
Sparsity of SVMs that use the ϵ -insensitive loss

Ingo Steinwart
Information Sciences Group CCS-3
Los Alamos National Laboratory
Los Alamos, NM 87545, USA
ingo@lanl.gov

Andreas Christmann
University of Bayreuth
Department of Mathematics
D-95440 Bayreuth
Andreas.Christmann@uni-bayreuth.de

Abstract

In this paper lower and upper bounds for the number of support vectors are derived for support vector machines (SVMs) based on the ϵ -insensitive loss function. It turns out that these bounds are asymptotically tight under mild assumptions on the data generating distribution. Finally, we briefly discuss a trade-off in ϵ between sparsity and accuracy if the SVM is used to estimate the conditional median.

1 Introduction

Given a reproducing kernel Hilbert space (RKHS) of a kernel $k : X \times X \rightarrow \mathbb{R}$ and training set $D := ((x_1, y_1), \dots, (x_n, y_n)) \in (X \times \mathbb{R})^n$, the ϵ -insensitive SVM proposed by Vapnik and his co-workers [10, 11] for regression tasks finds the unique minimizer $f_{D,\lambda} \in H$ of the regularized empirical risk

$$\lambda \|f\|_H^2 + \frac{1}{n} \sum_{i=1}^n L_\epsilon(y_i, f(x_i)), \quad (1)$$

where L_ϵ denotes the ϵ -insensitive loss defined by $L_\epsilon(y, t) := \max\{0, |y - t| - \epsilon\}$ for all $y, t \in \mathbb{R}$ and some fixed $\epsilon \geq 0$. It is well known, see e.g. [2, Proposition 6.21], that the solution is of the form

$$f_{D,\lambda} = \sum_{i=1}^n \beta_i^* k(x_i, \cdot), \quad (2)$$

where the coefficients β_i^* are a solution of the optimization problem

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^n y_i \beta_i - \epsilon \sum_{i=1}^n |\beta_i| - \frac{1}{2} \sum_{i,j=1}^n \beta_i \beta_j k(x_i, x_j) \end{aligned} \quad (3)$$

$$\text{subject to} \quad -C \leq \beta_i \leq C \quad \text{for all } i = 1, \dots, n. \quad (4)$$

Here we set $C := 1/(2\lambda n)$. Note that the equality constraint $\sum_{i=1}^n \beta_i = 0$ needed in [2, Proposition 6.21] is superfluous since we do not include an offset term b in the primal problem (1). In the following, we write $SV(f_{D,\lambda}) := \{i : \beta_i^* \neq 0\}$ for the set of indices that belong to the support vectors of $f_{D,\lambda}$. Furthermore, we write $\#$ for the counting measure, and hence $\#SV(f_{D,\lambda})$ denotes the number of support vectors of $f_{D,\lambda}$.

It is obvious from (2) that $\#SV(f_{D,\lambda})$ has a crucial influence on the time needed to compute $f_{D,\lambda}(x)$. Due to this fact, the ϵ -insensitive loss was originally motivated by the goal to achieve *sparse* decision functions, i.e., decision functions $f_{D,\lambda}$ with $\#SV(f_{D,\lambda}) < n$. Although empirically it is well-known that the ϵ -insensitive SVM achieves this sparsity, there is, so far, no theoretical explanation in the sense of [5]. The goal of this work is to provide such an explanation by establishing asymptotically tight lower and upper bounds for the number of support vectors. Based on these bounds we then investigate the trade-off between sparsity and estimation accuracy of the ϵ -insensitive SVM.

2 Main results

Before we can formulate our main results we need to introduce some more notations. To this end, let P be a probability measure on $X \times \mathbb{R}$, where X is some measurable space. Given a measurable $f : X \rightarrow \mathbb{R}$, we then define the L_ϵ -risk of f by $\mathcal{R}_{L_\epsilon, P}(f) := \mathbb{E}_{(x,y) \sim P} L_\epsilon(y, f(x))$. Moreover, recall that P can be split into the marginal distribution P_X on X and the regular conditional probability $P(\cdot | x)$. Given a RKHS H of a bounded kernel k , [1] then showed that

$$f_{P,\lambda} := \arg \inf_{f \in H} \lambda \|f\|_H^2 + \mathcal{R}_{L_\epsilon, P}(f)$$

exists and is uniquely determined whenever $\mathcal{R}_{L_\epsilon, P}(0) < \infty$. Let us write $\delta_{(x,y)}$ for the Dirac measure at some $(x, y) \in X \times \mathbb{R}$. By considering the empirical measure $D := \frac{1}{n} \sum_{i=1}^n \delta_{(x_i, y_i)}$ of a training set $D := ((x_1, y_1), \dots, (x_n, y_n)) \in (X \times \mathbb{R})^n$, we then see that the corresponding $f_{D,\lambda}$ is the solution of (1). Finally, we need to introduce the sets

$$\begin{aligned} A_{\text{low}}^\delta(f) &:= \{(x, y) \in X \times \mathbb{R} : |f(x) - y| > \epsilon + \delta\} \\ A_{\text{up}}^\delta(f) &:= \{(x, y) \in X \times \mathbb{R} : |f(x) - y| \geq \epsilon - \delta\}, \end{aligned}$$

where $f : X \rightarrow \mathbb{R}$ is an arbitrary function and $\delta \in \mathbb{R}$. Moreover, we use the short forms $A_{\text{low}}(f) := A_{\text{low}}^0(f)$ and $A_{\text{up}}(f) := A_{\text{up}}^0(f)$. Now we can formulate our first main result.

Theorem 2.1 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Then, for all $n \geq 1$, $\rho > 0$, $\delta > 0$, and $\lambda > 0$ satisfying $\delta\lambda \leq 4$, we have*

$$P^n \left(D \in (X \times \mathbb{R})^n : \frac{\#SV(f_{D,\lambda})}{n} > P(A_{\text{low}}^\delta(f_{P,\lambda})) - \rho \right) \geq 1 - 3e^{-\frac{\delta^2 \lambda^2 n}{16}} - e^{-2\rho^2 n}$$

and

$$P^n \left(D \in (X \times \mathbb{R})^n : \frac{\#SV(f_{D,\lambda})}{n} < P(A_{\text{up}}^\delta(f_{P,\lambda})) + \rho \right) \geq 1 - 3e^{-\frac{\delta^2 \lambda^2 n}{16}} - e^{-2\rho^2 n}.$$

Before we present our second main result, we briefly illustrate Theorem 2.1 for the case where we fix the regularization parameter λ and let $n \rightarrow \infty$.

Corollary 2.2 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Then, for all $\rho > 0$ and $\lambda > 0$, we have*

$$\lim_{n \rightarrow \infty} P^n \left(D \in (X \times \mathbb{R})^n : P(A_{\text{low}}(f_{P,\lambda})) - \rho \leq \frac{\#SV(f_{D,\lambda})}{n} \leq P(A_{\text{up}}(f_{P,\lambda})) + \rho \right) = 1.$$

Note that the above corollary exactly describes the asymptotic behavior of the fraction of support vectors modulo the probability of the set

$$A_{\text{up}}(f_{P,\lambda}) \setminus A_{\text{low}}(f_{P,\lambda}) = \{(x, f_{P,\lambda}(x) - \epsilon) : x \in X\} \cup \{(x, f_{P,\lambda}(x) + \epsilon) : x \in X\}.$$

In particular, if the conditional distributions $P(\cdot | x)$, $x \in X$, have no discrete components, then the above corollary gives an exact description.

Of course, in almost no situation it is realistic to assume that λ stays fixed if the sample size n grows. Instead, it is well-known, see [1], that the regularization parameter should vanish in order to achieve consistency. To investigate this case, we need to introduce some additional notations from [6] that are related to the L_ϵ -risk. Let us begin by denoting the Bayes L_ϵ -risk by $\mathcal{R}_{L_\epsilon, P}^* := \inf \mathcal{R}_{L_\epsilon, P}(f)$, where P is a distribution and the infimum is taken over all measurable functions $f : X \rightarrow \mathbb{R}$. In addition, given a distribution Q on \mathbb{R} , [6] and [7, Chapter 3] defined the *inner L_ϵ -risks* by

$$\mathcal{C}_{L_\epsilon, Q}(t) := \int_{\mathbb{R}} L_\epsilon(y, t) dQ(y), \quad t \in \mathbb{R},$$

and the *minimal inner L_ϵ -risks* were denoted by $\mathcal{C}_{L_\epsilon, Q}^* := \inf_{t \in \mathbb{R}} \mathcal{C}_{L_\epsilon, Q}(t)$. Obviously, we have

$$\mathcal{R}_{L_\epsilon, P}(f) = \int_X \mathcal{C}_{L_\epsilon, P(\cdot | x)}(f(x)) dP_X(x), \quad (5)$$

and [6, Lemma 2.5], see also [7, Lemma 3.4], further established the intuitive formula $\mathcal{R}_{L_\epsilon, P}^* = \int_X \mathcal{C}_{L_\epsilon, P(\cdot | x)}^* dP_X(x)$. Moreover, we need the sets of *conditional minimizers*

$$\mathcal{M}^*(x) := \{t \in \mathbb{R} : \mathcal{C}_{L_\epsilon, P(\cdot | x)}(t) = \mathcal{C}_{L_\epsilon, P(\cdot | x)}^*\}.$$

The following lemma collects some useful properties of these sets.

Lemma 2.3 *Let P be a probability measure on $X \times \mathbb{R}$ with $\mathcal{R}_{L_\epsilon, P}^* < \infty$. Then $\mathcal{M}^*(x)$ is a non-empty and compact interval for P_X -almost all $x \in X$.*

Given a function $f : X \rightarrow \mathbb{R}$, Lemma 2.3 shows that for P_X -almost all $x \in X$ there exists a unique $t^*(x) \in \mathcal{M}^*(x)$ such that

$$|t^*(x) - f(x)| \leq |t - f(x)| \quad \text{for all } t \in \mathcal{M}^*(x). \quad (6)$$

In other words, $t^*(x)$ is the element in $\mathcal{M}^*(x)$ that has the smallest distance to $f(x)$. In the following, we sometimes write $t_\lambda^*(x) := t^*(x)$ if $f = f_{P, \lambda}$ and we wish to emphasize the dependence of $t^*(x)$ on λ . With the help of these elements, we finally introduce the sets

$$\begin{aligned} M_{\text{low}}^\delta(f) &:= \{(x, y) \in X \times \mathbb{R} : |t^*(x) - y| > \epsilon + \delta\} \\ M_{\text{up}}^\delta(f) &:= \{(x, y) \in X \times \mathbb{R} : |t^*(x) - y| \geq \epsilon - \delta\}, \end{aligned}$$

where $\delta \in \mathbb{R}$. Moreover, we again use the short forms $M_{\text{low}}(f) := M_{\text{low}}^0(f)$ and $M_{\text{up}}(f) := M_{\text{up}}^0(f)$. Now we can formulate our second main result.

Theorem 2.4 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Assume that $\mathcal{R}_{L_\epsilon, P}(0) < \infty$ and that H is dense in $L_1(P_X)$. Then, for all $\rho > 0$, there exist a $\delta_\rho > 0$ and a $\lambda_\rho > 0$ such that for all $\lambda \in (0, \lambda_\rho]$ and all $n \geq 1$ we have*

$$P^n\left(D \in (X \times \mathbb{R})^n : P(M_{\text{low}}(f_{P, \lambda})) - \rho \leq \frac{\#SV(f_{D, \lambda})}{n} \leq P(M_{\text{up}}(f_{P, \lambda})) + \rho\right) \geq 1 - 8e^{-\delta_\rho^2 \lambda^2 n}.$$

If we choose a sequence of regularization parameters λ_n such that $\lambda_n \rightarrow 0$ and $\lambda_n^2 n \rightarrow \infty$, then the resulting SVM is L_ϵ -risk consistent under the assumptions of Theorem 2.4, see [1]. For this case, the following obvious corollary of Theorem 2.4 establishes lower and upper bounds on the number of support vectors.

Corollary 2.5 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Assume that $\mathcal{R}_{L_\epsilon, P}(0) < \infty$ and that H is dense in $L_1(P_X)$. Furthermore, let $(\lambda_n) \in (0, \infty)$ be a sequence with $\lambda_n \rightarrow 0$ and $\lambda_n^2 n \rightarrow \infty$. Then, for all $\rho > 0$, the probability P^n of $D \in (X \times \mathbb{R})^n$ satisfying*

$$\liminf_{m \rightarrow \infty} P(M_{\text{low}}(f_{P, \lambda_m})) - \rho \leq \frac{\#SV(f_{D, \lambda_n})}{n} \leq \limsup_{m \rightarrow \infty} P(M_{\text{up}}(f_{P, \lambda_m})) + \rho$$

converges to 1 for $n \rightarrow \infty$.

In general, the probabilities of the sets $M_{\text{low}}(f_{P, \lambda})$ and $M_{\text{up}}(f_{P, \lambda})$ are hard to control since, e.g., for fixed $x \in X$ and $\lambda \rightarrow 0$ it seems difficult to show that $f_{P, \lambda}(x)$ is *not* “flipping” from the left hand side of $\mathcal{M}^*(x)$ to the right hand side. Indeed, for general $\mathcal{M}^*(x)$, such flipping would give different values $t_\lambda^*(x) \in \mathcal{M}^*(x)$ for $\lambda \rightarrow 0$, and hence would result in significantly different sets $M_{\text{low}}(f_{P, \lambda})$ and $M_{\text{up}}(f_{P, \lambda})$. As a consequence, it seems hard to show that, for probability measures P whose conditional distributions $P(\cdot | x)$, $x \in X$, have no discrete components, we always have

$$\liminf_{\lambda \rightarrow 0} P(M_{\text{low}}(f_{P, \lambda})) = \limsup_{\lambda \rightarrow 0} P(M_{\text{up}}(f_{P, \lambda})). \quad (7)$$

However, there are situations in which this equality can easily be established. For example, assume that the sets $\mathcal{M}^*(x)$ are P_X -almost surely singletons. In this case, $t_\lambda^*(x)$ is in fact *independent* of λ , and hence so are $M_{\text{low}}(f_{P, \lambda})$ and $M_{\text{up}}(f_{P, \lambda})$. Namely, in this case these sets contain the pairs (x, y)

for which y is *not* contained in the closed or open ϵ -tube around $\mathcal{M}^*(x)$, respectively. Consequently, (7) holds provided that the conditional distributions $P(\cdot|x)$, $x \in X$, have no discrete components, and hence Corollary 2.5 gives a tight bound on the number of support vectors. Moreover, if in this case we additionally assume $\epsilon = 0$, i.e., we consider the absolute loss, then we easily find $P(M_{\text{low}}(f_{P,\lambda})) = P(M_{\text{up}}(f_{P,\lambda})) = 1$, and hence Corollary 2.5 shows that the corresponding SVM does *not* tend to produce sparse decision functions. Finally, recall that for this specific loss function, $\mathcal{M}^*(x)$ equals the median of $P(\cdot|x)$, and hence $\mathcal{M}^*(x)$ is a singleton whenever the median of $P(\cdot|x)$ is unique.

Let us now illustrate Corollary 2.5 for $\epsilon > 0$. To this end, we assume in the following that the conditional distributions $P(\cdot|x)$ are *symmetric*, i.e., for P_X -almost all $x \in X$ there exists a *conditional center* $c(x) \in \mathbb{R}$ such that $P(c(x) + A|x) = P(c(x) - A|x)$ for all measurable $A \subset \mathbb{R}$. Note that by considering $A := [0, \infty)$ it is easy to see that $c(x)$ is a median of $P(\cdot|x)$. Furthermore, the assumption $\mathcal{R}_{L_\epsilon, P}(0) < \infty$ imposed in the results above ensures that the conditional mean $f_P^*(x) := \mathbb{E}(Y|x)$ of $P(\cdot|x)$ exists P_X -almost surely, and from this it is easy to conclude that $c(x) = f_P^*(x)$ for P_X -almost all $x \in X$. Moreover, from [8, Proposition 3.2 and Lemma 3.3] we immediately obtain the following lemma.

Lemma 2.6 *Let P be a probability measure on $X \times \mathbb{R}$ such that $\mathcal{R}_{L_\epsilon, P}(0) < \infty$. Assume that the conditional distributions $P(\cdot|x)$, $x \in X$, are symmetric and that for P_X -almost all $x \in X$ there exists a $\delta(x) > 0$ such that for all $\delta \in (0, \delta(x)]$ we have*

$$P(f_P^*(x) + [-\delta, \delta]|x) > 0, \quad (8)$$

$$P(f_P^*(x) + [\epsilon - \delta, \epsilon + \delta]|x) > 0. \quad (9)$$

Then, for P_X -almost all $x \in X$, we have $\mathcal{M}^(x) = \{f_P^*(x)\}$ and $f_P^*(x)$ equals P_X -almost surely the unique median of $P(\cdot|x)$.*

Obviously, condition (8) means that the conditional distributions have some mass around their median f_P^* , whereas (9) means that the conditional distributions have some mass around $f_P^* \pm \epsilon$. Moreover, [8] showed that under the assumptions of Lemma 2.6, the corresponding ϵ -insensitive SVM can be used to estimate the conditional median. Let us now illustrate how the value of ϵ influences both the accuracy of this estimate and the sparsity. To this end, let us assume for the sake of simplicity that the conditional distributions $P(\cdot|x)$ have *continuous* Lebesgue densities $p(\cdot|x) : \mathbb{R} \rightarrow [0, \infty)$. By the symmetry of the conditional distributions it is then easy to see that these densities are symmetric around $f_P^*(x)$. Now, it follows from the continuity of the densities, that (8) is satisfied if $p(f_P^*(x)|x) > 0$, whereas (9) is satisfied if $p(f_P^*(x) + \epsilon|x) > 0$. Let us first consider the case where the conditional distributions are equal modulo translations. In other words, we assume that there exists a continuous Lebesgue density $q : \mathbb{R} \rightarrow [0, \infty)$ which is symmetric around 0 such that for P_X -almost all $x \in X$ we have

$$q(y) = p(f_P^*(x) + y|x), \quad y \in \mathbb{R}.$$

Note that this assumption is essentially identical to a classical ‘‘signal plus noise’’ assumption. In the following we further assume that q is unimodal, i.e., q has its only local and global maximum at 0. From this we easily see that (8) is satisfied, and (9) is satisfied if $q(\epsilon) > 0$. By Lemma 2.6 and the discussion around (7) we then conclude that under the assumptions of Corollary 2.5 the fraction of support vectors asymptotically approaches $2Q([\epsilon, \infty))$, where Q is the probability measure defined by q . This confirms the intuition that *larger values of ϵ lead to sparser decision functions*. In particular, if $Q([\epsilon, \infty)) = 0$, the corresponding SVM produces *super sparse* decision functions, i.e., decision functions whose number of support vectors does *not* grow linearly in the sample size. However, not surprisingly, there is a price to be paid for this sparsity. Indeed, [8, Lemma 3.3] indicates that the size of $q(\epsilon)$ has a direct influence on the ability of $f_{D,\lambda}$ to estimate the conditional median f_P^* . Let us describe this in a little more detail. To this end, we first find by [8, Lemma 3.3] and the convexity of $t \mapsto \mathcal{C}_{L_\epsilon, Q}(t)$ that

$$\mathcal{C}_{L_\epsilon, Q}(t) - \mathcal{C}_{L_\epsilon, Q}^* \geq q(\epsilon) \cdot \begin{cases} t^2/2 & \text{if } t \in [0, \epsilon] \\ t\epsilon - \epsilon^2/2 & \text{if } t \geq \epsilon. \end{cases}$$

By a literal repetition of the proof of [8, Theorem 2.5] we then find the self-calibration inequality

$$\|f - f_P^*\|_{L_1(P_X)} \leq \sqrt{2/q(\epsilon)} \sqrt{\mathcal{R}_{L_\epsilon, P}(f) - \mathcal{R}_{L_\epsilon, P}^*}, \quad (10)$$

which holds for all $f : X \rightarrow \mathbb{R}$ with $\mathcal{R}_{L_\epsilon, P}(f) - \mathcal{R}_{L_\epsilon, P}^* \leq \epsilon^2/2$. Now, if we are in the situation of Corollary 2.5, then we know that $\mathcal{R}_{L_\epsilon, P}(f_{D, \lambda_n}) \rightarrow \mathcal{R}_{L_\epsilon, P}^*$ in probability for $n \rightarrow \infty$, and thus (10) shows that f_{D, λ_n} approximates the conditional median f_P^* with respect to the $L_1(P_X)$ -norm. However, the guarantee for this approximation becomes worse the smaller $q(\epsilon)$ becomes, i.e., the larger ϵ is. In other words, *the sparsity of the decision functions may be paid by less accurate estimates of the conditional median*. On the other hand, our results also show that *moderate values for ϵ can lead to both reasonable estimates of the conditional median and relatively sparse decision functions*. In this regard we further note that one can also use [8, Lemma 3.3] to establish self-calibration inequalities that measure the distance of f to f_P^* only up to ϵ . In this case, however, it is obvious that such self-calibration inequalities are worse the larger ϵ is, and hence the informal conclusions above remain unchanged.

Finally, we like to mention that, if the conditional distributions are not equal modulo translations, then the situation may become more involved. In particular, if we are in a situation with $p(f_P^*(x)|x) > 0$ and $p(f_P^*(x) + \epsilon|x) > 0$ but $\inf_x p(f_P^*(x)|x) = \inf_x p(f_P^*(x) + \epsilon|x) = 0$, self-calibration inequalities of the form (10) are in general impossible, and weaker self-calibration inequalities require additional assumptions on P . We refer to [8] where the case $\epsilon = 0$ is considered.

3 Proofs

Setting $C := \frac{1}{2\lambda n}$ and introducing slack variables, we can restate the optimization problem (1) as

$$\begin{aligned} \text{minimize} \quad & \frac{1}{2} \|f\|_H^2 + C \sum_{i=1}^n (\xi_i + \tilde{\xi}_i) \\ \text{subject to} \quad & f(x_i) - y_i \leq \epsilon + \xi_i, \\ & y_i - f(x_i) \leq \epsilon + \tilde{\xi}_i, \\ & \xi_i, \tilde{\xi}_i \geq 0 \quad \text{for all } i = 1, \dots, n. \end{aligned} \quad (11)$$

In the following we denote the (unique) solution of (11) by $(f^*, \xi^*, \tilde{\xi}^*)$, where we note that we have $f^* = f_{D, \lambda}$. It is well-known, see e.g. [2, p. 117], that the dual optimization problem of (11) is

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^n y_i (\tilde{\alpha}_i - \alpha_i) - \epsilon \sum_{i=1}^n (\tilde{\alpha}_i + \alpha_i) - \frac{1}{2} \sum_{i,j=1}^n (\tilde{\alpha}_i - \alpha_i)(\tilde{\alpha}_j - \alpha_j) k(x_i, x_j) \\ \text{subject to} \quad & 0 \leq \alpha_i, \tilde{\alpha}_i \leq C \quad \text{for all } i = 1, \dots, n, \end{aligned} \quad (12)$$

where k is the kernel of the RKHS H . Furthermore, if $(\alpha_1^*, \tilde{\alpha}_1^*, \dots, \alpha_n^*, \tilde{\alpha}_n^*)$ denotes a solution of (12), then we can recover the primal solution $(f^*, \xi^*, \tilde{\xi}^*)$ by

$$f^* = \sum_{i=1}^n (\tilde{\alpha}_i^* - \alpha_i^*) k(x_i, \cdot), \quad (13)$$

$$\xi_i^* = \max\{0, f^*(x_i) - y_i - \epsilon\}, \quad (14)$$

$$\tilde{\xi}_i^* = \max\{0, y_i - f^*(x_i) - \epsilon\}, \quad (15)$$

for all $i = 1, \dots, n$. Moreover, the Karush-Kuhn-Tucker conditions of (12) are

$$\alpha_i^* (f^*(x_i) - y_i - \epsilon - \xi_i^*) = 0, \quad (16)$$

$$\tilde{\alpha}_i^* (y_i - f^*(x_i) - \epsilon - \tilde{\xi}_i^*) = 0, \quad (17)$$

$$(\alpha_i^* - C) \xi_i^* = 0, \quad (18)$$

$$(\tilde{\alpha}_i^* - C) \tilde{\xi}_i^* = 0, \quad (19)$$

$$\xi_i^* \tilde{\xi}_i^* = 0, \quad (20)$$

$$\alpha_i^* \tilde{\alpha}_i^* = 0, \quad (21)$$

where $i = 1, \dots, n$. Finally, note that by setting $\beta_i := \tilde{\alpha}_i - \alpha_i$ the problem (12) can be simplified to (3), and consequently, a solution β^* of (3) is of the form $\beta^* = \tilde{\alpha}^* - \alpha^*$. The following simple lemma provides lower and upper bounds for the set of support vectors.

Lemma 3.1 *Using the above notations we have*

$$\{i : |f_{D,\lambda}(x_i) - y_i| > \epsilon\} \subset \{i : \beta_i^* \neq 0\} \subset \{i : |f_{D,\lambda}(x_i) - y_i| \geq \epsilon\}.$$

Proof: Let us first prove the inclusion on the left hand side. To this end, we begin by fixing an index i with $f_{D,\lambda}(x_i) - y_i > \epsilon$. By $f_{D,\lambda} = f^*$ and (14), we then find $\xi_i^* > 0$, and hence (18) implies $\alpha_i^* = C$. From (21) we conclude $\tilde{\alpha}_i^* = 0$ and hence we have $\beta_i^* = \tilde{\alpha}_i^* - \alpha_i^* = -C \neq 0$. The case $y_i - f_{D,\lambda}(x_i) > \epsilon$ can be shown analogously, and hence we obtain the first inclusion. In order to show the second inclusion we fix an index i with $\beta_i^* \neq 0$. By $\beta_i^* = \tilde{\alpha}_i^* - \alpha_i^*$ and (21) we then have either $\alpha_i^* \neq 0$ or $\tilde{\alpha}_i^* \neq 0$. Let us first consider the case $\alpha_i^* \neq 0$ and $\tilde{\alpha}_i^* = 0$. The KKT condition (16) together with $f_{D,\lambda} = f^*$ implies $f_{D,\lambda}(x_i) - y_i - \epsilon = \xi_i^*$ and since $\xi_i^* \geq 0$ we get $f_{D,\lambda}(x_i) - y_i \geq \epsilon$. The second case $\tilde{\alpha}_i^* = 0$ can be shown analogously. ■

We further need the following Hilbert space version of Hoeffding's inequality from [12, Chapter 3], see also [7, Chapter 6.2] for a slightly sharper inequality.

Theorem 3.2 *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space and H be a separable Hilbert space. Moreover, let $\eta_1, \dots, \eta_n : \Omega \rightarrow H$ be independent random variables satisfying $\mathbb{E}_{\mathbb{P}}\eta_i = 0$ and $\|\eta_i\|_{\infty} \leq 1$ for all $i = 1, \dots, n$. Then, for all $\tau \geq 1$ and all $n \geq \tau$, we have*

$$\mathbb{P}\left(\left\|\frac{1}{n}\sum_{i=1}^n \eta_i\right\|_H < 4\sqrt{\frac{\tau}{n}}\right) \geq 1 - 3e^{-\tau}.$$

Finally, we need the following theorem, see [7, Corollary 5.10], which was essentially shown by [13, 5, 3].

Theorem 3.3 *Let \mathbb{P} be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_{\infty} \leq 1$. We write $\Phi : X \rightarrow H$ for the canonical feature map of H , i.e., $\Phi(x) := k(\cdot, x)$, $x \in X$. Then for all $\lambda > 0$ there exists a function $h : X \times \mathbb{R} \rightarrow [-1, 1]$ such that for all $n \geq 1$ and all $D \in (X \times \mathbb{R})^n$ we have*

$$\|f_{D,\lambda} - f_{\mathbb{P},\lambda}\|_H \leq \lambda^{-1} \|\mathbb{E}_D h\Phi - \mathbb{E}_{\mathbb{P}} h\Phi\|_H,$$

where \mathbb{E}_D denotes the empirical average with respect to D .

Proof of of Theorem 2.1: In order to show the first estimate we fix a $\delta > 0$ and a $\lambda > 0$ such that $\delta\lambda \leq 4$. Let $\tau := \lambda^2\delta^2n/16$ which implies $n \geq \tau$. Combining Theorems 3.2 and 3.3 we then obtain

$$\begin{aligned} 1 - 3e^{-\tau} &\leq \mathbb{P}^n(D \in (X \times \mathbb{R})^n : \|\mathbb{E}_D h\Phi - \mathbb{E}_{\mathbb{P}} h\Phi\|_H \leq 4\sqrt{\tau/n}) \\ &\leq \mathbb{P}^n(D \in (X \times \mathbb{R})^n : \|f_{D,\lambda} - f_{\mathbb{P},\lambda}\|_H \leq \delta). \end{aligned} \quad (22)$$

Let us now assume that we have a training set $D \in (X \times \mathbb{R})^n$ such that $\|f_{\mathbb{P},\lambda} - f_{D,\lambda}\|_H \leq \delta$. Given a pair $(x, y) \in A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda})$, we then have

$$\epsilon + \delta < |f_{\mathbb{P},\lambda}(x) - y| \leq |f_{D,\lambda}(x) - y| + |f_{\mathbb{P},\lambda}(x) - f_{D,\lambda}(x)| \leq |f_{D,\lambda}(x) - y| + \delta$$

by the triangle inequality and $\|k\|_{\infty} \leq 1$ which implies $\|\cdot\|_{\infty} \leq \|\cdot\|_H$. In other words, we have $A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda}) \subset A_{\text{low}}(f_{D,\lambda})$. Consequently, Lemma 3.1 yields

$$\begin{aligned} \#SV(f_{D,\lambda}) &\geq \#\{i : |f_{D,\lambda}(x_i) - y_i| > \epsilon\} \geq \#\{i : |f_{\mathbb{P},\lambda}(x_i) - y_i| > \epsilon + \delta\} \\ &= \sum_{i=1}^n \mathbf{1}_{A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda})}(x_i, y_i). \end{aligned}$$

Combining this estimate with (22) we then obtain

$$\mathbb{P}^n\left(D \in (X \times \mathbb{R})^n : \frac{\#SV(f_{D,\lambda})}{n} \geq \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda})}(x_i, y_i)\right) \geq 1 - 3e^{-\frac{\delta^2\lambda^2n}{16}}.$$

Moreover, Hoeffding's inequality, see, e.g. [4, Theorem 8.1], shows

$$\mathbb{P}^n\left(D \in (X \times \mathbb{R})^n : \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda})}(x_i, y_i) > \mathbb{P}(A_{\text{low}}^{\delta}(f_{\mathbb{P},\lambda})) - \rho\right) \geq 1 - e^{-2\rho^2n}$$

for all $\rho > 0$ and $n \geq 1$. From these estimates and a union bound we conclude the first inequality. In order to show the second estimate we first observe that for training sets $D \in (X \times \mathbb{R})^n$ with $\|f_{P,\lambda} - f_{D,\lambda}\|_H \leq \delta$ we have $A_{\text{up}}(f_{D,\lambda}) \subset A_{\text{up}}^\delta(f_{P,\lambda})$. Lemma 3.1 then shows

$$\#SV(f_{D,\lambda}) \leq \sum_{i=1}^n \mathbf{1}_{A_{\text{up}}^\delta(f_{P,\lambda})}(x_i, y_i),$$

and hence (22) yields

$$\mathbb{P}^n \left(D \in (X \times \mathbb{R})^n : \frac{\#SV(f_{D,\lambda})}{n} \leq \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{A_{\text{up}}^\delta(f_{P,\lambda})}(x_i, y_i) \right) \geq 1 - 3e^{-\frac{\delta^2 \lambda^2 n}{16}}.$$

Using Hoeffding's inequality analogously to the proof of the first estimate we then obtain the second estimate. \blacksquare

Proof of Corollary 2.2: We first observe that we have $A_{\text{low}}^\delta(f_{P,\lambda}) \subset A_{\text{low}}^{\delta'}(f_{P,\lambda})$ for $0 \leq \delta' \leq \delta$. Let us show

$$\bigcup_{\delta > 0} A_{\text{low}}^\delta(f_{P,\lambda}) = A_{\text{low}}(f_{P,\lambda}). \quad (23)$$

Obviously, the inclusion “ \subset ” directly follows from the above monotonicity. Conversely, for $(x, y) \in A_{\text{low}}(f_{P,\lambda})$ we have $|f(x) - y| > \epsilon$ and hence $|f(x) - y| > \epsilon + \delta$ for some $\delta > 0$, i.e., we have shown $(x, y) \in A_{\text{low}}^\delta(f_{P,\lambda})$. From (23) we now conclude

$$\lim_{\delta \searrow 0} \mathbb{P}(A_{\text{low}}^\delta(f_{P,\lambda})) = \mathbb{P}(A_{\text{low}}(f_{P,\lambda})). \quad (24)$$

In addition, we have $A_{\text{up}}^{\delta'}(f_{P,\lambda}) \subset A_{\text{up}}^\delta(f_{P,\lambda})$ for $0 \leq \delta' \leq \delta$, and it is easy to check that

$$\bigcap_{\delta > 0} A_{\text{up}}^\delta(f_{P,\lambda}) = A_{\text{up}}(f_{P,\lambda}). \quad (25)$$

Indeed, if $(x, y) \in A_{\text{up}}^\delta(f_{P,\lambda})$ for all $\delta > 0$ we have $|f(x) - y| \geq \epsilon - \delta$ for all $\delta > 0$, from which we conclude $|f(x) - y| \geq \epsilon$, i.e. $(x, y) \in A_{\text{up}}(f_{P,\lambda})$. Conversely, the inclusion “ \supset ” directly follows from the above monotonicity of the sets A_{up}^δ . From (25) we then conclude

$$\lim_{\delta \searrow 0} \mathbb{P}(A_{\text{up}}^\delta(f_{P,\lambda})) = \mathbb{P}(A_{\text{up}}(f_{P,\lambda})). \quad (26)$$

Let us now fix a decreasing sequence $(\delta_n) \subset (0, 1)$ with $\delta_n \rightarrow 0$ and $\delta_n^2 n \rightarrow \infty$. Combining (24) and (26) with the estimates of Theorem 2.1, we then obtain the assertion. \blacksquare

Proof of Lemma 2.3: Since the loss function L_ϵ is Lipschitz continuous and convex in t , it is easy to verify that $t \mapsto \mathcal{C}_{L_\epsilon, P(\cdot|x)}(t)$ is Lipschitz continuous and convex for P_X -almost all $x \in X$, and hence $\mathcal{M}^*(x)$ is a closed interval. In order to prove the remaining assertions it suffices to show that $\lim_{t \rightarrow \pm\infty} \mathcal{C}_{L_\epsilon, P(\cdot|x)}(t) = \infty$ for P_X -almost all $x \in X$. To this end, we first observe that $\mathcal{R}_{L_\epsilon, P}^* < \infty$ implies $\mathcal{C}_{L_\epsilon, P(\cdot|x)}^* < \infty$ for P_X -almost all $x \in X$. Let us fix such an x , a $B > 0$, and a sequence $(t_n) \subset \mathbb{R}$ with $t_n \rightarrow -\infty$. By the shape of L_ϵ , there then exists an $r_0 > 0$ such that $L_\epsilon(y, t) \geq 2B$ for all $y, t \in \mathbb{R}$ with $|y - t| \geq r_0$. Furthermore, there exists an $M > 0$ with $\mathbb{P}([-M, M] | x) \geq 1/2$, and since $t_n \rightarrow -\infty$ there further exists an $n_0 \geq 1$ such that $t_n \leq -M - r_0$ for all $n \geq n_0$. For $y \in [-M, M]$ we thus have $y - t_n \geq r_0$, and hence we finally find

$$\mathcal{C}_{L_\epsilon, P(\cdot|x)}(t_n) \geq \int_{[-M, M]} L_\epsilon(y, t_n) d\mathbb{P}(y|x) \geq B$$

for all $n \geq n_0$. The case $t_n \rightarrow \infty$ can be shown analogously. \blacksquare

For the proof of Theorem 2.4 we need the following two intermediate results.

Theorem 3.4 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Assume that $\mathcal{R}_{L_\epsilon, P}(0) < \infty$ and that H is dense in $L_1(P_X)$. Then, for all $\delta > 0$ and $\rho > 0$, there exists a $\lambda_0 > 0$ such that for all $\lambda \in (0, \lambda_0]$ we have*

$$\mathbb{P}_X(\{x \in X : |f_{P,\lambda}(x) - t| > \delta \text{ for all } t \in \mathcal{M}^*(x)\}) < \rho.$$

Proof: Since H is dense in $L_1(P_X)$ we have $\inf_{f \in H} \mathcal{R}_{L_\epsilon, P}(f) = \mathcal{R}_{L_\epsilon, P}^*$ by [9, Theorem 3], and hence $\lim_{\lambda \rightarrow 0} \mathcal{R}_{L_\epsilon, P}(f_{P, \lambda}) = \mathcal{R}_{L_\epsilon, P}^*$. Now we obtain the assertion from [6, Theorem 3.16]. ■

Lemma 3.5 *Let P be a probability measure on $X \times \mathbb{R}$ and H be a separable RKHS with bounded measurable kernel satisfying $\|k\|_\infty \leq 1$. Assume that $\mathcal{R}_{L_\epsilon, P}(0) < \infty$ and that H is dense in $L_1(P_X)$. Then, for all $\delta > 0$ and $\rho > 0$, there exists a $\lambda_0 > 0$ such that for all $\lambda \in (0, \lambda_0]$ we have*

$$P(M_{\text{low}}^{2\delta}(f_{P, \lambda})) \leq P(A_{\text{low}}^\delta(f_{P, \lambda})) + \rho \quad \text{and} \quad P(M_{\text{up}}^{2\delta}(f_{P, \lambda})) \geq P(A_{\text{up}}^\delta(f_{P, \lambda})) - \rho.$$

Proof: We write $t_\lambda^*(x)$ for the real number defined by (6) for $f(x) := f_{P, \lambda}(x)$. Then we have

$$M_{\text{low}}^{2\delta}(f_{P, \lambda}) \subset \left(M_{\text{low}}^{2\delta}(f_{P, \lambda}) \cap \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t_\lambda^*(x)| \leq \delta\} \right) \cup \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t(x)| > \delta \text{ for all } t(x) \in \mathcal{M}^*(x)\}.$$

Moreover, given an $(x, y) \in M_{\text{low}}^{2\delta}(f_{P, \lambda}) \cap \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t_\lambda^*(x)| \leq \delta\}$, we find

$$\epsilon + 2\delta < |t_\lambda^*(x) - y| \leq |f_{P, \lambda}(x) - t_\lambda^*(x)| + |f_{P, \lambda}(x) - y| \leq \delta + |f_{P, \lambda}(x) - y|,$$

i.e., we have $(x, y) \in A_{\text{low}}^\delta(f_{P, \lambda})$. Estimating the probability of the remaining set by Theorem 3.4 then yields the first assertion. In order to prove the second estimate we first observe that

$$A_{\text{up}}^\delta(f_{P, \lambda}) \subset \left(A_{\text{up}}^\delta(f_{P, \lambda}) \cap \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t_\lambda^*(x)| \leq \delta\} \right) \cup \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t(x)| > \delta \text{ for all } t(x) \in \mathcal{M}^*(x)\}.$$

For $(x, y) \in A_{\text{up}}^\delta(f_{P, \lambda}) \cap \{(x, y) \in X \times \mathbb{R} : |f_{P, \lambda}(x) - t_\lambda^*(x)| \leq \delta\}$ we further have

$$\epsilon - \delta \leq |f_{P, \lambda}(x) - y| \leq |f_{P, \lambda}(x) - t_\lambda^*(x)| + |t_\lambda^*(x) - y| \leq \delta + |t_\lambda^*(x) - y|,$$

i.e., we have $(x, y) \in M_{\text{up}}^{2\delta}(f_{P, \lambda})$. Again, the assertion now follows from Theorem 3.4. ■

Proof of Theorem 2.4: Analogously to the proofs of (24) and (26), we find

$$\lim_{\delta \searrow 0} P(M_{\text{low}}^\delta(f_{P, \lambda})) = P(M_{\text{low}}(f_{P, \lambda})) \quad \text{and} \quad \lim_{\delta \searrow 0} P(M_{\text{up}}^\delta(f_{P, \lambda})) = P(M_{\text{up}}(f_{P, \lambda}))$$

Combining these equations with Theorem 2.1 and Lemma 3.5, we then obtain the assertion. ■

References

- [1] A. Christmann and I. Steinwart. Consistency and robustness of kernel based regression. *Bernoulli*, 13:799–819, 2007.
- [2] N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines*. Cambridge University Press, Cambridge, 2000.
- [3] E. De Vito, L. Rosasco, A. Caponnetto, M. Piana, and A. Verri. Some properties of regularized kernel methods. *J. Mach. Learn. Res.*, 5:1363–1390, 2004.
- [4] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1996.
- [5] I. Steinwart. Sparseness of support vector machines. *J. Mach. Learn. Res.*, 4:1071–1105, 2003.
- [6] I. Steinwart. How to compare different loss functions. *Constr. Approx.*, 26:225–287, 2007.
- [7] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.
- [8] I. Steinwart and A. Christmann. How SVMs can estimate quantiles and the median. In J.C. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 305–312. MIT Press, Cambridge, MA, 2008.
- [9] I. Steinwart, D. Hush, and C. Scovel. Function classes that approximate the Bayes risk. In G. Lugosi and H. U. Simon, editors, *Proceedings of the 19th Annual Conference on Learning Theory*, pages 79–93. Springer, New York, 2006.
- [10] V. Vapnik, S. Golowich, and A. Smola. Support vector method for function approximation, regression estimation, and signal processing. In M. Mozer, M. Jordan, and T. Petsche, editors, *Advances in Neural Information Processing Systems 9*, pages 81–287. MIT Press, Cambridge, MA, 1997.
- [11] V. N. Vapnik. *Statistical Learning Theory*. John Wiley & Sons, New York, 1998.
- [12] V. Yurinsky. *Sums and Gaussian Vectors*. Lecture Notes in Math. 1617. Springer, Berlin, 1995.
- [13] T. Zhang. Convergence of large margin separable linear classification. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 357–363. MIT Press, Cambridge, MA, 2001.