M^2) \\
&= \inf_{\theta} \mathcal{R}_S(g_\theta) + \lambda^* \|\theta\|^2 - \lambda^* M^2.
\end{aligned}$$

where (θ^*, λ^*) is the solution of the Lagrangian problem. In other terms, for any $M > 0$, there exists a $\lambda \geq 0$ such that $\theta_\lambda = \theta_M$, where θ_λ is the minimizer of the regularized problem, and θ_M is the minimizer of the constrained problem.

Similarly to the proof of Theorem 1, one can prove concentration inequality on $\mathcal{R}_S(g_{\mathcal{D}_n,\lambda}) - \mathcal{R}_S(g_{\theta_\lambda})$, which can then be translated into exponential convergence rates on the original problem.

B ADDITIONAL CONTEXT ON SVM

In this section, we review the geometrical motivation behind support vector machines, as well as their hinge loss characterization. Suppose that we are given data $(x_i, y_i)_{i \leq n}$. We would like to find a linear separating hyperplane in the

features space $\mathcal{H} \supset \varphi(\mathcal{X})$ between the points $\{\varphi(x_i) \mid y_i = 1\}$ and $\{\varphi(x_i) \mid y_i = -1\}$. This can be formulated as

$$\begin{aligned} \text{find } & \theta \\ \text{s.t. } & y_i \theta^\top \varphi(x_i) > 0 \quad \forall i \leq n. \end{aligned}$$

For a feasible θ such that $\|\theta\| = 1$, one can compute the margin that separates the positive and negative labeled points along the θ -axis. It reads $\min_{y_i=1, y_j=-1} \theta^\top (\varphi(x_i) - \varphi(x_j))$. One can also compute the minimal displacement that would make a point change of assigned label according to the classification rule induced by θ , it reads $\min y_i \theta^\top \varphi(x_i)$. Hence, among the feasible θ , the most robust to point displacement, is defined through the maximization of the margin between points and the origin along the θ -axis.

$$\begin{aligned} \text{max } & c \\ \text{s.t. } & y_i \theta^\top \varphi(x_i) > c \|\theta\| \quad \forall i \leq n. \end{aligned}$$

Of course, as soon as there is noise, or if the model is not well-specified, this maximization problem is infeasible. One way to overcome this is to introduce slack variables that act as budget for points that are too close, or on the wrong side of the separating hyperplane,

$$\begin{aligned} \text{max } & \sum_{i \leq n} \xi_i \\ \text{s.t. } & \xi_i < 0 \quad \forall i \leq n \\ & y_i \theta^\top \varphi(x_i) > 1 + \xi_i \quad \forall i \leq n. \end{aligned}$$

This maximization problem can be rewritten as the minimization problem

$$\arg \min_{\theta} \sum_{i \leq n} \max \{0, 1 - y_i \theta^\top \varphi(x_i)\}.$$

This is exactly the empirical risk minimization of the hinge with the linear class of function considered in this paper. Indeed, our work completely forgets about the geometrical point of view of SVM, it uses the classical framework of statistical learning, where a variational objective is provided by a loss ℓ and the minimization is done over functions from inputs to outputs. Interestingly, the maximum margin principle, which was crucial in the introduction of SVM (Vapnik, 1995), reappears in Assumption 5 under a weaker form: we do not ask for a clear margin between points that have different labels, but for a margin between points where the optimal classifier should be positive and the ones where it should be negative. The subtle difference resides in the fact that we allow for labeling noise. In particular, we allow for much more labeling noise than previous works that have only shown exponential convergence rates under the hard low-noise condition.

C EXPERIMENTAL DETAILS

In our experiments, we used the SVM implementation of Chang and Lin (2011) through its *Scikit-learn* wrapper (Pedregosa et al., 2011) in *Python*. We used *Numpy* (Harris et al., 2020) to reduce our work to high-level array instructions, and *Matplotlib* for visualization (Hunter, 2007). Randomness in experiments was controlled with the random seed provided by *Numpy*, which we initialized at zero.

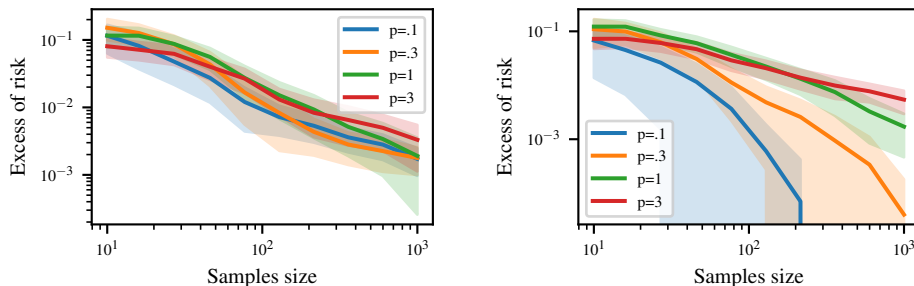


Figure 5: (Left) Similar setting as Figure 2 but with X uniform on $[-1, 1]$. The behavior of the excess of risk is quite different without the separation in \mathcal{X} : no exponential convergence rate is kicking in after a thousand of samples. (Right) Similar setting as Figure 2, using kernel ridge regression with the least-squares surrogate. Exponential convergence rates are observed with a slight delay compared to the hinge loss, and are explained by the hard low-noise condition 1.

Figures 2 and 5 are derived by averaging 100 trials of the following procedure. We draw uniformly at random n independent samples uniformly distributed on $\mathcal{X} \in \{[-1, 1], [-1, -1] \cup [1, 1]\}$. We draw randomly one output y_i for each input x_i ,

according to $\eta(x_i)$. We consider the Gaussian kernel $k(x, x') = \exp(-\|x - x'\|^2 / 2\sigma^2)$ for $\sigma = .2$, and solve the empirical risk minimization associated to the hinge loss with the penalization $\lambda \|\theta\|^2$ (rather than the hard constraint $\|\theta\| < M$) for $\lambda = 10^{-4}$. The generalization error is measured through the formula $\mathbb{E}[\|\eta(x)\| \mathbf{1}_{f(x) \neq f^*(x)}]$, with an empirical approximation of this sum with the points $(x_i)_{i \leq n}$ chosen such that $\rho_{\mathcal{X}}([x_i, x_{i+1}]) = 1/n$ and $\rho_{\mathcal{X}}([x_n, +\infty)) = 1/n$, with $n = 10^4$ (which makes sure that the exponential behavior observed is not due to the lack of testing samples). For each x , the height of each dark part corresponds to one standard deviation of the generalization error computed from the 100 trials, and the solid line corresponds to the empirical average. The fact that the dark parts are not centered around the averages is due to the fact that we have drawn log-plots but centered the interval for linearly-scaled plots.

Figure 3 is obtained by considering $\mathcal{X} = [0, 1]^2$ with uniform input distribution, the Gaussian kernel with $\sigma = .2$, and the penalty parameter $\lambda = 10^{-3}$ (instead of a hard constraint leading to a parameter M as in the main text derivations). We take $n = 10^4 = 100^2$ points uniformly spread out on \mathcal{X} (on the regular lattice $\frac{1}{\sqrt{n}} \cdot \mathbb{Z}^2 \cap \mathcal{X}$) to approximate $g_{\lambda, \sigma}$ with empirical risk minimization on this curated dataset. We consider $\eta(x) = \pi_{[-1, 1]}(2x_2 - .5 \sin(2\pi x_1) - 1)$, and assign to each x in the dataset a sample $(x, 1)$ weighted by $\mathbb{P}(Y = 1 | X = x) = (\eta(x) - 1)/2$, and a sample $(x, -1)$, weighted by $\mathbb{P}(Y = -1 | X = x)$. The “noiseless” setting denotes the setting where $(Y | X)$ is deterministic, but with the same decision frontier between the classes \mathcal{X}_1 and \mathcal{X}_{-1} characterized by $\{(x, .5 + .25 \sin(2\pi x)) | x \in [0, 1]\}$. Once we fit the support vector machine with this dataset, we test it with $n = 2.5 \cdot 10^5 = 500^2$ data points uniformly spread out on \mathcal{X} , and use *Matplotlib* to automatically draw level lines.

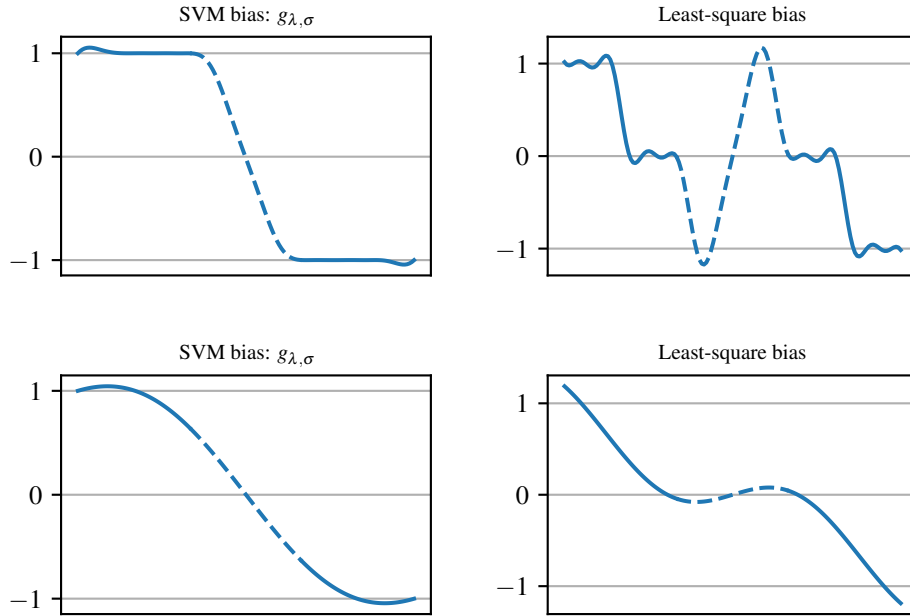


Figure 6: Same setting as Figure 4, with $\sigma = .2$ and $\lambda = 10^{-6}$ (top), and with $\sigma = 1$ and $\lambda = 10^{-3}$ (bottom).

Figures 4 and 6 correspond to $\mathcal{X} = [0, 3]$ with the input distribution uniform on $[0, 1] \cup [2, 3]$. Figure 4 is obtained with $\sigma = .1$ and $\lambda = 10^{-6}$. We derive it by considering $n = 100$ points uniformly spread out on the domain of η , solving the equivalent curated empirical risk minimization, that approximates both

$$g_{\lambda, \sigma} = \arg \min_{g: \mathcal{X} \rightarrow \mathbb{R}} \mathbb{E}_{\rho} \left[\left(1 - Y \left\langle \theta, \varphi \left(\frac{x}{\sigma} \right) \right\rangle \right)_+ \right] + \lambda \|\theta\|^2, \quad (25)$$

$$g_{(\text{LS})} = \arg \min_{g: \mathcal{X} \rightarrow \mathbb{R}} \mathbb{E}_{\rho} \left[\left| \left\langle \theta, \varphi \left(\frac{x}{\sigma} \right) \right\rangle - Y \right|^2 \right] + \lambda \|\theta\|^2. \quad (26)$$

The robustness of SVM might be understood from its geometrical definition: when trying to find the maximum separating margin, infinitesimal modifications that change the regularity properties of η do not really matter. The picture is different for the least-squares surrogate with kernel methods, where from few point evaluations, the system reconstructs a function by assuming regularity and inferring information on high-order derivatives. This is similar to the Runge phenomenon with

Hermite interpolation. More precisely, the Gaussian kernel is linked to a space of functions with rapidly decreasing Fourier coefficients (see, for example, Bach, 2023, for a more precise link). The function η that needs to be approximated on Figure 4 is similar to the Heaviside function, whose Fourier coefficients are of the form $(\frac{1}{i\pi k})_{k \in \mathbb{N}^*}$ and do not decrease fast enough to be all reconstructed. This leads to some high-frequency oscillations missing in the reconstruction as it appears on Figure 4.