
Fast-rate PAC-Bayes Generalization Bounds via Shifted Rademacher Processes

Jun Yang*

Department of Statistical Sciences
University of Toronto,
Vector Institute
jun@utstat.toronto.edu

Shengyang Sun*

Department of Computer Science
University of Toronto,
Vector Institute
ssy@cs.toronto.edu

Daniel M. Roy

Department of Statistical Sciences
University of Toronto,
Vector Institute
droy@utstat.toronto.edu

Abstract

The developments of Rademacher complexity and PAC-Bayesian theory have been largely independent. One exception is the PAC-Bayes theorem of Kakade, Sridharan, and Tewari [21], which is established via Rademacher complexity theory by viewing Gibbs classifiers as linear operators. The goal of this paper is to extend this bridge between Rademacher complexity and state-of-the-art PAC-Bayesian theory. We first demonstrate that one can match the fast rate of Catoni’s PAC-Bayes bounds [8] using shifted Rademacher processes [27, 43, 44]. We then derive a new fast-rate PAC-Bayes bound in terms of the “flatness” of the empirical risk surface on which the posterior concentrates. Our analysis establishes a new framework for deriving fast-rate PAC-Bayes bounds and yields new insights on PAC-Bayesian theory.

1 Introduction

PAC-Bayes theory [33, 38] was developed to provide *probably approximately correct* (PAC) guarantees for supervised learning algorithms whose outputs can be expressed as a weighted majority vote. Its uses have expanded considerably since [3, 6, 14, 17, 18, 28, 39, 40]. See [12, 25, 32] for gentle introductions. Indeed, there has been a surge of interest and work in PAC-Bayes theory and its application to large-scale neural networks, especially towards studying generalization in overparametrized neural networks trained by variants of gradient descent [9–11, 30, 36, 37].

PAC-Bayes bounds are one of several tools available for the study of the generalization and risk properties of learning algorithms. One advantage of the PAC-Bayes framework is its ease of use: one can obtain high-probability risk bounds for arbitrary (“posterior”) Gibbs classifiers provided one can compute or bound relative entropies with respect to some fixed (“prior”) Gibbs classifier. Another tool for studying generalization is Rademacher complexity, a distribution-dependent complexity measure for classes of real-valued functions [4, 5, 23, 29, 34, 44].

The literature on PAC-Bayes bounds and bounds based on Rademacher complexity are essentially disjoint. One point of contact is the work of Kakade, Sridharan, and Tewari [21], which builds the first bridge between PAC-Bayes theory and Rademacher complexity. By viewing Gibbs classifiers as

*These authors contributed equally.

linear operators and relative entropy as a strictly convex regularizer, they were able to use their general Rademacher complexity bounds on strictly convex linear classes to develop a slightly sharper version of McAllester’s PAC-Bayes bound [33]. This result offers new insight on PAC-Bayes theory, including potential roles for data-dependent complexity estimates and stability. However, even within the PAC-Bayes community, this result is relatively unknown.

While the PAC-Bayes bound established by Kakade, Sridharan, and Tewari improves on McAllester’s bound, it still converges at a slow $1/\sqrt{m}$ rate, where m denotes the number of data used to form the empirical risk estimate. This observation raises the question of whether one can match state-of-the-art PAC-Bayes bounds via a Rademacher-process argument. In particular, can one match Catoni’s bound [8, Thm. 1.2.6], which can obtain a fast $1/m$ rate of convergence?

There is an extensive literature on the problem of obtaining fast $1/m$ rates of convergence for the generalization error of (approximate) empirical risk minimization (ERM). Available approaches include the use of local Rademacher complexity [4, 22], shifted empirical processes [27], offset Rademacher complexities [29], and local empirical entropy [44]. See also [16, 19, 20, 26, 31, 35] and [13] for an extensive survey. To date, these techniques have not been connected to PAC-Bayesian theory, which presents the opportunity to obtain new PAC-Bayes theory for ERM.

1.1 Contributions

In this paper, we extend the bridge between Rademacher process theory and PAC-Bayes theory by constructing new bounds using Rademacher process techniques. Among our contributions:

- i) We show how to recover Catoni’s fast-rate PAC-Bayes bound [8], up to constants, using tail bounds on shifted Rademacher processes, which are special cases of shifted empirical processes [27, 43, 44]; See Section 3.
- ii) We derive a new fast-rate PAC-Bayes bound, building on our shifted-Rademacher-process approach. This bound is determined by the “flatness” of the empirical risk surface on which the posterior Gibbs classifier concentrates. The notion of “flatness” is inspired by the proposal by Dziugaite and Roy [9] to formalize the empirical connection between “flat minima” and generalization using PAC-Bayes bounds; See Section 4.
- iii) More generally, we introduce a new approach to derive fast-rate PAC-Bayes bounds and, in turn, offer new insight on PAC-Bayesian theory.

2 Background

Let \mathcal{D} be an unknown distribution over a space \mathcal{Z} of labeled examples, and let \mathcal{H} be a hypothesis class. Relative to a binary loss function $\ell : \mathcal{H} \times \mathcal{Z} \rightarrow \{0, 1\}$, we define the associated loss class $\mathcal{F} := \{\ell(h, \cdot) : h \in \mathcal{H}\}$ of functions from $\mathcal{Z} \rightarrow \{0, 1\}$, each associated to one or more hypotheses. Let $\mathcal{L}_{\mathcal{D}}(f) := \mathbb{E}_{z \sim \mathcal{D}} f(z)$ denote the expected loss, i.e., risk, of every hypothesis associated to f . Let $S = (z_1, \dots, z_m) \sim \mathcal{D}^m$ be a sequence of i.i.d. random variables. Let $\hat{\mathcal{L}}_S(f) = \frac{1}{m} \sum_{i=1}^m f(z_i)$ denote the empirical risk of every hypothesis associated to f .

We will be primarily interested in Gibbs classifiers, i.e., distributions P on \mathcal{F} which are interpreted as randomized classifiers that classify each new example according to a hypothesis drawn independently from P . (It is more common to work with distributions over \mathcal{H} , but these lead to looser results.) For a Gibbs classifier P and labeled example $z \in \mathcal{Z}$, let $\mathbb{E}_P f(z) = \mathbb{E}_{f \sim P}[f(z)]$ be the expected loss P suffers when labeling z . For Gibbs classifiers, the (expected) risk is defined to be $\mathcal{L}_{\mathcal{D}}(P) := \mathbb{E}_{f \sim P} \mathcal{L}_{\mathcal{D}}(f) = \mathbb{E}_{z \sim \mathcal{D}} \mathbb{E}_P f(z)$. The (expected) empirical risk is $\hat{\mathcal{L}}_S(P) := \mathbb{E}_{f \sim P} \hat{\mathcal{L}}_S(f) = \frac{1}{m} \sum_{i=1}^m \mathbb{E}_P f(z_i)$.

2.1 PAC-Bayes

The PAC-Bayes framework [33] provides data-dependent generalization guarantees for Gibbs classifiers. Each bound is specified in terms of a Gibbs classifier P called the *prior*, as it must be independent of the training sample. The bound then holds for all *posterior* distributions, i.e., Gibbs classifiers that may be defined in terms of the training sample.

Theorem 2.1 (PAC-Bayes [33]). *For any prior distribution P over \mathcal{F} , for any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over draws of training data $S \sim \mathcal{D}^m$, for all distributions Q over \mathcal{F} ,*

$$\mathcal{L}_Q(Q) \leq \hat{\mathcal{L}}_S(Q) + \sqrt{\frac{\text{KL}(Q||P) + \log \frac{m}{\delta}}{2(m-1)}}. \quad (1)$$

Note in Theorem 2.1, the generalization bound scales as $\mathcal{O}(m^{-\frac{1}{2}})$. Catoni [8] presents a *fast rate* PAC-Bayesian bound, in which the generalization bound scales as $\mathcal{O}(m^{-1})$.

Theorem 2.2 (Fast-Rate PAC-Bayes [8, Thm 1.2.6]). *For any prior distribution P over \mathcal{F} , for any $\delta \in (0, 1)$ and $C > 0$, with probability at least $1 - \delta$ over draws of training data $S \sim \mathcal{D}^m$, for all distributions Q over \mathcal{F} ,*

$$\mathcal{L}_Q(Q) \leq \frac{1}{1 - e^{-C}} \left[C \hat{\mathcal{L}}_S(Q) + \frac{\text{KL}(Q||P) + \log \frac{1}{\delta}}{m} \right]. \quad (2)$$

Because the constant $C/(1 - e^{-C}) > 1$ holds for any $C > 0$, the generalization bound in Theorem 2.2 will always be bounded below by the empirical risk. Usually for a well-trained distribution Q over training set, the empirical risk $\hat{\mathcal{L}}_S(Q)$ is small, therefore the generalization bound is dominated by the KL term. Compared to the standard PAC-Bayes bound in Theorem 2.1, where the KL term decreases at a rate $\mathcal{O}(m^{-\frac{1}{2}})$, the KL term of Catoni’s bound decreases at a rate $\mathcal{O}(m^{-1})$. For this reason, we say that Catoni’s bound achieves a *fast rate* of convergence. Note that fast-rate bounds can lead to much tighter bounds. Of course, $C/(1 - e^{-C}) \rightarrow 1$ as $C \rightarrow 0$, but, in that limit, the constants ignored in the asymptotic rate $\mathcal{O}(m^{-1})$ degrade. (See [28] for more discussion.)

2.2 Rademacher Viewpoint

Fix a prior Gibbs classifier P on \mathcal{F} . Then, for measurable functions g, h , consider the inner product $\langle g, h \rangle = \int g(f)h(f)P(df)$. The key observation of Kakade, Sridharan, and Tewari is that one can view $\mathcal{L}_Q(Q)$ (resp., $\hat{\mathcal{L}}_S(Q)$) as the inner product $\langle dQ/dP, \mathcal{L}_Q(\cdot) \rangle$ (resp., $\langle dQ/dP, \hat{\mathcal{L}}_S(\cdot) \rangle$) between the posterior Q , represented by its Radon–Nikodym derivative with P , and the risk (resp., empirical risk), viewed as measurable function on \mathcal{F} . Thus, Gibbs classifiers can be viewed as linear predictors. Using their distribution-independent bounds on the Rademacher complexity of certain classes of linear predictors, Kakade, Sridharan, and Tewari [21] derive a PAC-Bayes bound similar to Theorem 2.1. We refer to this as the “Rademacher viewpoint” on PAC-Bayes.

We now summarize their argument in more detail. Let $\mathcal{Q}(\kappa) := \{Q : \text{KL}(Q||P) \leq \kappa\}$. One can follow the classical steps for controlling the generalization error uniformly over $\mathcal{Q}(\kappa)$ using Rademacher complexity. Their first step is to connect $\sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_Q(Q) - \hat{\mathcal{L}}_S(Q)]$ to $\mathbb{E}_S \sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_Q(Q) - \hat{\mathcal{L}}_S(Q)]$ by the bounded difference inequality (McDiarmid’s inequality). In particular, with probability at least $1 - \delta$,

$$\sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_Q(Q) - \hat{\mathcal{L}}_S(Q)] \leq \mathbb{E}_S \sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_Q(Q) - \hat{\mathcal{L}}_S(Q)] + \sqrt{\frac{\log(1/\delta)}{m}}. \quad (3)$$

Then they apply a symmetrization argument to obtain an upper bound in terms of Rademacher complexity [5]. In particular, recalling that $S = (z_1, \dots, z_m)$ is our training data,

$$\mathbb{E}_S \sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_Q(Q) - \hat{\mathcal{L}}_S(Q)] \leq 2\mathbb{E}_S \mathbb{E}_\epsilon \sup_{Q \in \mathcal{Q}(\kappa)} \left[\frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{E}_Q f(z_i) \right], \quad (4)$$

where $\{\epsilon_i\}$ are i.i.d. Rademacher random variables, i.e., $\mathbb{P}(\epsilon_i = +1) = \mathbb{P}(\epsilon_i = -1) = 1/2$. Their last step is to bound the Rademacher complexity $\mathbb{E}_S \mathbb{E}_\epsilon \sup_{Q \in \mathcal{Q}(\kappa)} [\frac{1}{m} \sum_{i=1}^m \epsilon_i \mathbb{E}_Q f(z_i)]$, which can be seen as the Rademacher complexity of a linear class with a (strongly) convex constraint [21]. According to [21], the Rademacher complexity in Eq. (4) is of order $\sqrt{\kappa/m}$, which eventually leads to a term of order $\sqrt{\text{KL}(Q||P)/m}$ after applying a union bound argument on κ .

In the end, using the above arguments and their sharp bounds on the Rademacher and Gaussian complexities of (constrained) linear classes [21, Thm. 1], Kakade, Sridharan, and Tewari obtain the

following PAC-Bayes bound [21, Cor. 8]: for every prior P over \mathcal{F} , with probability at least $1 - \delta$ over draws of training data $S \sim \mathcal{D}^m$, for all distribution Q over \mathcal{F} ,

$$\mathcal{L}_{\mathcal{D}}(Q) \leq \hat{\mathcal{L}}_S(Q) + 4.5 \sqrt{\frac{\max\{\text{KL}(Q||P), 2\}}{m}} + \sqrt{\frac{\log(1/\delta)}{m}}. \quad (5)$$

Note that this PAC-Bayes bound has a slow rate of $\sqrt{1/m}$, but it slightly improves the rate in the term $\sqrt{\log(m/\delta)/m}$ of McAllester’s bound [33] to $\sqrt{\log(1/\delta)/m}$.

Since McAllester’s bound is far from the state-of-art in PAC-Bayesian theory, this raises the question whether one can extend the “Rademacher viewpoint” of PAC-Bayes to derive more advanced bounds, such as one matching the fast rate of Catoni’s bound.

3 Extending the Rademacher Viewpoint

There are at least two difficulties in the “Rademacher viewpoint” that prevent fast rates. First, if we connect the generalization error to Rademacher complexity using the bounded difference inequality, a slow rate term $\sqrt{\log(1/\delta)/m}$ will appear. Second, as is shown by Kakade, Sridharan, and Tewari [21], the standard Rademacher complexity of (constraint) linear classes leads to an upper bound with a slow rate of order $\mathcal{O}(\sqrt{\text{KL}(Q||P)/m})$. Therefore, in order to derive fast rate PAC-Bayes bounds, we need to extend the “Rademacher viewpoint”.

In order to obtain fast rates, we work with so-called shifted Rademacher processes, i.e., processes of the form $\{\frac{1}{m} \sum_{i=1}^m \varepsilon'_i f(z_i)\}_{f \in \mathcal{F}}$ where the variables $\{\varepsilon'_i\}$ are independent from S , i.i.d., and take two values with equal probability. (These shifted Rademacher variables, $\{\varepsilon'_i\}$, are not necessarily zero mean. When they take values in $\{\pm 1\}$, we obtain a standard Rademacher process.) Shifted Rademacher processes are examples of shifted empirical processes [27, 43, 44].

Recall that Rademacher complexity is the *expected value* of the supremum of Rademacher processes over a class [5]. In order to get a fast rate, we connect the *tail probabilities* of the supremum of the generalization error to the *tail probabilities* of shifted Rademacher processes via a symmetrization-in-deviation argument instead of the symmetrization-in-expectation argument. The key is that we can avoid using the bounded difference inequality by bounding the deviation. This removes the slow rate term of $\sqrt{\log(1/\delta)/m}$. It remains to bound the deviation of shifted Rademacher processes to get a fast rate bound of order $\mathcal{O}(\text{KL}(Q||P)/m)$.

In the following, we demonstrate how the extended “Rademacher viewpoint” via shifted Rademacher processes can be applied to derive a fast rate PAC-Bayes bound that matches the fast rate of Catoni’s bound. Note that, since $C/(1 - e^{-C}) > 1$ for fixed $C > 0$ in Catoni’s bound in Eq. (2), we can write $C/(1 - e^{-C}) = 1 + c$ for some constant $c > 0$. Furthermore, note that our goal in this section is not to derive new PAC-Bayes bounds. Therefore, we do not make attempts to optimize the constants.

Proposition 3.1 (Matching Catoni’s Fast Rate via Shifted Rademacher Processes). *For any given $c > 0$ and prior P over \mathcal{F} , there exists constants C_1 , C_2 , and C_3 such that, with probability at least $1 - \delta$, for all distributions Q over \mathcal{F}*

$$\mathcal{L}_{\mathcal{D}}(Q) \leq (1 + c) \hat{\mathcal{L}}_S(Q) + C_1 \frac{\text{KL}(Q||P)}{m} + C_2 \frac{\log \frac{1}{\delta}}{m} + C_3 \frac{1}{m}. \quad (6)$$

Outline of the proof. We wish to emphasize two key differences from traditional machinery for deriving Rademacher-complexity-based generalization bounds. The complete proof is given in Appendix A.1.

Fix P and let $\mathcal{Q}(\kappa) := \{Q : \text{KL}(Q||P) \leq \kappa\}$ be defined as in Section 2.2. Rather than control $\sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_{\mathcal{D}}(Q) - \hat{\mathcal{L}}_S(Q)]$ in terms of its *expectation* via the bounded difference inequality and Rademacher complexity, we bound the *tail/deviation* of $\sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_{\mathcal{D}}(Q) - (1 + c) \hat{\mathcal{L}}_S(Q)]$, thus avoiding the use of the bounded differences inequality altogether. In particular, we can obtain fast rates by bounding the tail in terms of tail of supremum of shifted Rademacher processes [27, 43, 44].

Define $\mathcal{G}_\kappa := \{\mathbb{E}_Q f(\cdot) : Q \in \mathcal{Q}(\kappa)\}$ and, by an abuse of notation, let $\mathcal{L}_\mathcal{D}(g)$ denote $\mathbb{E}_{z \sim \mathcal{D}}[g(z)]$. Then we can write $\sup_{Q \in \mathcal{Q}(\kappa)} [\mathcal{L}_\mathcal{D}(Q) - (1+c)\mathcal{L}_S(Q)]$ as $\sup_{g \in \mathcal{G}_\kappa} [\mathcal{L}_\mathcal{D}(g) - (1+c)\mathcal{L}_S(g)]$. We start from bounding the tail probability $\mathbb{P}_S(\sup_{g \in \mathcal{G}_\kappa} \mathcal{L}_\mathcal{D}(g) - (1+c)\mathcal{L}_S(g) \geq t)$. For fixed constants $c > c_2 > 0$, let $c' = \frac{c-c_2}{1+c_2}$ and $t' = \frac{t}{2(1+c_2)}$. Then, by [44, Cor. 1], we have

$$\mathbb{P}_S \left(\sup_{g \in \mathcal{G}_\kappa} \mathcal{L}_\mathcal{D}(g) - (1+c)\mathcal{L}_S(g) \geq t \right) \leq 4\mathbb{P}_{S,\varepsilon} \left(\sup_{g \in \mathcal{G}_\kappa} \left[\frac{1+\frac{c'}{2}}{m} \sum_{i=1}^m \left(\varepsilon_i - \frac{c'}{2+c'} \right) g(z_i) \right] \geq \frac{t'}{2} \right). \quad (7)$$

Letting $\varepsilon'_i := \varepsilon_i - \frac{c'}{2+c'}$, one can see that $\{\varepsilon'_i\}$ are i.i.d. “shifted” Rademacher random variables with mean $-\frac{c'}{2+c'}$. For any $g \in \mathcal{G}_\kappa$, there exists $Q \in \mathcal{Q}(\kappa)$ such that

$$\frac{1}{m} \sum_{i=1}^m \varepsilon'_i g(z_i) = \frac{1}{m} \sum_{i=1}^m \varepsilon'_i \mathbb{E}_Q f(z_i) = \mathbb{E}_Q \left[\frac{1}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right], \quad (8)$$

which can be viewed as a linear function of Q . Further, it can be verified that the set $\mathcal{Q}(\kappa)$ is (strongly) convex. Therefore, $\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \varepsilon'_i \mathbb{E}_Q f(z_i)$ is a convex optimization problem. By duality [7, Chp. 5], and, in this particular case, the Legendre transform of Kullback–Leibler divergence (see, e.g., [18]), we have

$$\sup_{g \in \mathcal{G}_\kappa} \frac{1}{m} \sum_{i=1}^m \varepsilon'_i g(z_i) = \sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \varepsilon'_i \mathbb{E}_Q f(z_i) = \inf_{\lambda > 0} \left\{ \frac{\kappa}{\lambda} + \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right] \right\}. \quad (9)$$

Combining the shifted symmetrization in deviation in Eq. (7) and the dual problem in Eq. (9), Markov’s inequality yields, for every $\lambda > 0$,

$$\mathbb{P}_S \left(\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q [\mathcal{L}_\mathcal{D}(f) - (1+c)\mathcal{L}_S(f)] \geq t \right) \leq 4e^{\kappa - \frac{\lambda t'}{2+c'}} \mathbb{E}_S \mathbb{E}_\varepsilon \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right]. \quad (10)$$

We then exploit the shifted property of ε'_i to bound the expectation term on the right-hand side and obtain fast rates. In particular, we show that, so long as $k \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}$,

$$\mathbb{E}_P \mathbb{E}_S \mathbb{E}_\varepsilon \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m (\varepsilon_i - k) f(z_i) \right) \right] \leq 1. \quad (11)$$

In our case, $k = \frac{c'}{2+c'}$, which leads to constraints relating λ , c , and c_2 . In particular, when $c = 0$, the required condition for the above result, $k \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}$, does not hold. Therefore, this approach obtains fast rates only if $c > 0$, i.e., if we shift. Combining Eqs. (10) and (11), there exists a constant C' , depending only on c , c_2 and δ , such that, with probability at least $1 - \delta$,

$$\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q [\mathcal{L}_\mathcal{D}(f) - (1+c)\mathcal{L}_S(f)] \leq \frac{C'}{m} (\kappa + \log(4/\delta)). \quad (12)$$

Finally, we may apply the same union-bound argument as in the proof of [21, Cor. 7] in order to cover all possible values of κ . This completes the proof. \square

4 New Fast Rate PAC-Bayes Bound based on “Flatness”

The extended “Rademacher viewpoint” of PAC-Bayes provides a new approach for deriving fast-rate PAC-Bayes bounds. In this section, we demonstrate the use of shifted Rademacher processes to derive a new fast-rate PAC-Bayes bound using a notion of “flatness”. This notion is inspired by the proposal by Dziugaite and Roy [9] to formalize the empirical connection between “flat minima” and generalization using PAC-Bayes bounds, and, in particular, posterior distributions which concentrate in these “flat minima”.

Definition 4.1 (Notion of “Flatness”). For given $h \in [0, 1]$, the “ h -flatness” of Q (w.r.t. S) is

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2. \quad (13)$$

One way to understand this new notion is to observe that, under zero-one loss, h -flatness can be written as the difference between the empirical risk and the quadratic empirical risk:

$$\frac{1}{m} \sum_{i=1}^m \mathbb{E}_Q[f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2 = \hat{\mathcal{L}}_S(Q) - \frac{1-h^2}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2. \quad (14)$$

Note that, for $[0, 1]$ -valued (bounded) loss, equality is replaced by an inequality: the r.h.s. is an upper bound of the l.h.s.

Remark 4.2. To see that optimizing h -flatness prefers “flat minima”, consider the following simplified case: Call a posterior Q “completely flat” if $f = g$ on S a.s., when $f, g \sim Q$. It can be verified that, if the posterior is “completely flat”, then under the zero-one loss, the “ h -flatness” is $h^2 \hat{\mathcal{L}}_S(Q)$. That is, given a “completely flat” posterior, the “ h -flatness” goes to zero as $h \rightarrow 0$. For $h > 0$, the “ h -flatness” is zero when Q is “completely flat” and $\hat{\mathcal{L}}_S(Q) = 0$. \triangleleft

The following PAC-Bayes theorem establishes favorable bounds for h -flat posteriors:

Theorem 4.3 (Fast Rate PAC-Bayes using “Flatness”). *For any given $c > 0$ and $h \in (0, 1)$, with probability at least $1 - \delta$ over random draws of training set $S \sim \mathcal{D}^m$, for all distributions Q over \mathcal{F} ,*

$$\mathcal{L}_{\mathcal{D}}(Q) \leq \hat{\mathcal{L}}_S(Q) + \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q[f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2 + \frac{4}{Cm} \left[3\text{KL}(Q||P) + \log \frac{1}{\delta} + 5 \right], \quad (15)$$

where $C = \frac{2h^4 c}{1+16h^2 c}$.

This bound can be tighter than Catoni’s bound under certain conditions. We delay the comparison with Catoni’s bound to Section 4.1. We now give an outline of the proof of Theorem 4.3, highlighting the technical differences from the proof of Proposition 3.1. The complete proof is given in Appendix A.2.

Outline of the proof of Theorem 4.3. By Eq. (14), we can write

$$\begin{aligned} \mathbb{E}_Q \mathcal{L}_{\mathcal{D}}(f) - \hat{\mathcal{L}}_S(Q) - \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q[f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2 \\ = \mathcal{L}_{\mathcal{D}}(Q) - (1+c)\hat{\mathcal{L}}_S(Q) + \frac{c(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2. \end{aligned} \quad (16)$$

There are at least two new challenges compared with the proof of Proposition 3.1. First, the shifted symmetrization in Eq. (7) cannot be applied because of the existence of the quadratic term $\frac{c(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2$. This means we need to derive a new shifted symmetrization involving the quadratic term. Second, the quadratic term $\frac{c(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2$ cannot be seen as a linear function of Q . Therefore, some technical arguments are required in order to apply the Legendre transform of Kullback–Leibler divergence.

First, we derive a new shifted symmetrization which involves quadratic terms. The proof is inspired by an argument due to Zhivotovskiy and Hanneke [44]. The result extends [44, Cor. 1], which is recovered as a special case when $h = 1$. For $\kappa > 0$, recall that we have defined $\mathcal{Q}(\kappa) = \{Q : \text{KL}(Q||P) \leq \kappa\}$ and $\mathcal{G}_{\kappa} = \{\mathbb{E}_Q f(\cdot) : Q \in \mathcal{Q}(\kappa)\}$. Then for any $g \in \mathcal{G}_{\kappa}$, there exists a $Q \in \mathcal{Q}(\kappa)$ such that $g = \mathbb{E}_Q f(\cdot)$. We can first show a tail bound that for any given $c_2 > 0$ and $g \in \mathcal{G}_{\kappa}$, if $t \geq \frac{(1+c_2)(1+c_2 h^2)}{mc_2 h^2}$, then

$$\mathbb{P}_S \left(\mathcal{L}_{\mathcal{D}}(g) - (1+c_2)\hat{\mathcal{L}}_S(g) + c_2(1-h^2)\hat{\mathcal{L}}_S(g^2) \geq \frac{t}{2} \right) \leq \frac{1}{2}. \quad (17)$$

Then, consider another independent random data set $S' = \{z'_1, \dots, z'_m\} \in \mathcal{D}^m$. For $c > c_2$, by taking the difference of $\mathcal{L}_{\mathcal{D}}(g) - (1+c)\hat{\mathcal{L}}_S(g) + c(1-h^2)\hat{\mathcal{L}}_S(g^2)$ and $\mathcal{L}_{\mathcal{D}}(g) - (1+c_2)\hat{\mathcal{L}}_{S'}(g) + c_2(1-h^2)\hat{\mathcal{L}}_{S'}(g^2)$

$h^2)\mathcal{L}_{S'}(g^2)$ and using Eq. (17), we obtain

$$\frac{1}{4}\mathbb{P}_S\left(\sup_{g\in\mathcal{G}_\kappa}\mathcal{L}_\mathcal{D}(g)-(1+c)\mathcal{L}_S(g)+c(1-h^2)\mathcal{L}_S(g^2)\geq t\right) \quad (18)$$

$$\leq \frac{1}{2}\mathbb{P}_{S,S'}\left(\sup_{g\in\mathcal{G}_\kappa}(1+c_2)\mathcal{L}_{S'}(g)-c_2(1-h^2)\mathcal{L}_{S'}(g^2)-(1+c)\mathcal{L}_S(g)+c(1-h^2)\mathcal{L}_S(g^2)\geq \frac{t}{2}\right). \quad (19)$$

Now by writing $(1+c_2)\mathcal{L}_{S'}(g)-c_2(1-h^2)\mathcal{L}_{S'}(g^2)-(1+c)\mathcal{L}_S(g)+c(1-h^2)\mathcal{L}_S(g^2)$ as

$$\begin{aligned} & \left(1+\frac{c+c_2}{2}\right)(\mathcal{L}_{S'}(g)-\mathcal{L}_S(g))-\frac{c+c_2}{2}(1-h^2)(\mathcal{L}_{S'}(g^2)-\mathcal{L}_S(g^2)) \\ & -\frac{c-c_2}{2}\mathcal{L}_S(g-(1-h^2)g^2)-\frac{c-c_2}{2}\mathcal{L}_{S'}(g-(1-h^2)g^2), \end{aligned} \quad (20)$$

one can apply the symmetrization argument to get

$$\begin{aligned} & \frac{1}{4}\mathbb{P}_S\left(\sup_{g\in\mathcal{G}_\kappa}\mathcal{L}_\mathcal{D}(g)-(1+c)\mathcal{L}_S(g)+c(1-h^2)\mathcal{L}_S(g^2)\geq t\right) \\ & \leq \mathbb{P}_{S,\epsilon}\left(\sup_{g\in\mathcal{G}_\kappa}\left[\frac{1}{m}\sum_{i=1}^m\epsilon_i((1+c')g(z_i)-c'(1-h^2)g^2(z_i))-c''\mathcal{L}_S(g-(1-h^2)g^2)\right]\geq \frac{t}{4}\right), \end{aligned} \quad (21)$$

where $c' = \frac{c+c_2}{2}$, $c'' = \frac{c-c_2}{2}$. Therefore, we have derived the new shifted symmetrization in deviation involving a quadratic term.

Recalling the definition of \mathcal{G}_κ , we have

$$\begin{aligned} & \sup_{g\in\mathcal{G}_\kappa}\frac{1}{m}\sum_{i=1}^m\epsilon_i((1+c')g(z_i)-c'(1-h^2)g^2(z_i))-c''\mathcal{L}_S(g-(1-h^2)g^2) \\ & = \sup_{Q\in\mathcal{Q}(\kappa)}\frac{1}{m}\sum_{i=1}^m[\epsilon_i(1+c')-c'']\mathbb{E}_Q f(z_i)-[\epsilon_i c'-c''](1-h^2)[\mathbb{E}_Q f(z_i)]^2. \end{aligned} \quad (22)$$

Note that there are two shifted Rademacher random variables $\epsilon_i(1+c')-c''$ and $\epsilon_i c'-c''$, which not only involve a shift term $-c''$ but also scale terms $(1+c')$ and c' , respectively. Furthermore, the term $[\mathbb{E}_Q f(z_i)]^2$ cannot be seen as a linear function of Q . This prevents the use of the key argument in [21] to formulate an upper bound using Rademacher complexities of constrained linear classes by considering the generalization error as a linear function of Q .

In order to sidestep this obstruction, define $\epsilon := \{\epsilon_i\}_{i=1}^m, \mathbf{z} := \{z_i\}_{i=1}^m$ and suppose $\hat{Q}(\epsilon, \mathbf{z})$ achieves the supremum above. (If the supremum cannot be achieved, one can use a carefully chosen sequence of $\{\hat{Q}_i(\epsilon, \mathbf{z})\}$ to prove the same statement as the supremum can be approximated arbitrarily closely.) The following inequality then holds:

$$\begin{aligned} & \sup_{Q\in\mathcal{Q}(\kappa)}\frac{1}{m}\sum_{i=1}^m[\epsilon_i(1+c')-c'']\mathbb{E}_Q f(z_i)-[\epsilon_i c'-c''](1-h^2)[\mathbb{E}_Q f(z_i)]^2 \\ & \leq \sup_{Q\in\mathcal{Q}(\kappa)}\frac{1}{m}\sum_{i=1}^m[\epsilon_i(1+c')-c'']\mathbb{E}_Q f(z_i)-[\epsilon_i c'-c''](1-h^2)\mathbb{E}_Q f(z_i)\mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i). \end{aligned} \quad (23)$$

To see this, note that, on the one hand, if we plug in $Q = \hat{Q}(\epsilon, \mathbf{z})$ the inequality is tight; on the other hand, by definition, $Q = \hat{Q}(\epsilon, \mathbf{z})$ already achieves the supremum of the l.h.s. Note that the r.h.s. can be seen as a linear function of Q , because $\hat{Q}(\epsilon, \mathbf{z})$ is a random variable which does not depend on Q .

Let $\epsilon_i'' := \epsilon_i c'-c'' = \epsilon_i \frac{c_1+c_2}{2} - \frac{c_1-c_2}{2}$. Then by keeping the term $\hat{Q}(\epsilon, \mathbf{z})$, one can apply the convex conjugate of relative entropy to get

$$\begin{aligned} & \mathbb{P}\left[\sup_{Q\in\mathcal{Q}(\kappa)}\mathcal{L}_\mathcal{D}(Q)-(1+c)\mathcal{L}_S(Q)+\frac{c(1-h^2)}{m}\sum_{i=1}^m(\mathbb{E}_Q f(z_i))^2\geq t\right] \\ & \leq 4\exp\left(\kappa-\frac{\lambda t}{4}\right)\mathbb{E}_S\mathbb{E}_\epsilon\mathbb{E}_P\left[\exp\left(\frac{\lambda}{m}\sum_{i=1}^m f(z_i)\left[(\epsilon_i+\epsilon_i'')-\epsilon_i''(1-h^2)\mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i)\right]\right)\right]. \end{aligned} \quad (24)$$

Therefore, the problem turns to bounding the expectation of a function involving shifted Rademacher processes. Although the expectation looks quite complicated since it involves two scaled and shifted Rademacher variables as well as the unknown $\hat{Q}(\epsilon, \mathbf{z})$, fortunately, we are able to show that, for any random variables $Y_i \in [0, 1]$, we have

$$\mathbb{E}_S \mathbb{E}_\epsilon \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m f(z_i) [(\epsilon_i + \epsilon_i'') - \epsilon_i''(1 - h^2)Y_i] \right) \right] \leq 1, \quad (25)$$

if $h \in (0, 1]$, $1 > h^2 c > c_2 > 0$ and $0 < \frac{\lambda}{m} < C = \frac{h^2 c - c_2}{2(1+h^2 c)(1+c_2)}$. This result removes the term $\hat{Q}(\epsilon, \mathbf{z})$ by letting $Y_i = \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i)$. Finally, we combine different values of κ by a union bound argument similar to the proof of Proposition 3.1 to complete the proof. \square

4.1 Comparison with Catoni's Bound

As we have shown in Proposition 3.1, using shifted Rademacher processes, we can match Catoni's fast-rate PAC-Bayesian bound (Theorem 2.2) up to constants. We have also presented a new fast-rate PAC-Bayes bound based on "flatness". Although both our bound and Catoni's bound show fast $\mathcal{O}(m^{-1})$ rates of convergence, our bound can exploit flatness in the posterior distribution.

In particular, our PAC-Bayes bound based on flatness (Eq. (15)) can be much tighter than Catoni's bound (Eq. (6)) when the posterior is chosen to concentrate on a "flat minimum" where $\frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2$ is very small yet $\mathcal{L}_S(Q)$ is nonzero. It can be verified that the "flatness" term $\frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2$ in Eq. (15) is smaller than the excess empirical risk term $c\mathcal{L}_S(Q)$ when $\frac{1-h^2}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2$ is greater than 0, which is precisely when the empirical risk is greater than zero. (See Eq. (14).)

Based on this observation, we expect our bound to be tighter for sufficient flat posteriors, nonzero empirical risk, and sufficient training data. In order to see this, note that Catoni's bound has the form $(1+c_c)\mathcal{L}_S Q + \frac{\mathcal{C}_c}{m}(\text{KL}(Q\|P) + \log \frac{1}{\delta})$, while our bound based on Eq. (14) can be written $(1+c_r)\mathcal{L}_S Q - \frac{c_r(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2 + \frac{\mathcal{C}_r}{m}(\text{KL}(Q\|P) + \log \frac{1}{\delta} + 1)$. Here c_c, c_r inflate the empirical risk and $\mathcal{C}_c, \mathcal{C}_r$ are constants. Let T_m be $\frac{c_r(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2$. Note that c_c and c_r must be fixed before seeing the data. Assuming we equate the inflation of the empirical risk, i.e., $c_c = c_r$, the proposed bound is tighter than Catoni's bound provided $m > \frac{1}{T_m} ((\mathcal{C}_r - \mathcal{C}_c)(\text{KL}(Q\|P) + \log \frac{1}{\delta}) + \mathcal{C}_r)$. If T_m converges to a positive number (a reasonable assumption), then our proposed bound will be tighter for sufficiently many samples. If we assume $c_c \neq c_r$, our bound can still be tighter than Catoni's bound under more involved conditions.

5 Related Work

There is a large literature on obtaining fast $1/m$ convergence rates for generalization error and excess risk using Rademacher processes and their generalizations [4, 22, 27, 29, 44]. As far as we know, this literature does not connect with the PAC-Bayesian literature. There do exist, however, PAC-Bayesian analyses for specific learning algorithms that achieve fast rates [2, 15, 24]. These specific analyses do not lead to general PAC-Bayes bounds, like those produced by Catoni [8].

Our new PAC-Bayes bound based on flatness bears a superficial resemblance to a number of bounds in the literature. However, our notion of flatness is *not* related to the variance of the randomized classifier caused by the randomness of the observed data. Therefore, our new bound is fundamentally different from existing PAC-Bayes bounds based on this type of variance [15, 24, 41].

For example, Tolstikhin and Seldin [41, Thm. 4] presents a generalization bound based on the "empirical variance", which is distinct from our "flatness". The "empirical variance" is $\mathbb{E}_Q \frac{1}{m} \sum_{i=1}^m [f(z_i) - \frac{1}{m} \sum_{i=1}^m f(z_i)]^2$, while our "flatness" is $\frac{1}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - \mathbb{E}_Q f(z_i)]^2$. Note that it is possible for flatness to be zero, even when empirical variance is large.

To the best of our knowledge, the closest work to ours in the literature is that by Audibert [2]. The bound given in [2, Thm. 6.1] uses a notion similar to our "flatness". The bound is, however, not comparable with ours for several reasons: First, [2, Theorem 6.1] holds only for the particular

algorithm proposed by Audibert, and so it is not a general PAC-Bayes bound like ours. Second, our notion of “flatness” is empirical, while the “flatness” term in [2, Theorem 6.1] is defined by an expectation over the data distribution, which is often presumed unknown. Finally, the proof techniques used to establish [2, Theorem 6.1] are specialized to the proposed algorithm and not based on the use of Rademacher processes. Our proof techniques via shifted Rademacher processes provides a blueprint for other approaches to deriving fast-rate PAC-Bayes bounds.

Grünwald and Mehta [17] establish new excess risk bounds in terms of a novel complexity measure based on “luckiness” functions. In the setting of randomized classifiers, particular choices of luckiness functions can be related to PAC-Bayesian notions of complexity based on “priors”. Indeed, in this setting, their complexity measure can be bounded in terms of a KL divergence, as in PAC-Bayesian bounds. In a setting with deterministic classifiers, the authors show that their complexity measure can be bounded in terms of Rademacher complexity. Thus, while their framework connects with both PAC-Bayesian and Rademacher-complexity bounds, it is not immediately clear whether it produces direct connections, as we have accomplished here. It is certainly interesting to consider whether our bounds can be achieved (or surpassed) by an appropriate use of their framework.

6 Conclusion

In this paper we exploit the connections between modern PAC-Bayesian theory and Rademacher complexities. Using shifted Rademacher processes [27, 43, 44], we derive a novel fast-rate PAC-Bayes bound that depends on the empirical “flatness” of the posterior. Our work provides new insights on PAC-Bayesian theory and opens up new avenues for developing stronger bounds.

It is worth highlighting some potentially interesting directions that may be worth further investigation:

We have “rederived” Catoni’s bound via shifted Rademacher processes, up to constants. It is interesting to ask whether the Rademacher approach can dominate the direct PAC-Bayes bound. In the other direction, we have not derived our flatness bound via a direct PAC-Bayes approach. Whether this is possible and what it achieves might shed light on the relative strengths of these two distinct approaches to PAC-Bayes bounds. It may also be interesting to pursue PAC-Bayes bounds via some adaptation of Talagrand’s concentration inequalities [42, Ch.3].

We have derived PAC-Bayes bounds for zero–one loss. While the extension to bounded loss is straightforward, the problem of extending our approach to unbounded loss relates to a growing body of work on this problem within the PAC-Bayesian framework. (See, for example, [1] and the references therein). Whether the Rademacher perspective is helpful or not in this regard is not clear at this point.

There has been a surge of interest in PAC-Bayes bounds and their application to the study of generalization in large-scale neural networks. One promising direction is to consider Rademacher-process techniques may aid in the development of PAC-Bayesian analyses of specific algorithms [2, 15, 24], especially in the case when the algorithms are related to large-scale neural networks trained by stochastic gradient descent [30, 36, 37].

It would be interesting to perform a careful empirical study of our flatness bound in the context of large-scale neural networks, in the vein of the work of Dziugaite and Roy [9]. Preliminary work suggests that the posteriors found by PAC-Bayes bound optimization are not flat in our sense. After some investigation, we believe the reason is that optimizing the PAC-Bayes bound results in underfitting, due in part to the distribution-independent prior. It would be interesting to compare various PAC-Bayes bounds under strict constraints on the empirical risk.

Acknowledgments

We would like to also thank Peter Bartlett, Gintare Karolina Dziugaite, Roger Grosse, Yasaman Mahdaviyeh, Zacharie Naulet, and Sasha Rakhlin for helpful discussions. In particular, the authors would like to thank Sasha Rakhlin for introducing us to the work of Kakade, Sridharan, and Tewari [21]. The work benefitted also from constructive feedback from anonymous referees. JY was supported by an Alexander Graham Bell Canada Graduate Scholarship (NSERC CGS D), Ontario Graduate Scholarship (OGS), and Queen Elizabeth II Graduate Scholarship in Science and Technology (QEII-GSST). SS was supported by a Borealis AI Global Fellowship Award, Connaught New

Researcher Award, and Connaught Fellowship. DMR was supported by an NSERC Discovery Grant and Ontario Early Researcher Award.

References

- [1] P. Alquier and B. Guedj. “Simpler PAC-Bayesian bounds for hostile data”. *Machine Learning* 107.5 (2018), pp. 887–902.
- [2] J.-Y. Audibert. “Fast learning rates in statistical inference through aggregation”. *The Annals of Statistics* 37.4 (2009), pp. 1591–1646.
- [3] J.-Y. Audibert and O. Bousquet. “Combining PAC-Bayesian and Generic Chaining Bounds”. *Journal of Machine Learning Research* 8 (2007), pp. 863–889.
- [4] P. L. Bartlett, O. Bousquet, and S. Mendelson. “Local Rademacher Complexities”. *The Annals of Statistics* 33.4 (2005), pp. 1497–1537.
- [5] P. L. Bartlett and S. Mendelson. “Rademacher and Gaussian complexities: Risk bounds and structural results”. *Journal of Machine Learning Research* 3 (2002), pp. 463–482.
- [6] L. Bégin, P. Germain, F. Laviolette, and J.-F. Roy. “PAC-Bayesian bounds based on the Rényi divergence”. In: *Artificial Intelligence and Statistics*. 2016, pp. 435–444.
- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [8] O. Catoni. *PAC-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*. Vol. 56. Lecture Notes – Monograph Series. Institute of Mathematical Statistics, 2007.
- [9] G. K. Dziugaite and D. M. Roy. “Computing Nonvacuous Generalization Bounds for Deep (Stochastic) Neural Networks with Many More Parameters than Training Data”. In: *Proceedings of the 33rd Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. 2017.
- [10] G. K. Dziugaite and D. M. Roy. “Data-dependent PAC-Bayes priors via differential privacy”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 8430–8441.
- [11] G. K. Dziugaite and D. M. Roy. “Entropy-SGD optimizes the prior of a PAC-Bayes bound: Generalization properties of Entropy-SGD and data-dependent priors”. In: *International Conference on Machine Learning*. 2018, pp. 1376–1385.
- [12] T. van Erven. *PAC-Bayes Mini-tutorial: A Continuous Union Bound*. 2014. arXiv: [1405.1580](#).
- [13] T. van Erven, P. D. Grünwald, N. A. Mehta, M. D. Reid, and R. C. Williamson. “Fast Rates in Statistical and Online Learning”. *Journal of Machine Learning Research* 16 (2015), pp. 1793–1861.
- [14] P. Germain, F. Bach, A. Lacoste, and S. Lacoste-Julien. “PAC-Bayesian theory meets Bayesian inference”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 1884–1892.
- [15] P. Germain, A. Lacasse, F. Laviolette, M. Marchand, and J.-F. Roy. “Risk bounds for the majority vote: From a PAC-Bayesian analysis to a learning algorithm”. *The Journal of Machine Learning Research* 16.1 (2015), pp. 787–860.
- [16] E. Giné and V. Koltchinskii. “Concentration inequalities and asymptotic results for ratio type empirical processes”. *The Annals of Probability* 34.3 (2006), pp. 1143–1216.
- [17] P. D. Grünwald and N. A. Mehta. “A tight excess risk bound via a unified PAC-Bayesian–Rademacher–Shtarkov–MDL complexity”. In: *Algorithmic Learning Theory*. 2019, pp. 433–465. arXiv: [1710.07732](#).
- [18] B. Guedj. *A primer on PAC-Bayesian learning*. 2019. arXiv: [1901.05353](#).
- [19] S. Hanneke. “Refined error bounds for several learning algorithms”. *The Journal of Machine Learning Research* 17.1 (2016), pp. 4667–4721.
- [20] S. Hanneke and L. Yang. “Minimax analysis of active learning”. *The Journal of Machine Learning Research* 16.1 (2015), pp. 3487–3602.

- [21] S. M. Kakade, K. Sridharan, and A. Tewari. “On the complexity of linear prediction: risk bounds, margin bounds, and regularization”. In: *Advances in Neural Information Processing Systems*. 2008, pp. 793–800.
- [22] V. Koltchinskii. “Local Rademacher complexities and oracle inequalities in risk minimization”. *The Annals of Statistics* 34.6 (2006), pp. 2593–2656.
- [23] V. Koltchinskii and D. Panchenko. “Empirical margin distributions and bounding the generalization error of combined classifiers”. *The Annals of Statistics* 30.1 (2002), pp. 1–50.
- [24] A. Lacasse, F. Laviolette, M. Marchand, P. Germain, and N. Usunier. “PAC-Bayes bounds for the risk of the majority vote and the variance of the Gibbs classifier”. In: *Advances in Neural Information Processing Systems*. 2007, pp. 769–776.
- [25] J. Langford. “Tutorial on practical prediction theory for classification”. *Journal of Machine Learning Research* 6.Mar (2005), pp. 273–306.
- [26] G. Lecué and S. Mendelson. *Learning subgaussian classes: Upper and minimax bounds*. 2013. arXiv: [1305.4825](#).
- [27] G. Lecué and C. Mitchell. “Oracle inequalities for cross-validation type procedures”. *Electronic Journal of Statistics* 6 (2012), pp. 1803–1837.
- [28] G. Lever, F. Laviolette, and J. Shawe-Taylor. “Tighter PAC-Bayes bounds through distribution-dependent priors”. *Theoretical Computer Science* 473 (2013), pp. 4–28.
- [29] T. Liang, A. Rakhlin, and K. Sridharan. “Learning with square loss: Localization through offset Rademacher complexity”. In: *Conference on Learning Theory*. 2015, pp. 1260–1285.
- [30] B. London. “A PAC-Bayesian analysis of randomized learning with application to stochastic gradient descent”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 2931–2940.
- [31] P. Massart and É. Nédélec. “Risk bounds for statistical learning”. *The Annals of Statistics* 34.5 (2006), pp. 2326–2366.
- [32] D. A. McAllester. *A PAC-Bayesian Tutorial with A Dropout Bound*. 2013. arXiv: [1307.2118](#).
- [33] D. A. McAllester. “PAC-Bayesian Model Averaging”. In: *Conference on Learning Theory*. 1999, pp. 164–170.
- [34] S. Mendelson. “Learning without concentration”. In: *Conference on Learning Theory*. 2014, pp. 25–39.
- [35] S. Mendelson. ““Local” vs. “global” parameters—breaking the Gaussian complexity barrier”. *The Annals of Statistics* 45.5 (2017), pp. 1835–1862.
- [36] B. Neyshabur, S. Bhojanapalli, D. McAllester, and N. Srebro. “Exploring generalization in deep learning”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 5947–5956.
- [37] B. Neyshabur, S. Bhojanapalli, and N. Srebro. *A PAC-Bayesian approach to spectrally-normalized margin bounds for neural networks*. 2017. arXiv: [1707.09564](#).
- [38] J. Shawe-Taylor and R. C. Williamson. “A PAC Analysis of a Bayesian Estimator”. In: *Conference on Learning Theory*. 1997, pp. 2–9.
- [39] S. L. Smith and Q. V. Le. *A Bayesian perspective on generalization and stochastic gradient descent*. 2017. arXiv: [1710.06451](#).
- [40] N. Thiemann, C. Igel, O. Wintenberger, and Y. Seldin. *A strongly quasiconvex PAC-Bayesian bound*. 2016. arXiv: [1608.05610](#).
- [41] I. O. Tolstikhin and Y. Seldin. “PAC-Bayes-empirical-Bernstein inequality”. In: *Advances in Neural Information Processing Systems*. 2013, pp. 109–117.
- [42] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.
- [43] M. Wegkamp. “Model selection in nonparametric regression”. *The Annals of Statistics* 31.1 (2003), pp. 252–273.
- [44] N. Zhivotovskiy and S. Hanneke. “Localization of VC classes: Beyond local Rademacher complexities”. *Theoretical Computer Science* 742 (2018), pp. 27–49.

A Proofs

A.1 Proof of Proposition 3.1

To match Catoni's bound, we need to control $\sup_{f \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(f) - (1+c)\hat{\mathcal{L}}_S(f)$, given $c > 0$. We apply an existing result due to Zhivotovskiy and Hanneke [44], which we quote here:

Lemma A.1 (Shifted Symmetrization in deviation [44, Cor. 7]). *Fix constants $c > c_2 > 0$,*

$$\begin{aligned} & \frac{1}{4} \mathbb{P}_S \left(\sup_{f \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(f) - (1+c)\hat{\mathcal{L}}_S(f) \geq t \right) \\ & \leq \mathbb{P}_{S,\varepsilon} \left(\sup_{f \in \mathcal{F}} \left[\frac{1+c'/2}{m} \sum_{i=1}^m \left(\varepsilon_i - \frac{c'/2}{1+c'/2} \right) f(z_i) \right] \geq \frac{t'}{2} \right). \end{aligned} \quad (26)$$

where $c' = \frac{c-c_2}{1+c_2}$, $t' = \frac{t}{2(1+c_2)}$, and $\{\varepsilon_i\}$ are Rademacher random variables.

Let $\varepsilon'_i := \varepsilon_i - \frac{c'}{2+c'}$. For $\kappa > 0$, define $\mathcal{Q}(\kappa) := \{Q : \text{KL}(Q||P) \leq \kappa\}$. Next, using convex conjugate of the Kullback–Leibler divergence (the change-measure inequality), for $\lambda > 0$,

$$\sup_{Q \in \mathcal{Q}(\kappa)} \left[\frac{1}{m} \sum_{i=1}^m \varepsilon'_i \mathbb{E}_Q f(z_i) - \frac{1}{\lambda} \text{KL}(Q||P) \right] \leq \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right], \quad (27)$$

which implies

$$\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \varepsilon'_i \mathbb{E}_Q f(z_i) \leq \frac{\kappa}{\lambda} + \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right]. \quad (28)$$

Therefore, we have

$$\mathbb{P}_S \left(\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q [\mathcal{L}_{\mathcal{D}}(f) - (1+c)\hat{\mathcal{L}}_S(f)] \geq t \right) \quad (29)$$

$$= \mathbb{P}_S \left(\sup_{Q \in \mathcal{Q}(\kappa)} \mathcal{L}_{\mathcal{D}}(\mathbb{E}_Q f) - (1+c)\hat{\mathcal{L}}_S(\mathbb{E}_Q f) \geq t \right) \quad (30)$$

$$\leq 4 \mathbb{P}_{S,\varepsilon} \left(\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \left(\varepsilon_i - \frac{c'}{2+c'} \right) \mathbb{E}_Q f(z_i) \geq \frac{t'}{2+c'} \right) \quad (31)$$

$$\leq 4 \mathbb{P}_{S,\varepsilon} \left(\frac{\kappa}{\lambda} + \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right] \geq \frac{t'}{2+c'} \right) \quad (32)$$

$$= 4 \mathbb{P}_{S,\varepsilon} \left(\log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right] \geq \frac{\lambda t'}{2+c'} - \kappa \right) \quad (33)$$

$$= 4 \mathbb{P}_{S,\varepsilon} \left(\mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right] \geq \exp \left(\frac{\lambda t'}{2+c'} - \kappa \right) \right) \quad (34)$$

$$\stackrel{\text{Markov}}{\leq} 4 \exp \left(\kappa - \frac{\lambda t'}{2+c'} \right) \mathbb{E}_S \mathbb{E}_{\varepsilon} \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m \varepsilon'_i f(z_i) \right) \right]. \quad (35)$$

Now we use the result of Lemma A.2. For any c_2 such that $0 < c_2 < c$, if $t \geq \frac{1}{m} \frac{(1+c_2)^2}{c_2}$ and $\frac{c'}{c'+2} \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}$ then we have

$$\mathbb{P} \left(\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q [\mathcal{L}_{\mathcal{D}}(f) - (1+c)\hat{\mathcal{L}}_S(f)] \geq t \right) \leq 4 \exp \left(\kappa - \frac{\lambda t'}{2+c'} \right), \quad (36)$$

where $c' = \frac{c-c_2}{1+c_2}$ and $t' = \frac{t}{2(1+c_2)}$. Now letting $4 \exp\left(\kappa - \frac{\lambda t'}{2+c'}\right)$ equals to δ , we have

$$t = 2(1+c_2)t' = 2(1+c_2) \left[\frac{2+c'}{\lambda} (\kappa + \log(4/\delta)) \right]. \quad (37)$$

Now let $\frac{\lambda}{m} = C$, noting that

$$t \geq 2(1+c_2) \left[\frac{2+c'}{\lambda} \log(4/\delta) \right] \quad (38)$$

then we can choose C small enough (clearly bounded away from 0) to satisfy

$$\frac{\log \cosh(C)}{C} \leq \frac{c'}{c'+2}, \quad C \leq \frac{2(1+c_2)(2+c') \log(4/\delta)}{(1+c_2)^2/c_2} \quad (39)$$

which guarantees both $t \geq \frac{1}{m} \frac{(1+c_2)^2}{c_2}$ and $\frac{c'}{c'+2} \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}$.

Therefore, using such C we have

$$\mathbb{P} \left(\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q[\mathcal{L}_Q(f) - (1+c)\mathcal{L}_S(f)] \geq \frac{2(1+c_2)(2+c')}{Cm} (\kappa + \log(4/\delta)) \right) \leq \delta, \quad (40)$$

which implies

$$\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q[\mathcal{L}_Q(f) - (1+c)\mathcal{L}_S(f)] \leq \frac{C'}{m} (\kappa + \log(4/\delta)), \quad \text{w.p. } 1 - \delta. \quad (41)$$

where $C' = \frac{2(1+c_2)(2+c')}{C}$.

Finally, we combine all possible κ using a union bound. Define $\Gamma_0 := \{Q : \text{KL}(Q||P) \leq 2\}$ and $\Gamma_j := \{Q : \text{KL}(Q||P) \in [2^j, 2^{j+1}]\}$ for $j \geq 1$. Let $\delta_j = 2^{-(j+1)}\delta$ so $\sum_{j=0}^{\infty} \delta_j = \delta$. Then for any Q there is a j_Q such that $Q \in \Gamma_{j_Q}$. Then by definition we have

$$\text{KL}(Q||P) \leq 2^{j_Q+1} \leq 2 \max\{\text{KL}(Q||P), 1\} \quad (42)$$

$$\delta_{j_Q} = 2^{-(j_Q+1)}\delta \geq \frac{\delta}{2 \max\{\text{KL}(Q||P), 1\}}. \quad (43)$$

Therefore, we have that for any Q , with probability $1 - \delta$, we have

$$\mathcal{L}_Q(Q) - (1+c)\mathcal{L}_S(Q) \leq \frac{C'}{m} (2 \max\{\text{KL}(Q||P), 1\} + \log(8 \max\{\text{KL}(Q||P), 1\}/\delta)) \quad (44)$$

$$\leq \frac{C'}{m} (2 \max\{\text{KL}(Q||P), 1\} + \log(\max\{\text{KL}(Q||P), 1\}) + \log(8/\delta)). \quad (45)$$

Now we simplify the order without optimizing the constants, which gives

$$\log(\max\{\text{KL}(Q||P), 1\}) \leq \max\{\text{KL}(Q||P), 1\} \leq \text{KL}(Q||P) + 1. \quad (46)$$

Therefore, we have

$$\mathcal{L}_Q(f) - (1+c)\mathcal{L}_S(f) \leq \frac{C_1}{m} \text{KL}(Q||P) + \frac{C_2}{m} \log(1/\delta) + \frac{C_3}{m}, \quad \text{w.p. } 1 - \delta. \quad (47)$$

Lemma A.2. If $k \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}$, then

$$\mathbb{E}_P \mathbb{E}_S \mathbb{E}_\epsilon \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m (\epsilon_i - k) f(z_i) \right) \right] \leq 1. \quad (48)$$

Proof. Noting that $\{f(z_i)\}$ are independent Bernoulli random variables, we have

$$\mathbb{E}_P \mathbb{E}_S \mathbb{E}_\varepsilon \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m (\varepsilon_i - k) f(z_i) \right) \right] \quad (49)$$

$$\stackrel{\text{indep}}{=} \mathbb{E}_P \prod_{i=1}^m \mathbb{E}_S \mathbb{E}_{\varepsilon_i} \left[\exp \left(\frac{\lambda}{m} (\varepsilon_i - k) f(z_i) \right) \right] \quad (50)$$

$$\stackrel{\text{Rademacher}}{=} \mathbb{E}_P \prod_{i=1}^m \mathbb{E}_S \left[\frac{e^{(1-k)\frac{\lambda}{m}f(z_i)} + e^{-(1+k)\frac{\lambda}{m}f(z_i)}}{2} \right] \quad (51)$$

$$\stackrel{\text{Bernoulli}}{=} \mathbb{E}_P \prod_{i=1}^m \left[(1 - \mathbb{E}_S[f(z_i)]) + \left(\frac{e^{\frac{\lambda}{m}} + e^{-\frac{\lambda}{m}}}{2e^{k\frac{\lambda}{m}}} \right) \mathbb{E}_S[f(z_i)] \right] \quad (52)$$

$$= \mathbb{E}_P \prod_{i=1}^m \left[(1 - \mathbb{E}_S[f(z_i)]) + \frac{\cosh\left(\frac{\lambda}{m}\right)}{e^{k\frac{\lambda}{m}}} \mathbb{E}_S[f(z_i)] \right], \quad (53)$$

which is upper bounded by 1 if we choose k such that

$$\cosh(\lambda/m) = \frac{e^{\lambda/m} + e^{-\lambda/m}}{2} \leq e^{k\lambda/m}. \quad (54)$$

That is

$$k \geq \frac{\log \cosh(\lambda/m)}{\lambda/m}. \quad (55)$$

□

A.2 Proof of Theorem 4.3

We first present some lemmas that will be used in the later proof.

Lemma A.3 (Shifted-Flatness Inequality). *Consider a function $f : \mathcal{X} \rightarrow [0, 1]$, constants $h \in [0, 1]$ and $c_2 > 0$, if $t \geq \frac{(1+c_2)(1+c_2h^2)}{mc_2h^2}$, we have*

$$\mathbb{P}_S \left(\mathcal{L}_{\mathcal{D}}(f) - (1+c_2)\hat{\mathcal{L}}_S(f) + c_2(1-h^2)\hat{\mathcal{L}}_S(f^2) \geq \frac{t}{2} \right) \leq \frac{1}{2}. \quad (56)$$

Proof. Let $v = c_2\mathcal{L}_{\mathcal{D}}(f) - c_2(1-h^2)\mathcal{L}_{\mathcal{D}}(f^2) = c_2\mathcal{L}_{\mathcal{D}}(f - (1-h^2)f^2)$. Then, we have

$$\mathbb{P} \left(\mathcal{L}_{\mathcal{D}}(f) - (1+c_2)\hat{\mathcal{L}}_S(f) + c_2(1-h^2)\hat{\mathcal{L}}_S(f^2) \geq \frac{t}{2} \right) \quad (57)$$

$$= \mathbb{P} \left(\mathcal{L}_{\mathcal{D}}(f) - \frac{c_2(1-h^2)}{1+c_2}\mathcal{L}_{\mathcal{D}}(f^2) - \hat{\mathcal{L}}_S(f) + \frac{c_2(1-h^2)}{1+c_2}\hat{\mathcal{L}}_S(f^2) \geq \frac{t/2+v}{1+c_2} \right) \quad (58)$$

$$= \mathbb{P} \left(\mathbb{E}_{z \sim \mathcal{D}}(f(z) - \frac{c_2(1-h^2)}{1+c_2}f(z)^2) - \frac{1}{m} \sum_{i=1}^m (f(z_i) - \frac{c_2(1-h^2)}{1+c_2}f(z_i)^2) \geq \frac{t/2+v}{1+c_2} \right). \quad (59)$$

Because $f(z_i) - \frac{c_2(1-h^2)}{1+c_2}f(z_i)^2, i = 1, \dots, m$ are i.i.d. random samples, using Chebyshev's inequality together with $4ab \leq (a+b)^2$ and $f \in [0, 1]$, the formula above is upper bounded by

$$\frac{(1+c_2)^2 \text{Var}(f - \frac{c_2(1-h^2)}{1+c_2}f^2)}{m(t/2+v)^2} \leq \frac{(1+c_2)^2 \mathcal{L}_{\mathcal{D}}(f - \frac{c_2(1-h^2)}{1+c_2}f^2)^2}{2mvt} \leq \frac{(1+c_2)^2 \mathcal{L}_{\mathcal{D}}(f - \frac{c_2(1-h^2)}{1+c_2}f^2)}{2mvt}. \quad (60)$$

We can further decompose the term in the numerator by

$$\mathcal{L}_{\mathcal{D}}(f - \frac{c_2(1-h^2)}{1+c_2}f^2) = \frac{c_2}{1+c_2} \mathcal{L}_{\mathcal{D}}(f - (1-h^2)f^2) + \frac{1}{1+c_2} \mathcal{L}_{\mathcal{D}}(f) \quad (61)$$

$$\leq \frac{1}{1+c_2}v + \frac{1}{1+c_2} \frac{1}{c_2h^2}v = \frac{c_2h^2+1}{(1+c_2)c_2h^2}v, \quad (62)$$

Therefore the lemma follows directly from $t \geq \frac{(1+c_2)(1+c_2h^2)}{mc_2h^2}$. □

Lemma A.4 (New Shifted Symmetrization in Deviation). *Fix constants $c > c_2 > 0$, $h \in [0, 1]$, if $t \geq \frac{(1+c_2)(1+c_2h^2)}{mc_2h^2}$, we have*

$$\begin{aligned} & \frac{1}{4} \mathbb{P}_S \left(\sup_{f \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(f) - (1+c)\mathcal{L}_S(f) + c(1-h^2)\mathcal{L}_S(f^2) \geq t \right) \\ & \leq \mathbb{P}_{S,\epsilon} \left(\sup_{f \in \mathcal{F}} \left[\frac{1}{m} \sum_{i=1}^m \epsilon_i ((1+c')f(z_i) - c'(1-h^2)f^2(z_i)) - c''\mathcal{L}_S(f - (1-h^2)f^2) \right] \geq \frac{t}{4} \right), \end{aligned} \quad (63)$$

where $c' = \frac{c+c_2}{2}$, $c'' = \frac{c-c_2}{2}$, $\epsilon := \{\epsilon_i\}_{i=1}^m$, in which $\{\epsilon_i\}$ are independent Rademacher random variables.

Proof. Consider a random set $S' = \{z'_i\}_{i=1}^m \in \mathcal{D}^m$, in Lemma A.3 we have shown that

$$\mathbb{P}_{S'} \left(\mathcal{L}_{\mathcal{D}}(f) - (1+c_2)\mathcal{L}_{S'}(f) + c_2(1-h^2)\mathcal{L}_{S'}(f^2) \geq \frac{t}{2} \right) \leq \frac{1}{2}. \quad (64)$$

Therefore, we can get

$$\frac{1}{4} \mathbb{P}_S \left(\sup_{f \in \mathcal{F}} \mathcal{L}_{\mathcal{D}}(f) - (1+c)\mathcal{L}_S(f) + c(1-h^2)\mathcal{L}_S(f^2) \geq t \right) \quad (65)$$

$$\leq \frac{1}{2} \mathbb{P}_{S,S'} \left(\sup_{f \in \mathcal{F}} (1+c_2)\mathcal{L}_{S'}(f) - c_2(1-h^2)\mathcal{L}_{S'}(f^2) - (1+c)\mathcal{L}_S(f) + c(1-h^2)\mathcal{L}_S(f^2) \geq \frac{t}{2} \right) \quad (66)$$

$$= \frac{1}{2} \mathbb{P}_{S,S'} \left[\sup_{f \in \mathcal{F}} \left(1 + \frac{c+c_2}{2} \right) (\mathcal{L}_S(f) - \mathcal{L}_{S'}(f)) - \frac{c+c_2}{2} (1-h^2) (\mathcal{L}_S(f^2) - \mathcal{L}_{S'}(f^2)) \right] \quad (67)$$

$$- \frac{c-c_2}{2} \mathcal{L}_{S'}(f - (1-h^2)f^2) - \frac{c-c_2}{2} \mathcal{L}_S(f - (1-h^2)f^2) \geq \frac{t}{2} \quad (68)$$

$$\leq \mathbb{P}_{S,\epsilon} \left(\sup_{f \in \mathcal{F}} \left[\frac{1}{m} \sum_{i=1}^m \epsilon_i ((1+c')f(z_i) - c'(1-h^2)f^2(z_i)) - c''\mathcal{L}_S(f - (1-h^2)f^2) \right] \geq \frac{t}{4} \right), \quad (69)$$

where $c' = \frac{c+c_2}{2}$, $c'' = \frac{c-c_2}{2}$, and the last inequality is by the symmetrization argument. \square

Lemma A.5. *For constants $h \in (0, 1]$, $h^2c > c_2 > 0$, let $C = \frac{h^2c-c_2}{2(1+h^2c)(1+c_2)}$, if $0 < \frac{\lambda}{m} < C$, then given independent Bernoulli random variables X_i , for any random variables $Y_i \in [0, 1]$,*

$$\mathbb{E}_{\epsilon} \mathbb{E}_X \mathbb{E}_{Y|X,\epsilon} \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m X_i [(\epsilon_i + \epsilon_i'') - \epsilon_i''(1-h^2)Y_i] \right) \right] \leq 1, \quad (70)$$

where $\{\epsilon_i\}$ are i.i.d. Rademacher random variables and $\epsilon_i'' = \epsilon_i \frac{c+c_2}{2} - \frac{c-c_2}{2}$.

Proof. Note when $X_i = 0$, the value of Y_i has no effect onto LHS. When $X_i = 1$,

$$(\epsilon_i + \epsilon_i'') - \epsilon_i''(1-h^2)Y_i = (1+c_2) - c_2(1-h^2)Y_i, \text{ if } \epsilon_i = 1, \quad (71)$$

$$(\epsilon_i + \epsilon_i'') - \epsilon_i''(1-h^2)Y_i = -(1+c) + c(1-h^2)Y_i, \text{ if } \epsilon_i = -1, \quad (72)$$

Therefore, we have

$$(\epsilon_i + \epsilon_i'') - \epsilon_i''(1-h^2)Y_i \leq (\epsilon_i + \epsilon_i'') - \epsilon_i''(1-h^2) \frac{1-\epsilon_i}{2}. \quad (73)$$

Denoting $\mu_i = \mathbb{E}[X_i]$, by the monotonicity of the exponential function, we have

$$\mathbb{E}_\varepsilon \mathbb{E}_X \mathbb{E}_{Y|X,\varepsilon} \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m X_i [(\varepsilon_i + \varepsilon_i'') - \varepsilon_i''(1-h^2)Y_i] \right) \right] \quad (74)$$

$$\leq \mathbb{E}_\varepsilon \mathbb{E}_X \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m X_i \left[(\varepsilon_i + \varepsilon_i'') - \varepsilon_i''(1-h^2) \frac{1-\varepsilon_i}{2} \right] \right) \right] \quad (75)$$

$$= \prod_{i=1}^m \left[1 - \mu_i + \frac{\mu_i}{2} \exp \left(\frac{\lambda}{m} (1+c_2) \right) + \frac{\mu_i}{2} \exp \left(-\frac{\lambda}{m} (1+h^2c) \right) \right], \quad (76)$$

For the formula upper bounded by 1, it is sufficient to prove

$$\exp \left(\frac{\lambda}{m} (1+c_2) \right) + \exp \left(-\frac{\lambda}{m} (1+h^2c) \right) \leq 2. \quad (77)$$

Because we have $e^x \geq x+1$, thus $e^{-x} \leq \frac{1}{1+x}$ for $x > -1$ and $e^x \leq \frac{1}{1-x}$ for $x < 1$. Therefore, it is sufficient to have

$$\frac{1}{1 - \frac{\lambda}{m} (1+c_2)} + \frac{1}{1 + \frac{\lambda}{m} (1+h^2c)} \leq 2, \quad \frac{\lambda}{m} (1+c_2) \leq 1. \quad (78)$$

Thus, we know the argument holds for

$$\frac{\lambda}{m} \leq \min \left(\frac{h^2c - c_2}{2(1+h^2c)(1+c_2)}, \frac{1}{1+c_2} \right) = \frac{h^2c - c_2}{2(1+h^2c)(1+c_2)}, \quad (79)$$

where the last equality holds when $c_2 < h^2c$. \square

Now we are ready for the proof of Theorem 4.3. Denoting $g(\cdot) = \mathbb{E}_Q f(\cdot)$, one can write

$$\mathcal{L}_\mathcal{Q}(Q) - \hat{\mathcal{L}}_S(Q) - \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h)\mathbb{E}_Q f(z_i)]^2 \quad (80)$$

$$= \mathcal{L}_\mathcal{Q}(Q) - (1+c)\hat{\mathcal{L}}_S(Q) + \frac{c(1-h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2 \quad (81)$$

$$= \mathcal{L}_\mathcal{Q}(g) - (1+c)\hat{\mathcal{L}}_S(g) + \frac{c(1-h^2)}{m} \sum_{i=1}^m (g(z_i))^2 \quad (82)$$

$$= \mathcal{L}_\mathcal{Q}(g) - (1+c)\hat{\mathcal{L}}_S(g) + c(1-h^2)\hat{\mathcal{L}}_S(g^2). \quad (83)$$

Recall that for $\kappa > 0$, we have defined $\mathcal{Q}(\kappa) = \{Q : \text{KL}(Q||P) \leq \kappa\}$. We start from the formula,

$$\mathbb{P}_S \left[\sup_{Q \in \mathcal{Q}(\kappa)} \mathcal{L}_\mathcal{Q}(g) - (1+c)\hat{\mathcal{L}}_S(g) + c(1-h^2)\hat{\mathcal{L}}_S(g^2) \geq t \right]. \quad (84)$$

By Lemma A.4, let $c' = \frac{c+c_2}{2}$, $c'' = \frac{c-c_2}{2}$, and $\{\varepsilon_i\}$ being i.i.d. Rademacher random variables, we have

$$\mathbb{P}_S \left[\sup_{Q \in \mathcal{Q}(\kappa)} \mathcal{L}_\mathcal{Q}(g) - (1+c)\hat{\mathcal{L}}_S(g) + c(1-h^2)\hat{\mathcal{L}}_S(g^2) \geq t \right] \quad (85)$$

$$\leq 4\mathbb{P}_{S,\varepsilon} \left[\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \varepsilon_i ((1+c')g(z_i) - c'(1-h^2)g(z_i)^2) - c''\hat{\mathcal{L}}_S(g - (1-h^2)g^2) \geq \frac{t}{4} \right]. \quad (86)$$

Plugging into $g = \mathbb{E}_Q f(\cdot)$ yields

$$\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m \varepsilon_i ((1+c')g(z_i) - c'(1-h^2)g(z_i)^2) - c''\hat{\mathcal{L}}_S(g - (1-h^2)g^2) \quad (87)$$

$$= \sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m [\varepsilon_i (1+c') - c''] \mathbb{E}_Q f(z_i) - [\varepsilon_i c' - c''] (1-h^2) [\mathbb{E}_Q f(z_i)]^2. \quad (88)$$

Given $\epsilon = \{\epsilon_i\}_{i=1}^m, \mathbf{z} = \{z_i\}_{i=1}^m$, we suppose $\hat{Q}(\epsilon, \mathbf{z})$ achieves the supremum above (if the supremum cannot be achieved, one can use a sequence of $\{\hat{Q}_i(\epsilon, \mathbf{z})\}$ to approximate arbitrarily close to the supremum). Using

$$\epsilon_i'' := \epsilon_i c' - c'' = \epsilon_i \frac{c_1 + c_2}{2} - \frac{c_1 - c_2}{2}, \quad (89)$$

we have

$$\sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m [\epsilon_i (1 + c') - c''] \mathbb{E}_Q f(z_i) - [\epsilon_i c' - c''] (1 - h^2) [\mathbb{E}_Q f(z_i)]^2 \quad (90)$$

$$\leq \sup_{Q \in \mathcal{Q}(\kappa)} \frac{1}{m} \sum_{i=1}^m [\epsilon_i (1 + c') - c''] \mathbb{E}_Q f(z_i) - [\epsilon_i c' - c''] (1 - h^2) \mathbb{E}_Q f(z_i) \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i) \quad (91)$$

$$= \sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q \left[\frac{1}{m} \sum_{i=1}^m (\epsilon_i + \epsilon_i'') f(z_i) - \epsilon_i'' (1 - h^2) f(z_i) \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i) \right] \quad (92)$$

$$\leq \frac{\kappa}{\lambda} + \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m f(z_i) \left[(\epsilon_i + \epsilon_i'') - \epsilon_i'' (1 - h^2) \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i) \right] \right) \right], \quad (93)$$

where the last inequality follows from duality for convex optimization [7, Chp. 5].

Therefore, we have

$$\mathbb{P} \left[\sup_{Q \in \mathcal{Q}(\kappa)} \mathcal{L}_{\mathcal{Q}}(Q) - (1 + c) \mathcal{L}_S(Q) + \frac{c(1 - h^2)}{m} \sum_{i=1}^m (\mathbb{E}_Q f(z_i))^2 \geq t \right] \quad (94)$$

$$\leq 4\mathbb{P} \left[\frac{\kappa}{\lambda} + \frac{1}{\lambda} \log \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m f(z_i) \left[(\epsilon_i + \epsilon_i'') - \epsilon_i'' (1 - h^2) \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i) \right] \right) \right] \geq \frac{t}{4} \right] \quad (95)$$

$$\leq 4 \exp \left(\kappa - \frac{\lambda t}{4} \right) \mathbb{E}_S \mathbb{E}_{\epsilon} \mathbb{E}_P \left[\exp \left(\frac{\lambda}{m} \sum_{i=1}^m f(z_i) \left[(\epsilon_i + \epsilon_i'') - \epsilon_i'' (1 - h^2) \mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i) \right] \right) \right] \quad (96)$$

$$\leq 4 \exp \left(\kappa - \frac{\lambda t}{4} \right), \quad (97)$$

where the last inequality comes from Lemma A.5 by considering X_i as $f(z_i)$ and Y_i as $\mathbb{E}_{\hat{Q}(\epsilon, \mathbf{z})} f(z_i)$, with

$$C := \frac{\lambda}{m} \leq \frac{h^2 c - c_2}{2(1 + h^2 c)(1 + c_2)}. \quad (98)$$

Now let $4 \exp \left(\kappa - \frac{\lambda t}{4} \right)$ equals to δ , we have

$$t = \frac{4}{\lambda} (\kappa + \log \frac{4}{\delta}) = \frac{4}{Cm} (\kappa + \log \frac{4}{\delta}). \quad (99)$$

Note that the shifted symmetrization inequality requires $t \geq \frac{(1+c_2)(1+c_2 h^2)}{m c_2 h^2}$ by Lemma A.3. Combining with the previous requirement for C together, we have

$$C \leq \min \left(\frac{h^2 c - c_2}{2(1 + h^2 c)(1 + c_2)}, \frac{4 c_2 h^2}{(1 + c_2)(1 + c_2 h^2)} (\kappa + \log \frac{4}{\delta}) \right). \quad (100)$$

Using such C we have with probability at least $1 - \delta$,

$$\sup_{Q \in \mathcal{Q}(\kappa)} \mathbb{E}_Q \mathcal{L}_{\mathcal{Q}}(f) - \mathbb{E}_Q \hat{\mathcal{L}}_S(f) - \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1 + h) \mathbb{E}_Q f(z_i)]^2 \leq \frac{4}{Cm} (\kappa + \log \frac{4}{\delta}). \quad (101)$$

Finally we combine all possible κ using a union bound. Define $\Gamma_0 = \{Q : \text{KL}(Q||P) \leq 2\}$ and $\Gamma_j = \{Q : \text{KL}(Q||P) \in [2^j, 2^{j+1}]\}$ for $j \geq 1$. Let $\delta_j = 2^{-(j+1)} \delta$ so that $\sum_{j=0}^{\infty} \delta_j = \delta$. Then for any Q

there is a j_Q such that $Q \in \Gamma_{j_Q}$. Then we have

$$\text{KL}(Q||P) \leq 2^{j_Q+1} \leq 2 \max(\text{KL}(Q||P), 1) \quad (102)$$

$$\delta_{j_Q} = 2^{-(j_Q+1)} \delta \geq \frac{\delta}{2 \max(\text{KL}(Q||P), 1)}, \quad (103)$$

Therefore with probability at least $1 - \delta$ over draws of S , for any Q ,

$$\begin{aligned} \mathbb{E}_Q \mathcal{L}_{\mathcal{D}}(f) &\leq \mathbb{E}_Q \hat{\mathcal{L}}_S(f) + \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h) \mathbb{E}_Q f(z_i)]^2 \\ &\quad + \frac{4}{Cm} \left[2 \max(\text{KL}(Q||P), 1) + \log \frac{8 \max(\text{KL}(Q||P), 1)}{\delta} \right] \end{aligned} \quad (104)$$

$$\leq \mathbb{E}_Q \hat{\mathcal{L}}_S(f) + \frac{c}{m} \sum_{i=1}^m \mathbb{E}_Q [f(z_i) - (1+h) \mathbb{E}_Q f(z_i)]^2 + \frac{4}{Cm} \left[3 \text{KL}(Q||P) + \log \frac{1}{\delta} + 5 \right], \quad (105)$$

provided that

$$C \leq \min\left(\frac{h^2 c - c_2}{2(1+h^2 c)(1+c_2)}, \frac{4c_2 h^2}{(1+c_2)(1+c_2 h^2)} (\kappa + \log \frac{4}{\delta})\right). \quad (106)$$

Therefore, it is sufficient if

$$C = \frac{2h^4 c}{1+16h^2 c}, \quad c_2 = \frac{h^2 c}{1+16h^2 c}. \quad (107)$$