
Split-kl and PAC-Bayes-split-kl Inequalities for Ternary Random Variables

Yi-Shan Wu
University of Copenhagen
yswu@di.ku.dk

Yevgeny Seldin
University of Copenhagen
seldin@di.ku.dk

Abstract

We present a new concentration of measure inequality for sums of independent bounded random variables, which we name a split-kl inequality. The inequality is particularly well-suited for ternary random variables, which naturally show up in a variety of problems, including analysis of excess losses in classification, analysis of weighted majority votes, and learning with abstention. We demonstrate that for ternary random variables the inequality is simultaneously competitive with the kl inequality, the Empirical Bernstein inequality, and the Unexpected Bernstein inequality, and in certain regimes outperforms all of them. It resolves an open question by Tolstikhin and Seldin [2013] and Mhammedi et al. [2019] on how to match simultaneously the combinatorial power of the kl inequality when the distribution happens to be close to binary and the power of Bernstein inequalities to exploit low variance when the probability mass is concentrated on the middle value. We also derive a PAC-Bayes-split-kl inequality and compare it with the PAC-Bayes-kl, PAC-Bayes-Empirical-Bennett, and PAC-Bayes-Unexpected-Bernstein inequalities in an analysis of excess losses and in an analysis of a weighted majority vote for several UCI datasets. Last, but not least, our study provides the first direct comparison of the Empirical Bernstein and Unexpected Bernstein inequalities and their PAC-Bayes extensions.

1 Introduction

Concentration of measure inequalities for sums of independent random variables are the most fundamental analysis tools in statistics and many other domains [Boucheron et al., 2013]. Their history stretches almost a century back, and inequalities such as Hoeffding’s [Hoeffding, 1963] and Bernstein’s [Bernstein, 1946] are the main work horses of learning theory.

For binary random variables, one of the tightest concentration of measure inequalities is the kl inequality [Maurer, 2004, Langford, 2005, Foong et al., 2021, 2022], which is based on combinatorial properties of a sum of n independent random variables.¹ However, while being extremely tight for binary random variables and applicable to any bounded random variables, the kl inequality is not necessarily a good choice for sums of bounded random variables that can take more than two values. In the latter case, the Empirical Bernstein [Mnih et al., 2008, Audibert et al., 2009, Maurer and Pontil, 2009] and the Unexpected Bernstein [Cesa-Bianchi et al., 2007, Mhammedi et al., 2019] inequalities can be significantly tighter due to their ability to exploit low variance, as shown by Tolstikhin and Seldin [2013]. However, the Empirical and Unexpected Bernstein inequalities are loose for binary random variables [Tolstikhin and Seldin, 2013].

¹The Binomial tail bound is slightly tighter, but it does not extend to the PAC-Bayes setting [Langford, 2005]. Our split-kl approach can be directly applied to obtain a “split-Binomial-tail” inequality.

The challenge of exploiting low variance and, at the same time, matching the tightness of the kl inequality if a distribution happens to be close to binary, was faced by multiple prior works [Tolstikhin and Seldin, 2013, Mhammedi et al., 2019, Wu et al., 2021], but remained an open question. We resolve this question for the case of ternary random variables. Such random variables appear in a variety of applications, and we illustrate two of them. One is a study of excess losses, which are differences between the zero-one losses of a prediction rule h and a reference prediction rule h^* , $Z = \ell(h(X), Y) - \ell(h^*(X), Y) \in \{-1, 0, 1\}$. Mhammedi et al. [2019] have applied the PAC-Bayes-Unexpected-Bernstein bound to excess losses in order to improve generalization bounds for classification. Another example of ternary random variables is the tandem loss with an offset, defined by $\ell_\alpha(h(X), h'(X), Y) = (\ell(h(X), Y) - \alpha)(\ell(h'(X), Y) - \alpha) \in \{\alpha^2, -\alpha(1 - \alpha), (1 - \alpha)^2\}$. Wu et al. [2021] have applied the PAC-Bayes-Empirical-Bennett inequality to the tandem loss with an offset to obtain a generalization bound for the weighted majority vote. Yet another potential application, which we leave for future work, is learning with abstention [Cortes et al., 2018, Thulasidasan et al., 2019].

We present the split-kl inequality, which simultaneously matches the tightness of the Empirical/Unexpected Bernstein and the kl, and outperforms both for certain distributions. It works for sums of any bounded random variables Z_1, \dots, Z_n , not only the ternary ones, but it is best suited for ternary random variables, for which it is almost tight (in the same sense, as the kl is tight for binary random variables). The idea behind the split-kl inequality is to write a random variable Z as $Z = \mu + Z^+ - Z^-$, where μ is a constant, $Z^+ = \max\{0, Z - \mu\}$, and $Z^- = \max\{0, \mu - Z\}$. Then $\mathbb{E}[Z] = \mu + \mathbb{E}[Z^+] - \mathbb{E}[Z^-]$ and, given an i.i.d. sample Z_1, \dots, Z_n , we can bound the distance between $\frac{1}{n} \sum_{i=1}^n Z_i$ and $\mathbb{E}[Z]$ by using kl upper and lower bounds on the distances between $\frac{1}{n} \sum_{i=1}^n Z_i^+$ and $\mathbb{E}[Z^+]$, and $\frac{1}{n} \sum_{i=1}^n Z_i^-$ and $\mathbb{E}[Z^-]$, respectively. For ternary random variables $Z \in \{a, b, c\}$ with $a \leq b \leq c$, the best split is to take $\mu = b$, then both Z^+ and Z^- are binary and the kl upper and lower bounds for their rescaled versions are tight and, therefore, the split-kl inequality for Z is also tight. Thus, this approach provides the best of both worlds: the combinatorial tightness of the kl bound and exploitation of low variance when the probability mass on the middle value happens to be large, as in Empirical Bernstein inequalities. We further elevate the idea to the PAC-Bayes domain and derive a PAC-Bayes-split-kl inequality.

We present an extensive set of experiments, where we first compare the kl, Empirical Bernstein, Unexpected Bernstein, and split-kl inequalities applied to (individual) sums of independent random variables in simulated data, and then compare the PAC-Bayes-kl, PAC-Bayes-Unexpected-Bernstein, PAC-Bayes-split-kl, and, in some of the setups, PAC-Bayes-Empirical-Bennett, for several prediction models on several UCI datasets. In particular, we evaluate the bounds in the linear classification setup studied by Mhammedi et al. [2019] and in the weighted majority prediction setup studied by Wu et al. [2021]. To the best of our knowledge, this is also the first time when the Empirical Bernstein and the Unexpected Bernstein inequalities are directly compared, with and without the PAC-Bayesian extension. In Appendix A.2 we also show that an inequality introduced by Cesa-Bianchi et al. [2007] yields a relaxation of the Unexpected Bernstein inequality by Mhammedi et al. [2019].

2 Concentration of Measure Inequalities for Sums of Independent Random Variables

We start with the most basic question in probability theory and statistics: how far can an average of an i.i.d. sample Z_1, \dots, Z_n deviate from its expectation? We cite the major existing inequalities, the kl, Empirical Bernstein, and Unexpected Bernstein, then derive the new split-kl inequality, and then provide a numerical comparison.

2.1 Background

We use $\text{KL}(\rho||\pi)$ to denote the Kullback-Leibler divergence between two probability distributions, ρ and π [Cover and Thomas, 2006]. We further use $\text{kl}(p||q)$ as a shorthand for the Kullback-Leibler divergence between two Bernoulli distributions with biases p and q , namely $\text{kl}(p||q) = \text{KL}((1 - p, p)|| (1 - q, q))$. For $\hat{p} \in [0, 1]$ and $\varepsilon \geq 0$ we define the upper and lower inverse of kl, respectively, as $\text{kl}^{-1,+}(\hat{p}, \varepsilon) := \max\{p : p \in [0, 1] \text{ and } \text{kl}(\hat{p}||p) \leq \varepsilon\}$ and $\text{kl}^{-1,-}(\hat{p}, \varepsilon) := \min\{p : p \in [0, 1] \text{ and } \text{kl}(\hat{p}||p) \leq \varepsilon\}$.

The first inequality that we cite is the kl inequality.

Theorem 1 (kl Inequality [Langford, 2005, Foong et al., 2021, 2022]). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded in the $[0, 1]$ interval and with $\mathbb{E}[Z_i] = p$ for all i . Let $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$ be their empirical mean. Then, for any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(\text{kl}(\hat{p}||p) \geq \frac{\ln \frac{1}{\delta}}{n}\right) \leq \delta$$

and, by inversion of the kl,

$$\mathbb{P}\left(p \geq \text{kl}^{-1,+}\left(\hat{p}, \frac{1}{n} \ln \frac{1}{\delta}\right)\right) \leq \delta, \quad (1)$$

$$\mathbb{P}\left(p \leq \text{kl}^{-1,-}\left(\hat{p}, \frac{1}{n} \ln \frac{1}{\delta}\right)\right) \leq \delta. \quad (2)$$

We note that the PAC-Bayes-kl inequality (Theorem 5 below) is based on the inequality $\mathbb{E}\left[e^{n \text{kl}(\hat{p}||p)}\right] \leq 2\sqrt{n}$ [Maurer, 2004], which gives $\mathbb{P}\left(\text{kl}(\hat{p}||p) \geq \frac{\ln \frac{2\sqrt{n}}{\delta}}{n}\right) \leq \delta$. Foong et al. [2021, 2022] reduce the logarithmic factor down to $\ln \frac{1}{\delta}$ by basing the proof on Chernoff's inequality, but this proof technique cannot be combined with PAC-Bayes. Therefore, when we move on to PAC-Bayes we pay the extra $\ln 2\sqrt{n}$ factor in the bounds. It is a long-standing open question whether this factor can be reduced in the PAC-Bayesian setting [Foong et al., 2021].

Next we cite two versions of the Empirical Bernstein inequality.

Theorem 2 (Empirical Bernstein Inequality [Maurer and Pontil, 2009]). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded in a $[a, b]$ interval for some $a, b \in \mathbb{R}$, and with $\mathbb{E}[Z_i] = p$ for all i . Let $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$ be the empirical mean and let $\hat{\sigma} = \frac{1}{n-1} \sum_{i=1}^n (Z_i - \hat{p})^2$ be the empirical variance. Then for any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(p \geq \hat{p} + \sqrt{\frac{2\hat{\sigma} \ln \frac{2}{\delta}}{n}} + \frac{7(b-a) \ln \frac{2}{\delta}}{3(n-1)}\right) \leq \delta. \quad (3)$$

Theorem 3 (Unexpected Bernstein Inequality [Fan et al., 2015, Mhammedi et al., 2019]). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded from above by b for some $b > 0$, and with $\mathbb{E}[Z_i] = p$ for all i . Let $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$ be the empirical mean and let $\hat{\sigma} = \frac{1}{n} \sum_{i=1}^n Z_i^2$ be the empirical mean of the second moments. Let $\psi(u) := u - \ln(1+u)$ for $u > -1$. Then, for any $\gamma \in (0, 1/b)$ and any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(p \geq \hat{p} + \frac{\psi(-\gamma b)}{\gamma b^2} \hat{\sigma} + \frac{\ln \frac{1}{\delta}}{\gamma n}\right) \leq \delta. \quad (4)$$

To facilitate a comparison with other bounds, Theorem 3 provides a slightly different form of the Unexpected Bernstein inequality than the one used by Mhammedi et al. [2019]. We provide a proof of the theorem in Appendix A.1, which is based on the Unexpected Bernstein Lemma [Fan et al., 2015]. We note that an inequality proposed by Cesa-Bianchi et al. [2007] can be used to derive a relaxed version of the Unexpected Bernstein inequality, as discussed in Appendix A.2.

2.2 The Split-kl Inequality

Let Z be a random variable bounded in a $[a, b]$ interval for some $a, b \in \mathbb{R}$ and let $\mu \in [a, b]$ be a constant. We decompose $Z = \mu + Z^+ - Z^-$, where $Z^+ = \max(0, Z - \mu)$ and $Z^- = \max(0, \mu - Z)$. Let $p = \mathbb{E}[Z]$, $p^+ = \mathbb{E}[Z^+]$, and $p^- = \mathbb{E}[Z^-]$. For an i.i.d. sample Z_1, \dots, Z_n let $\hat{p}^+ = \frac{1}{n} \sum_{i=1}^n Z_i^+$ and $\hat{p}^- = \frac{1}{n} \sum_{i=1}^n Z_i^-$.

With these definitions we present the split-kl inequality.

Theorem 4 (Split-kl inequality). *Let Z_1, \dots, Z_n be i.i.d. random variables in a $[a, b]$ interval for some $a, b \in \mathbb{R}$, then for any $\mu \in [a, b]$ and $\delta \in (0, 1)$:*

$$\mathbb{P}\left(p \geq \mu + (b - \mu) \text{kl}^{-1,+}\left(\frac{\hat{p}^+}{b - \mu}, \frac{1}{n} \ln \frac{2}{\delta}\right) - (\mu - a) \text{kl}^{-1,-}\left(\frac{\hat{p}^-}{\mu - a}, \frac{1}{n} \ln \frac{2}{\delta}\right)\right) \leq \delta. \quad (5)$$

Proof.

$$\begin{aligned}
& \mathbb{P}\left(p \geq \mu + (b - \mu) \text{kl}^{-1,+} \left(\frac{\hat{p}^+}{b - \mu}, \frac{1}{n} \ln \frac{2}{\delta} \right) - (\mu - a) \text{kl}^{-1,-} \left(\frac{\hat{p}^-}{\mu - a}, \frac{1}{n} \ln \frac{2}{\delta} \right)\right) \\
& \leq \mathbb{P}\left(p^+ \geq (b - \mu) \text{kl}^{-1,+} \left(\frac{\hat{p}^+}{b - \mu}, \frac{1}{n} \ln \frac{2}{\delta} \right)\right) + \mathbb{P}\left(p^- \leq (\mu - a) \text{kl}^{-1,-} \left(\frac{\hat{p}^-}{\mu - a}, \frac{1}{n} \ln \frac{2}{\delta} \right)\right) \\
& \leq \delta,
\end{aligned}$$

where the last inequality follows by application of the kl upper and lower bounds from Theorem 1 to the first and second terms in the middle line, respectively. \blacksquare

For ternary random variables the best choice is to take μ to be the middle value, then the resulting Z^+ and Z^- are binary and the corresponding kl upper and lower bounds on p^+ and p^- are tight, and the resulting split-kl bound is tight. The inequality can be applied to any bounded random variables, but same way as the kl inequality is not necessarily a good choice for bounded random variables, if the distribution is not binary, the split-kl is not necessarily a good choice if the distribution is not ternary.

2.3 Empirical Comparison

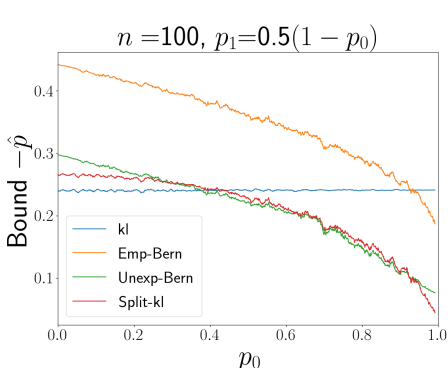
We present an empirical comparison of the tightness of the above four concentration inequalities: the kl, the Empirical Bernstein, the Unexpected Bernstein, and the split-kl. We take n i.i.d. samples Z_1, \dots, Z_n taking values in $\{-1, 0, 1\}$. The choice is motivated both by instructiveness of presentation and by subsequent applications to excess losses. We let $p_{-1} = \mathbb{P}(Z = -1)$, $p_0 = \mathbb{P}(Z = 0)$, and $p_1 = \mathbb{P}(Z = 1)$, where $p_{-1} + p_0 + p_1 = 1$. Then $p = \mathbb{E}[Z] = p_1 - p_{-1}$. We also let $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$.

In Figure 1 we plot the difference between the bounds on p given by the inequalities (1), (3), (4), and (5), and \hat{p} . Lower values in the plot correspond to tighter bounds. To compute the kl bound we first rescale the losses to the $[0, 1]$ interval, and then rescale the bound back to the $[-1, 1]$ interval. For the Empirical Bernstein bound we take $a = -1$ and $b = 1$. For the Unexpected Bernstein bound we take a grid of $\gamma \in \{1/(2b), \dots, 1/(2^k b)\}$ for $k = \lceil \log_2(\sqrt{n/\ln(1/\delta)}/2) \rceil$ and a union bound over the grid, as proposed by Mhammedi et al. [2019]. For the split-kl bound we take μ to be the middle value, 0, of the ternary random variable. In the experiments we take $\delta = 0.05$, and truncate the bounds at 1.

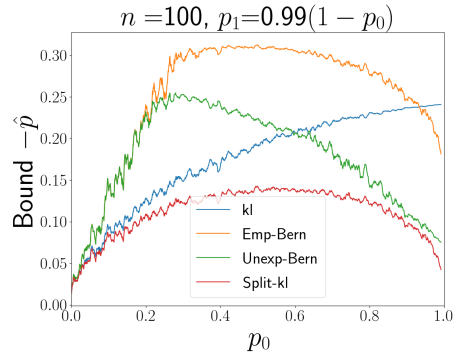
In the first experiment, presented in Figure 1a, we take $p_{-1} = p_1 = (1 - p_0)/2$ and plot the difference between the values of the bounds and \hat{p} as a function of p_0 . For $p_0 = 0$ the random variable Z is Bernoulli and, as expected, the kl inequality performs the best, followed by split-kl, and then Unexpected Bernstein. As p_0 grows closer to 1, the variance of Z decreases and, also as expected, the kl inequality falls behind, whereas split-kl and Unexpected Bernstein go closely together. Empirical Bernstein falls behind all other bounds throughout most of the range, except slightly outperforming kl when p_0 gets very close to 1.

In the second experiment, presented in Figure 1b, we take a skewed random variable with $p_1 = 0.99(1 - p_0)$ and $p_{-1} = 0.01(1 - p_0)$, and again plot the difference between the values of the bounds and \hat{p} as a function of p_0 . This time the kl also starts well for p_0 close to zero, but then falls behind due to its inability of properly handling the values inside the interval. Unexpected Bernstein exhibits the opposite trend due to being based on uncentered second moment, which is high when p_0 is close to zero, even though the variance is small in this case. Empirical Bernstein lags behind all other bounds for most of the range due to poor constants, whereas split-kl matches the tightest bounds, the kl and Unexpected Bernstein, at the endpoints of the range of p_0 , and outperforms all other bounds in the middle of the range, around $p_0 = 0.6$, due to being able to exploit the combinatorics of the problem.

The experiments demonstrate that for ternary random variables the split-kl is a powerful alternative to existing concentration of measure inequalities. To the best of our knowledge, this is also the first empirical evaluation of the Unexpected Bernstein inequality, and it shows that in many cases it is also a powerful inequality. We also observe that in most settings the Empirical Bernstein is weaker than the other three inequalities we consider. Numerical evaluations in additional settings are provided in Appendix D.



(a) Comparison of the concentration bounds with $n = 100$, $\delta = 0.05$ and $p_{-1} = p_1 = 0.5(1 - p_0)$.



(b) Comparison of the concentration bounds with $n = 100$, $\delta = 0.05$, $p_1 = 0.99(1 - p_0)$, and $p_{-1} = 0.01(1 - p_0)$.

Figure 1: Empirical comparison of the concentration bounds.

3 PAC-Bayesian Inequalities

Now we elevate the basic concentration of measure inequalities to the PAC-Bayesian domain. We start with the supervised learning problem setup, then provide a background on existing PAC-Bayesian inequalities, and finish with presentation of the PAC-Bayes-split-kl inequality.

3.1 Supervised Learning Problem Setup and Notations

Let \mathcal{X} be a sample space, \mathcal{Y} be a label space, and let $S = \{(X_i, Y_i)\}_{i=1}^n$ be an i.i.d. sample drawn according to an unknown distribution \mathcal{D} on the product-space $\mathcal{X} \times \mathcal{Y}$. Let \mathcal{H} be a hypothesis space containing hypotheses $h : \mathcal{X} \rightarrow \mathcal{Y}$. The quality of a hypothesis h is measured using the zero-one loss $\ell(h(X), Y) = \mathbf{1}(h(X) \neq Y)$, where $\mathbf{1}(\cdot)$ is the indicator function. The expected loss of h is denoted by $L(h) = \mathbb{E}_{(X,Y) \sim \mathcal{D}}[\ell(h(X), Y)]$, and the empirical loss of h on a sample S is denoted by $\hat{L}(h, S) = \frac{1}{|S|} \sum_{(X,Y) \in S} \ell(h(X), Y)$. We use $\mathbb{E}_{\mathcal{D}}[\cdot]$ as a shorthand for $\mathbb{E}_{(X,Y) \sim \mathcal{D}}[\cdot]$.

PAC-Bayesian bounds bound the generalization error of Gibbs prediction rules. For each input $X \in \mathcal{X}$, Gibbs prediction rule associated with a distribution ρ on \mathcal{H} randomly draws a hypothesis $h \in \mathcal{H}$ according to ρ and predicts $h(X)$. The expected loss of the Gibbs prediction rule is $\mathbb{E}_{h \sim \rho}[L(h)]$ and the empirical loss is $\mathbb{E}_{h \sim \rho}[\hat{L}(h, S)]$. We use $\mathbb{E}_{\rho}[\cdot]$ as a shorthand for $\mathbb{E}_{h \sim \rho}[\cdot]$.

3.2 PAC-Bayesian Analysis Background

Now we present a brief background on the relevant results from the PAC-Bayesian analysis.

PAC-Bayes-kl Inequality The PAC-Bayes-kl inequality cited below is one of the tightest known generalization bounds on the expected loss of the Gibbs prediction rule.

Theorem 5 (PAC-Bayes-kl Inequality, Seeger, 2002, Maurer, 2004). *For any probability distribution π on \mathcal{H} that is independent of S and any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(\exists \rho \in \mathcal{P} : \text{kl}\left(\mathbb{E}_{\rho}[\hat{L}(h, S)] \parallel \mathbb{E}_{\rho}[L(h)]\right) \geq \frac{\text{KL}(\rho \parallel \pi) + \ln(2\sqrt{n}/\delta)}{n}\right) \leq \delta, \quad (6)$$

where \mathcal{P} is the set of all possible probability distributions on \mathcal{H} that can depend on S .

The following relaxation of the PAC-Bayes-kl inequality based on Refined Pinsker's relaxation of the kl divergence helps getting some intuition about the bound [McAllester, 2003]. With probability at least $1 - \delta$, for all $\rho \in \mathcal{P}$ we have

$$\mathbb{E}_{\rho}[L(h)] \leq \mathbb{E}_{\rho}[\hat{L}(h, S)] + \sqrt{2\mathbb{E}_{\rho}[\hat{L}(h, S)] \frac{\text{KL}(\rho \parallel \pi) + \ln(2\sqrt{n}/\delta)}{n}} + \frac{2(\text{KL}(\rho \parallel \pi) + \ln(2\sqrt{n}/\delta))}{n}. \quad (7)$$

If $\mathbb{E}_\rho[\hat{L}(h, S)]$ is close to zero, the middle term in the inequality above vanishes, leading to so-called "fast convergence rates" (convergence of $\mathbb{E}_\rho[\hat{L}(h, S)]$ to $\mathbb{E}_\rho[L(h)]$ at the rate of $1/n$). However, achieving low $\mathbb{E}_\rho[\hat{L}(h, S)]$ is not always possible [Dziugaite and Roy, 2017, Zhou et al., 2019]. Subsequent research in PAC-Bayesian analysis has focused on two goals: (1) achieving fast convergence rates when the variance of prediction errors is low (and not necessarily the errors themselves), and (2) reducing the $\text{KL}(\rho||\pi)$ term, which may be quite large for large hypothesis spaces. For the first goal Tolstikhin and Seldin [2013] developed the PAC-Bayes-Empirical-Bernstein inequality and Mhammedi et al. [2019] proposed to use excess losses and also derived the alternative PAC-Bayes-Unexpected-Bernstein inequality. For the second goal Ambroladze et al. [2007] suggested to use informed priors and Mhammedi et al. [2019] perfected the idea by proposing to average over "forward" and "backward" construction with informed prior. Next we explain the ideas behind the excess losses and informed priors in more details.

Excess Losses Let h^* be a reference prediction rule that is independent of S . We define the excess loss of a prediction rule h with respect to the reference h^* by

$$\Delta_\ell(h(X), h^*(X), Y) = \ell(h(X), Y) - \ell(h^*(X), Y).$$

If ℓ is the zero-one loss, the excess loss naturally gives rise to ternary random variables, but it is well-defined for any real-valued loss function. We use $\Delta_L(h, h^*) = \mathbb{E}_D[\Delta_\ell(h(X), h^*(X), Y)] = L(h) - L(h^*)$ to denote the expected excess loss of h relative to h^* and $\Delta_{\hat{L}}(h, h', S) = \frac{1}{|S|} \sum_{(X, Y) \in S} \Delta_\ell(h(X), h'(X), Y) = \hat{L}(h) - \hat{L}(h')$ to denote the empirical excess loss of h relative to h' . The expected loss of a Gibbs prediction rule can then be written as

$$\mathbb{E}_\rho[L(h)] = \mathbb{E}_\rho[\Delta_L(h, h^*)] + L(h^*).$$

A bound on $\mathbb{E}_\rho[L(h)]$ can thus be decomposed into a summation of a PAC-Bayes bound on $\mathbb{E}_\rho[\Delta_L(h, h^*)]$ and a bound on $L(h^*)$. When the variance of the excess loss is small, we can use tools that exploit small variance, such as the PAC-Bayes-Empirical-Bernstein, PAC-Bayes-Unexpected-Bernstein, or PAC-Bayes-Split-kl inequalities proposed below, to achieve fast convergence rates for the excess loss. Bounding $L(h^*)$ involves just a single prediction rule and does not depend on the value of $\text{KL}(\rho||\pi)$. We note that it is essential that the variance and not just the magnitude of the excess loss is small. For example, if the excess losses primarily take values in $\{-1, 1\}$ and average out to zero, fast convergence rates are impossible.

Informed Priors The idea behind informed priors is to split the data into two subsets, $S = S_1 \cup S_2$, and to use S_1 to learn a prior π_{S_1} , and then use it to learn a posterior on S_2 Ambroladze et al. [2007]. Note that since the size of S_2 is smaller than the size of S , this approach gains in having potentially smaller $\text{KL}(\rho||\pi_{S_1})$, but loses in having a smaller sample size in the denominator of the PAC-Bayes bounds. The balance between the advantage and disadvantage depends on the data: for some data sets it strengthens the bounds, but for some it weakens them. Mhammedi et al. [2019] perfected the approach by proposing to use it in the "forward" and "backward" direction and average over the two. Let S_1 and S_2 be of equal size. The "forward" part uses S_1 to train π_{S_1} and then computes a posterior on S_2 , while the "backward" part uses S_2 to train π_{S_2} and then computes a posterior on S_1 . Finally, the two posteriors are averaged with equal weight and the KL term becomes $\frac{1}{2}(\text{KL}(\rho||\pi_{S_1}) + \text{KL}(\rho||\pi_{S_2}))$. See [Mhammedi et al., 2019] for the derivation.

Excess Losses and Informed Priors Excess losses and informed priors make an ideal combination. If we split S into two equal parts, $S = S_1 \cup S_2$, we can use S_1 to train both a reference prediction rule h_{S_1} and a prior π_{S_1} , and then learn a PAC-Bayes posterior on S_2 , and the other way around. By combining the "forward" and "backward" approaches we can write

$$\mathbb{E}_\rho[L(h)] = \frac{1}{2}\mathbb{E}_\rho[\Delta_L(h, h_{S_1})] + \frac{1}{2}\mathbb{E}_\rho[\Delta_L(h, h_{S_2})] + \frac{1}{2}(L(h_{S_1}) + L(h_{S_2})), \quad (8)$$

and we can use PAC-Bayes to bound the first term using the prior π_{S_1} and the data in S_2 , and to bound the second term using the prior π_{S_2} and the data in S_1 , and we can bound $L(h_{S_1})$ and $L(h_{S_2})$ using the "complementary" data in S_2 and S_1 , respectively.

PAC-Bayes-Empirical-Bernstein Inequalities The excess losses are ternary random variables taking values in $\{-1, 0, 1\}$ and, as we have already discussed, the kl inequality is not well-suited

for them. PAC-Bayesian inequalities tailored for non-binary random variables were derived by Seldin et al. [2012], Tolstikhin and Seldin [2013], Wu et al. [2021], and Mhammedi et al. [2019]. Seldin et al. [2012] derived the PAC-Bayes-Bernstein oracle bound, which assumes knowledge of the variance. Tolstikhin and Seldin [2013] made it into an empirical bound by deriving the PAC-Bayes-Empirical-Bernstein bound for the variance and plugging it into the PAC-Bayes-Bernstein bound of Seldin et al.. Wu et al. [2021] derived an oracle PAC-Bayes-Bennett inequality, which again assumes oracle knowledge of the variance, and showed that it is always at least as tight as the PAC-Bayes-Bernstein, and then also plugged in the PAC-Bayes-Empirical-Bernstein bound on the variance. Mhammedi et al. [2019] derived the PAC-Bayes-Unexpected-Bernstein inequality, which directly uses the empirical second moment. Since we have already shown that the Unexpected Bernstein inequality is tighter than the Empirical Bernstein, and since the approach of Wu et al. requires a combination of two inequalities, PAC-Bayes-Empirical-Bernstein for the variance and PAC-Bayes-Bennett for the loss, whereas the approach of Mhammedi et al. only makes a single application of PAC-Bayes-Unexpected-Bernstein, we only compare our work to the latter.

We cite the inequality of Mhammedi et al. [2019], which applies to an arbitrary loss function. We use $\tilde{\ell}$ and matching tilde-marked quantities to distinguish it from the zero-one loss ℓ . For any $h \in \mathcal{H}$, let $\tilde{L}(h) = \mathbb{E}_D[\tilde{\ell}(h(X), Y)]$ be the expected tilde-loss of h , and let $\hat{\tilde{L}}(h, S) = \frac{1}{|S|} \sum_{(X, Y) \in S} \tilde{\ell}(h(X), Y)$ be the empirical tilde-loss of h on a sample S .

Theorem 6 (PAC-Bayes-Unexpected-Bernstein inequality [Mhammedi et al., 2019]). *Let $\tilde{\ell}(\cdot, \cdot)$ be an arbitrary loss function bounded from above by b for some $b > 0$, and assume that $\hat{\tilde{V}}(h, S) = \frac{1}{|S|} \sum_{(X, Y) \in S} \tilde{\ell}(h(X), Y)^2$ is finite for all h . Let $\psi(u) := u - \ln(1 + u)$ for $u > -1$. Then for any distribution π on \mathcal{H} that is independent of S , any $\gamma \in (0, 1/b)$, and any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(\exists \rho \in \mathcal{P} : \mathbb{E}_\rho[\tilde{L}(h)] \geq \mathbb{E}_\rho[\hat{\tilde{L}}(h, S)] + \frac{\psi(-\gamma b)}{\gamma b^2} \mathbb{E}_\rho[\hat{\tilde{V}}(h, S)] + \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{1}{\delta}}{\gamma n}\right) \leq \delta,$$

where \mathcal{P} is the set of all possible probability distributions on \mathcal{H} that can depend on S .

In optimization of the bound, we take the same grid of $\gamma \in \{1/(2b), \dots, 1/(2^k b)\}$ for $k = \lceil \log_2(\sqrt{n}/\ln(1/\delta)/2) \rceil$ and a union bound over the grid, as we did for Theorem 3.

3.3 PAC-Bayes-Split-kl Inequality

Now we present our PAC-Bayes-Split-kl inequality. For an arbitrary loss function $\tilde{\ell}$ taking values in a $[a, b]$ interval for some $a, b \in \mathbb{R}$, let $\tilde{\ell}^+ := \max\{0, \tilde{\ell} - \mu\}$ and $\tilde{\ell}^- := \max\{0, \mu - \tilde{\ell}\}$ for some $\mu \in [a, b]$. For any $h \in \mathcal{H}$, let $\tilde{L}^+(h) = \mathbb{E}_D[\tilde{\ell}^+(h(X), Y)]$ and $\tilde{L}^-(h) = \mathbb{E}_D[\tilde{\ell}^-(h(X), Y)]$. The corresponding empirical losses are denoted by $\hat{\tilde{L}}^+(h, S) = \frac{1}{n} \sum_{i=1}^n \tilde{\ell}^+(h(X_i), Y_i)$ and $\hat{\tilde{L}}^-(h, S) = \frac{1}{n} \sum_{i=1}^n \tilde{\ell}^-(h(X_i), Y_i)$.

Theorem 7 (PAC-Bayes-Split-kl Inequality). *Let $\tilde{\ell}(\cdot, \cdot)$ be an arbitrary loss function taking values in a $[a, b]$ interval for some $a, b \in \mathbb{R}$. Then for any distribution π on \mathcal{H} that is independent of S , any $\mu \in [a, b]$, and any $\delta \in (0, 1)$:*

$$\mathbb{P}\left[\exists \rho \in \mathcal{P} : \mathbb{E}_\rho[\tilde{L}(h)] \geq \mu + (b - \mu) \text{kl}^{-1,+} \left(\frac{\mathbb{E}_\rho[\hat{\tilde{L}}^+(h, S)]}{b - \mu}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{4\sqrt{n}}{\delta}}{n} \right) - (\mu - a) \text{kl}^{-1,-} \left(\frac{\mathbb{E}_\rho[\hat{\tilde{L}}^-(h, S)]}{\mu - a}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{4\sqrt{n}}{\delta}}{n} \right) \right] \leq \delta,$$

where \mathcal{P} is the set of all possible probability distributions on \mathcal{H} that can depend on S .

Proof. We have $\mathbb{E}_\rho[\tilde{L}(h)] = \mu + \mathbb{E}_\rho[\tilde{L}^+(h)] - \mathbb{E}_\rho[\tilde{L}^-(h)]$. Similar to the proof of Theorem 4, we take a union bound of PAC-Bayes-kl upper bound on $\mathbb{E}_\rho[\tilde{L}^+(h)]$ and PAC-Bayes-kl lower bound on $\mathbb{E}_\rho[\tilde{L}^-(h)]$. ■

3.4 PAC-Bayes-split-kl with Excess Loss and Informed Prior

Looking back at the expected loss decomposition in equation (8), we can use PAC-Bayes-split-kl to bound the first two terms and a bound on the binomial tail distribution to bound the last term. For n i.i.d. Bernoulli random variables Z_1, \dots, Z_n with bias $p \in (0, 1)$, we define the binomial tail distribution $\text{Bin}(n, k, p) = \mathbb{P}(\sum_{i=1}^n X_i \leq k)$ and its inverse $\text{Bin}^{-1}(n, k, \delta) = \max\{p : p \in [0, 1] \text{ and } \text{Bin}(n, k, p) \geq \delta\}$. The following theorem relates $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$ and p .

Theorem 8 (Test Set Bound [Langford, 2005]). *Let Z_1, \dots, Z_n be n i.i.d. Bernoulli random variables with bias $p \in (0, 1)$ and let $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$ be the empirical mean. Then for any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(p \geq \text{Bin}^{-1}(n, n\hat{p}, \delta)\right) \leq \delta.$$

By applying Theorems 7 and 8 to equation (8) we obtain the following result.

Theorem 9. *For any $\mu \in [-1, 1]$ and any $\delta \in (0, 1)$:*

$$\mathbb{P}\left(\exists \rho \in \mathcal{P} : \mathbb{E}_\rho[L(h)] \geq \mu + (1 - \mu)(a) - (\mu + 1)(b) + \frac{1}{2}(c)\right) \leq \delta,$$

where \mathcal{P} is the set of all possible probability distributions on \mathcal{H} that can depend on S ,

$$(a) = \text{kl}^{-1,+} \left(\frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\hat{L}}^+(h, h_{S_1}, S_2)]}{1 - \mu} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\hat{L}}^+(h, h_{S_2}, S_1)]}{1 - \mu}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{8\sqrt{n/2}}{\delta}}{n/2} \right),$$

$$(b) = \text{kl}^{-1,-} \left(\frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\hat{L}}^-(h, h_{S_1}, S_2)]}{\mu + 1} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\hat{L}}^-(h, h_{S_2}, S_1)]}{\mu + 1}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{8\sqrt{n/2}}{\delta}}{n/2} \right),$$

in which $\pi = \frac{1}{2}\pi_{S_1} + \frac{1}{2}\pi_{S_2}$, and

$$(c) = \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_1}, S_2), \frac{\delta}{4} \right) + \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_2}, S_1), \frac{\delta}{4} \right).$$

The proof is postponed to Appendix C.

4 Experiments

We evaluate the performance of the PAC-Bayes-split-kl inequality in linear classification and in weighted majority vote using several data sets from UCI and LibSVM repositories [Dua and Graff, 2019, Chang and Lin, 2011]. An overview of the data sets is provided in Appendix E.1. For linear classification we reproduce the experimental setup of Mhammedi et al. [2019], and for the weighted majority vote we reproduce the experimental setup of Wu et al. [2021].

4.1 The Experimental Setup of Mhammedi et al. [2019]: Linear Classifiers

In the first experiment we follow the experimental setup of Mhammedi et al. [2019], who consider binary classification problems with linear classifiers in \mathbb{R}^d and Gaussian priors and posteriors. A classifier h_w associated with a vector $w \in \mathbb{R}^d$ makes a prediction on an input X by $h_w(X) = \mathbb{1}(w^\top X > 0)$. The posteriors have the form of Gaussian distributions centered at $w_S \in \mathbb{R}^d$, with covariance Σ_S that depends on a sample S , $\rho = \mathcal{N}(w_S, \Sigma_S)$. The informed priors $\pi_{S_1} = \mathcal{N}(w_{S_1}, \Sigma_{S_1})$ and $\pi_{S_2} = \mathcal{N}(w_{S_2}, \Sigma_{S_2})$ are also taken to be Gaussian distributions centered at w_{S_1} and w_{S_2} , with covariance Σ_{S_1} and Σ_{S_2} , respectively. We take the classifier associated with w_{S_1} as the reference classifier h_{S_1} and the classifier associated with w_{S_2} as the reference classifier h_{S_2} . More details on the construction are provided in Appendix E.2.

Figure 2 compares the PAC-Bayes-Unexpected-Bernstein bound PBUB and the PAC-Bayes-split-kl bound PBSkl with excess losses and informed priors. The ternary random variables in this setup take values in $\{-1, 0, 1\}$, and we select μ to be the middle value 0. Since the PAC-Bayes-kl bound (PBkl) is one of the tightest known generalization bounds, we take PBkl with informed priors as a baseline.

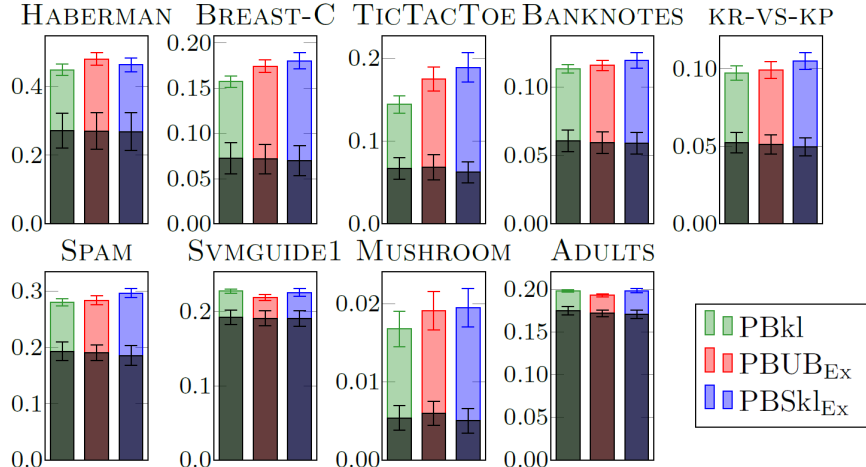


Figure 2: Comparison of the bounds and the test losses of the optimized Gaussian posterior ρ^* generated by PBkl with informed priors, PBUB with excess losses and informed priors, and PBSkl with excess losses and informed priors. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 20 runs of the experiments.

The details on bound calculation and optimization are provided in Appendix E.2. In this experiment all the three bounds, PBkl, PBUB, and PBSkl performed comparably. We believe that the reason is that with informed priors the $\text{KL}(\rho \parallel \pi)$ term is small. From the relaxation of the PBkl bound in equation (7), we observe that a small $\text{KL}(\rho \parallel \pi)$ term implies smaller difference between fast and slow convergence rates, and thus smaller advantage to bounding the excess loss instead of the raw loss. In other words, we believe that the effect of using informed priors dominates the effect of using excess losses. We note that in order to use excess losses we need to train the reference hypothesis h^* on part of the data and, therefore, training an informed prior on the same data comes at no extra cost.

4.2 The Experimental Setup of Wu et al. [2021]: Weighted Majority Vote

In the second experiment we reproduce the experimental setup of Wu et al. [2021], who consider multiclass classification by a weighted majority vote. Given an input $X \in \mathcal{X}$, a hypothesis space \mathcal{H} , and a distribution ρ on \mathcal{H} , a ρ -weighted majority vote classifier predicts $\text{MV}_\rho(X) = \arg \max_{y \in \mathcal{Y}} \mathbb{E}_\rho[\mathbb{1}(h(X) = y)]$. One of the tightest bound for the majority vote is the tandem bound (TND) proposed by Masegosa et al. [2020], which is based on tandem losses for pairs of hypotheses, $\ell(h(X), h'(X), Y) = \mathbb{1}(h(X) \neq Y) \mathbb{1}(h'(X) \neq Y)$, and the second order Markov's inequality. Wu et al. [2021] proposed two improved forms of the bound, both based on a parametric form of the Chebyshev-Cantelli inequality. The first, CCTND, using Chebyshev-Cantelli with the tandem losses and the PAC-Bayes-kl bound for bounding the tandem losses. The second, CCPBB, using tandem losses with an offset, defined by $\ell_\alpha(h(X), h'(X), Y) = (\mathbb{1}(h(X) \neq Y) - \alpha)(\mathbb{1}(h'(X) \neq Y) - \alpha)$ for $\alpha < 0.5$, and PAC-Bayes-Empirical-Bennett inequality for bounding the tandem losses with an offset. We note that while the tandem losses are binary random variables, tandem losses with an offset are ternary random variables taking values in $\{\alpha^2, -\alpha(1 - \alpha), (1 - \alpha)^2\}$ and, therefore, application of Empirical Bernstein type inequalities makes sense. However, in the experiments of Wu et al. CCPBB lagged behind TND and CCTND. We replaced PAC-Bayes-Empirical-Bennett with PAC-Bayes-Unexpected-Bernstein (CCPBUB) and PAC-Bayes-split-kl (CCPBSkl) and showed that the weakness of CCPBB was caused by looseness of PAC-Bayes-Empirical-Bernstein, and that CCPBUB and CCPBSkl lead to tighter bounds that are competitive and sometimes outperforming TND and CCTND. For the PAC-Bayes-split-kl bound we took μ to be the middle value of the tandem loss with an offset, namely, for $\alpha \geq 0$ we took $\mu = \alpha^2$, and for $\alpha < 0$ we took $\mu = -\alpha(1 - \alpha)$.

In Figure 3 we present a comparison of the TND, CCTND, CCPBB, CCPBUB, and CCPBSkl bounds on weighted majority vote of heterogeneous classifiers (Linear Discriminant Analysis, k -Nearest Neighbors, Decision Tree, Logistic Regression, and Gaussian Naive Bayes), which adds the

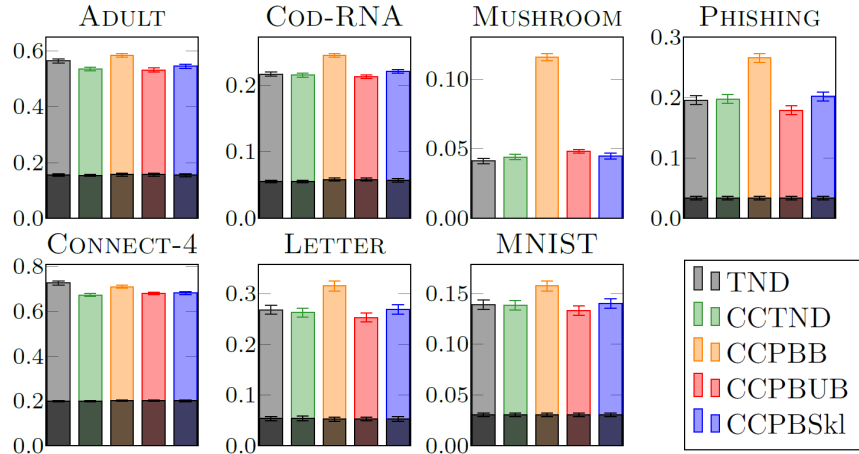


Figure 3: Comparison of the bounds and the test losses of the weighted majority vote on ensembles of heterogeneous classifiers with optimized posterior ρ^* generated by TND, CCTND, CCPBB, CCPBUB, and CCPBSkl. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 10 runs of the experiments.

two new bounds, CCPBUB and CCPBSkl to the experiment done by Wu et al. [2021]. A more detailed description of the experiment and results for additional data sets are provided in Appendix E.3. We note that CCPBUB and CCPBSkl consistently outperform CCPBB, demonstrating that they are more appropriate for tandem losses with an offset. The former two bounds perform comparably to TND and CCTND, which operate on tandem losses without an offset. In Appendix E.4 we replicate another experiment of Wu et al., where we use the bounds to reweigh trees in a random forest classifier. The results are similar to the results for heterogeneous classifiers.

5 Discussion

We have presented the split-kl and PAC-Bayes-split-kl inequalities. The inequalities answer a long-standing open question on how to exploit the structure of ternary random variables in order to provide tight concentration bounds. The proposed split-kl and PAC-Bayes-split-kl inequalities are as tight for ternary random variables, as the kl and PAC-Bayes-kl inequalities are tight for binary random variables.

In our empirical evaluation the split-kl inequality was always competitive with the kl and Unexpected Bernstein inequalities and outperformed both in certain regimes, whereas Empirical Bernstein typically lagged behind. In our experiments in the PAC-Bayesian setting the PAC-Bayes-split-kl was always comparable to PAC-Bayes-Unexpected-Bernstein, whereas PAC-Bayes-Empirical-Bennett most often lagged behind. The first two inequalities were usually comparable to PAC-Bayes-kl, although in some cases the attempt to exploit low variance did not pay off and PAC-Bayes-kl outperformed, which is also the trend observed earlier by Mhammedi et al. [2019]. To the best of our knowledge, this is the first time when the various approaches to exploitation of low variance were directly compared, and the proposed split-kl emerged as a clear winner in the basic setting, whereas in the PAC-Bayes setting in our experiments the PAC-Bayes-Unexpected-Bernstein and PAC-Bayes-split-kl were comparable, and preferable over PAC-Bayes-Empirical-Bernstein and PAC-Bayes-Empirical-Bennett.

Acknowledgments and Disclosure of Funding

This project has received funding from European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 801199. The authors also acknowledge partial support by the Independent Research Fund Denmark, grant number 0135-00259B.

References

- Amiran Ambroladze, Emilio Parrado-Hernández, and John Shawe-Taylor. Tighter PAC-Bayes bounds. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2007.
- Jean Yves Audibert, Rémi Munos, and Csaba Szepesvári. Exploration-exploitation trade-off using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 2009.
- Sergei N. Bernstein. *Probability Theory*. Moscow-Leningrad, 4th edition, 1946. In Russian.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Nicolò Cesa-Bianchi, Yishay Mansour, and Gilles Stoltz. Improved second-order bounds for prediction with expert advice. *Machine Learning*, 66, 2007.
- Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2, 2011.
- Corinna Cortes, Giulia DeSalvo, Claudio Gentile, Mehryar Mohri, and Scott Yang. Online learning with abstention. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2018.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing, 2nd edition, 2006.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2019. URL <http://archive.ics.uci.edu/ml>.
- Gintare Karolina Dziugaite and Daniel M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *UAI*, 2017.
- Xiequan Fan, Ion Grama, and Quansheng Liu. Exponential inequalities for martingales with applications. *Electronic Journal of Probability*, 20:1–22, 2015.
- Ciprian Florescu and Christian Igel. Resilient backpropagation (rprop) for batch-learning in tensorflow. In *International Conference on Learning Representations (ICLR) Workshop*, 2018.
- Andrew Foong, Wessel Bruinsma, David Burt, and Richard Turner. How tight can pac-bayes be in the small data regime? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Andrew Y. K. Foong, Wessel P. Bruinsma, and David R. Burt. A note on the chernoff bound for random variables in the unit interval. *arXiv preprint arXiv:2205.07880*, 2022.
- W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Christian Igel and Michael Hüsken. Empirical evaluation of the improved rprop learning algorithms. *Neurocomputing*, 50, 2003.
- John Langford. Tutorial on practical prediction theory for classification. *Journal of Machine Learning Research*, 6, 2005.
- Andrés R. Masegosa, Stephan S. Lorenzen, Christian Igel, and Yevgeny Seldin. Second order PAC-Bayesian bounds for the weighted majority vote. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- Andreas Maurer. A note on the PAC-Bayesian theorem. arXiv preprint cs/0411099, 2004.
- Andreas Maurer and Massimiliano Pontil. Empirical Bernstein bounds and sample variance penalization. In *Proceedings of the Conference on Learning Theory (COLT)*, 2009.
- David McAllester. PAC-Bayesian stochastic model selection. *Machine Learning*, 51, 2003.
- Zakaria Mhammedi, Peter Grünwald, and Benjamin Guedj. PAC-Bayes un-expected Bernstein inequality. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. Empirical Bernstein stopping. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2008.
- Matthias Seeger. PAC-Bayesian generalization error bounds for Gaussian process classification. *Journal of Machine Learning Research*, 3, 2002.

- Yevgeny Seldin, François Laviolette, Nicolò Cesa-Bianchi, John Shawe-Taylor, and Peter Auer. PAC-Bayesian inequalities for martingales. *IEEE Transactions on Information Theory*, 58, 2012.
- Niklas Thiemann, Christian Igel, Olivier Wintenberger, and Yevgeny Seldin. A strongly quasiconvex PAC-Bayesian bound. In *Proceedings of the International Conference on Algorithmic Learning Theory (ALT)*, 2017.
- Sunil Thulasidasan, Tanmoy Bhattacharya, Jeff A. Bilmes, Gopinath Chennupati, and Jamal Mohd-Yusof. Combating label noise in deep learning using abstention. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2019.
- Ilya Tolstikhin and Yevgeny Seldin. PAC-Bayes-Empirical-Bernstein inequality. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2013.
- Olivier Wintenberger. Optimal learning with Bernstein online aggregation. *Machine Learning*, 106, 2017.
- Yi-Shan Wu, Andres Masegosa, Stephan Lorenzen, Christian Igel, and Yevgeny Seldin. Chebyshev-cantelli pac-bayes-bennett inequality for the weighted majority vote. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- Wenda Zhou, Victor Veitch, Morgane Austern, Ryan P. Adams, and Peter Orbanz. Non-vacuous generalization bounds at the imagenet scale: a PAC-bayesian compression approach. In *ICLR*, 2019.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] All assumptions are stated in the statements of the theorems.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A] The work mainly discussed the fundamental concentration inequalities and the potential applications. The main applications discussed in the work are PAC-Bayes analysis on classification problems, and the potential negative societal impacts follow the problems being studied.
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See Section 4, Appendix E.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 4, Appendix E.
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] Appendix E.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes] In the code and in the paper.
 - (b) Did you mention the license of the assets? [Yes]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] New code in the supplementary.
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [No] We use publicly available datasets from the UCI and LibSVM repositories.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [No] We use publicly available datasets from the UCI and LibSVM repositories. The data does not contain personally identifiable information or offensive content.
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Unexpected Bernstein Inequality

A.1 A Proof of the Unexpected Bernstein Inequality (Theorem 3)

The proof is based on the Unexpected Bernstein lemma.

Lemma 10 (Unexpected Bernstein lemma [Fan et al., 2015, Mhammedi et al., 2019]). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded from above by $b > 0$, and assume that $\sum_{i=1}^n Z_i^2$ is finite. Let $\psi(u) := u - \ln(1 + u)$ for $u \in \mathbb{R}$. Then for any $\gamma \in (0, \frac{1}{b})$:*

$$\mathbb{E} \left[e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i) - \frac{\psi(-\gamma)}{b^2} \sum_{i=1}^n Z_i^2} \right] \leq 1.$$

Proof of Theorem 3. Recall that by the assumption of the theorem Z_1, \dots, Z_n are i.i.d., bounded from above by $b > 0$, and that $p = \mathbb{E}[Z_i]$ for all i , $\hat{p} = \frac{1}{n} \sum_{i=1}^n Z_i$, and $\hat{\sigma} = \frac{1}{n} \sum_{i=1}^n Z_i^2$. For any $\gamma \in (0, 1/b)$ we have:

$$\begin{aligned} \mathbb{P} \left(p - \hat{p} - \frac{\psi(-\gamma b)}{\gamma b^2} \hat{\sigma} \geq \varepsilon \right) &= \mathbb{P} \left(\sum_{i=1}^n \left(\mathbb{E}[Z_i] - Z_i - \frac{\psi(-\gamma b)}{\gamma b^2} Z_i^2 \right) \geq n\varepsilon \right) \\ &= \mathbb{P} \left(e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i - \frac{\psi(-\gamma b)}{\gamma b^2} Z_i^2)} \geq e^{\gamma n\varepsilon} \right) \\ &\leq \mathbb{E} \left[e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i - \frac{\psi(-\gamma b)}{\gamma b^2} Z_i^2)} \right] / e^{\gamma n\varepsilon} \\ &\leq e^{-\gamma n\varepsilon}, \end{aligned}$$

where the first inequality is by application of Markov's inequality and the second inequality is by application of Lemma 10. By taking $\delta = e^{-\gamma n\varepsilon}$ and solving for ε we complete the proof. \blacksquare

A.2 A relaxation of the Unexpected Bernstein lemma

We show that a concentration inequality introduced by Cesa-Bianchi et al. [2007] yields a relaxation of the Unexpected Bernstein Lemma. The inequality of Cesa-Bianchi et al. can be used to directly derive a relaxed version of the Unexpected Bernstein lemma, but as we show the result is weaker than the Unexpected Bernstein lemma. Cesa-Bianchi et al. [2007, Lemma 1] have shown that

$$\forall \gamma \geq -1/2 : \quad \gamma - \gamma^2 \leq \ln(1 + \gamma). \quad (9)$$

Thus,

$$\forall \gamma \leq 1/2 : \quad -\gamma^2 \leq \gamma + \ln(1 - \gamma) = -\psi(-\gamma). \quad (10)$$

This gives a relaxed version of the Unexpected Bernstein lemma. For simplicity, we present it with $b = 1$.

Lemma 11 (Relaxed Unexpected Bernstein lemma). *Let Z_1, \dots, Z_n be i.i.d. random variables bounded from above by 1, and assume that $\sum_{i=1}^n Z_i^2$ is finite. Then for any $\gamma \in [0, \frac{1}{2}]$:*

$$\mathbb{E} \left[e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i) - \gamma^2 \sum_{i=1}^n Z_i^2} \right] \leq 1.$$

Proof. By (10) and Lemma 10 we have

$$\mathbb{E} \left[e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i) - \gamma^2 \sum_{i=1}^n Z_i^2} \right] \leq \mathbb{E} \left[e^{\gamma \sum_{i=1}^n (\mathbb{E}[Z_i] - Z_i) - \psi(-\gamma) \sum_{i=1}^n Z_i^2} \right] \leq 1. \quad \blacksquare$$

We note that it is possible to prove Lemma 11 directly by using inequality (9) and without using Lemma 10, as done by Wintenberger [2017]. The first inequality in our proof of Lemma 11 shows that it is a relaxation of Lemma 10.

B A Proof of the PAC-Bayes Unexpected Bernstein Inequality (Theorem 6)

The proof is based on using the Unexpected Bernstein lemma within a standard change of measure argument cited in Lemma 12 below. We cite the version before the expectations of S' and π are exchanged, which is an intermediate step in the proof of Tolstikhin and Seldin [2013, Lemma 1].

Lemma 12 (PAC-Bayes Lemma [Tolstikhin and Seldin, 2013]). *For any function $f_n : \mathcal{H} \times (\mathcal{X} \times \mathcal{Y})^n \rightarrow \mathbb{R}$ and for any distribution π on \mathcal{H} , with probability at least $1 - \delta$ over a random draw of S , for all distributions ρ on \mathcal{H} simultaneously:*

$$\mathbb{E}_\rho[f_n(h, S)] \leq \text{KL}(\rho \|\pi) + \ln \frac{1}{\delta} + \ln \mathbb{E}_{S'}[\mathbb{E}_\pi[e^{f_n(h, S')}]].$$

Proof of Theorem 6. Let $f_n(h, S) = \gamma n \left(\tilde{L}(h) - \hat{L}(h, S) \right) - \frac{\psi(-b\gamma)}{b^2} n \hat{\mathbb{V}}(h, S)$. Since π is independent of S by assumption, we can exchange the expectations of S and π . Then by Lemma 10 we have $\mathbb{E}[e^{f_n(h, S)}] \leq 1$. By plugging this into Lemma 12 and dividing both sides by γn , we complete the proof. \blacksquare

C Proof of Theorem 9

To prove the theorem, we need the test set bound (Theorem 8), the PAC-Bayes Lemma (Lemma 12), and the following lemma.

Lemma 13 ([Maurer, 2004]). *Let X_1, \dots, X_n be i.i.d. random variables with mean p and bounded in the $[0, 1]$ interval. Let $\hat{p} = \frac{1}{n} \sum_{i=1}^n X_i$ be the empirical mean. Then:*

$$\mathbb{E} \left[e^{n \text{kl}(\hat{p} \| p)} \right] \leq 2\sqrt{n}.$$

Proof of Theorem 9. Recall that

$$\mathbb{E}_\rho[L(h)] = \frac{1}{2} \mathbb{E}_\rho[\Delta_L(h, h_{S_1})] + \frac{1}{2} \mathbb{E}_\rho[\Delta_L(h, h_{S_2})] + \frac{1}{2} (L(h_{S_1}) + L(h_{S_2})). \quad (11)$$

First, by applying Theorem 8 to $L(h_{S_1})$ and $L(h_{S_2})$, respectively, we have:

$$\mathbb{P} \left(L(h_{S_1}) \geq \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_1}, S_2), \delta \right) \right) \leq \delta \quad (12)$$

and

$$\mathbb{P} \left(L(h_{S_2}) \geq \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_2}, S_1), \delta \right) \right) \leq \delta. \quad (13)$$

Next, since

$$\mathbb{E}_\rho[\Delta_L(h, h_{S_1})] = \mu + \mathbb{E}_\rho[\Delta_L^+(h, h_{S_1})] - \mathbb{E}_\rho[\Delta_L^-(h, h_{S_1})]$$

and

$$\mathbb{E}_\rho[\Delta_L(h, h_{S_2})] = \mu + \mathbb{E}_\rho[\Delta_L^+(h, h_{S_2})] - \mathbb{E}_\rho[\Delta_L^-(h, h_{S_2})],$$

for any $\mu \in [a, b]$ we have

$$\begin{aligned} & \frac{1}{2} \mathbb{E}_\rho[\Delta_L(h, h_{S_1})] + \frac{1}{2} \mathbb{E}_\rho[\Delta_L(h, h_{S_2})] \\ &= \mu + \left(\frac{1}{2} \mathbb{E}_\rho[\Delta_L^+(h, h_{S_1})] + \frac{1}{2} \mathbb{E}_\rho[\Delta_L^+(h, h_{S_2})] \right) - \left(\frac{1}{2} \mathbb{E}_\rho[\Delta_L^-(h, h_{S_1})] + \frac{1}{2} \mathbb{E}_\rho[\Delta_L^-(h, h_{S_2})] \right). \end{aligned} \quad (14)$$

Let $\pi = \frac{1}{2} \pi_{S_1} + \frac{1}{2} \pi_{S_2}$, and let S_* be either S_1 or S_2 and $\bar{S}_* = S \setminus S_*$. If h is sampled from π_{S_*} , we take h_{S_*} as a reference hypothesis and estimate the excess loss on \bar{S}_* . Then,

$$\begin{aligned}
\mathbb{E}_S \mathbb{E}_\pi \left[e^{\frac{n}{2} \text{kl} \left(\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right)} \right] &= \frac{1}{2} \sum_{i=1,2} \mathbb{E}_S \mathbb{E}_{\pi_{S_i}} \left[e^{\frac{n}{2} \text{kl} \left(\frac{\Delta_L^+(h, h_{S_i}, \bar{S}_i)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_i})}{1-\mu} \right)} \right] \\
&= \frac{1}{2} \sum_{i=1,2} \mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} \left[e^{\frac{n}{2} \text{kl} \left(\frac{\Delta_L^+(h, h_{S_i}, \bar{S}_i)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_i})}{1-\mu} \right)} \right] \\
&\leq 2\sqrt{n/2},
\end{aligned}$$

where the second equality is due to the fact that π_{S_*} is independent of \bar{S}_* so they are exchangeable, and the inequality follows by Lemma 13.

Therefore, by applying Lemma 12 with $f(h, S) = \frac{n}{2} \text{kl} \left(\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right)$, we have with probability at least $1 - \delta$ over S , for all ρ on \mathcal{H} simultaneously:

$$\mathbb{E}_\rho \left[\frac{n}{2} \text{kl} \left(\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right) \right] \leq \text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n/2}}{\delta}.$$

By the convexity of KL, we further have

$$\text{kl} \left(\mathbb{E}_\rho \left[\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \right] \parallel \mathbb{E}_\rho \left[\frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right] \right) \leq \mathbb{E}_\rho \left[\text{kl} \left(\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \parallel \frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right) \right],$$

which together gives with probability at least $1 - \delta$ over S , for all ρ on \mathcal{H} simultaneously:

$$\text{kl} \left(\mathbb{E}_\rho \left[\frac{\Delta_L^+(h, h_{S_*}, \bar{S}_*)}{1-\mu} \right] \parallel \mathbb{E}_\rho \left[\frac{\Delta_L^+(h, h_{S_*})}{1-\mu} \right] \right) \leq \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n/2}}{\delta}}{n/2}.$$

Similarly, $\Delta_L^-(h, h_{S_*}, \bar{S}_*)$ and $\Delta_L^-(h, h_{S_*})$ also satisfy with probability at least $1 - \delta$ over S , for all ρ on \mathcal{H} simultaneously:

$$\text{kl} \left(\mathbb{E}_\rho \left[\frac{\Delta_L^-(h, h_{S_*}, \bar{S}_*)}{\mu+1} \right] \parallel \mathbb{E}_\rho \left[\frac{\Delta_L^-(h, h_{S_*})}{\mu+1} \right] \right) \leq \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n/2}}{\delta}}{n/2}.$$

Let $\rho = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$ be constructed in a similar way to π , where ρ_1 and ρ_2 are probability distributions on \mathcal{H} . If h is sampled from ρ_* , then we take h_{S_*} as a reference hypothesis and estimate the excess loss on \bar{S}_* . In our case, $\rho_1 = \rho_2 = \rho$. Let Δ° denote either Δ^+ or Δ^- . Then,

$$\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_*})] = \frac{1}{2}\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_1})] + \frac{1}{2}\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_2})]$$

and

$$\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_*}, \bar{S}_*)] = \frac{1}{2}\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_1}, S_2)] + \frac{1}{2}\mathbb{E}_\rho[\Delta_L^\circ(h, h_{S_2}, S_1)].$$

By taking the inverse of kl, we obtain that with probability at least $1 - \delta$ over S , for all ρ on \mathcal{H} simultaneously:

$$\begin{aligned}
&\frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_L^+(h, h_{S_1})]}{1-\mu} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_L^+(h, h_{S_2})]}{1-\mu} \\
&\leq \text{kl}^{-1,+} \left(\frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_L^+(h, h_{S_1}, S_2)]}{1-\mu} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_L^+(h, h_{S_2}, S_1)]}{1-\mu}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n/2}}{\delta}}{n/2} \right), \quad (15)
\end{aligned}$$

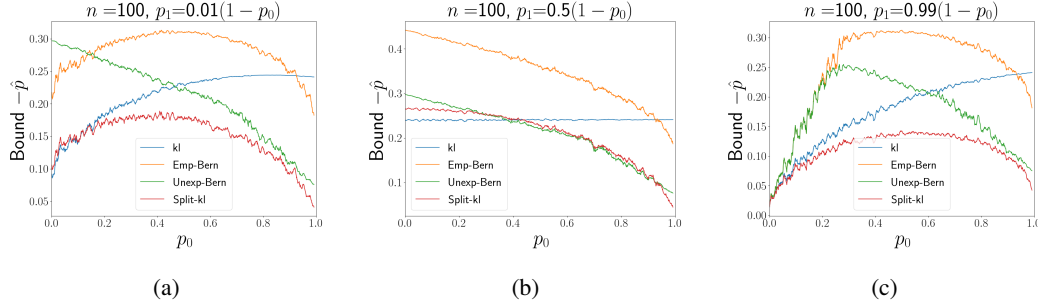


Figure 4: Comparison of the concentration bounds with $n = 100$, $\delta = 0.05$, and (a) $p_1 = 0.01(1 - p_0)$ and $p_{-1} = 0.99(1 - p_0)$, (b) $p_{-1} = p_1 = 0.5(1 - p_0)$, (c) $p_1 = 0.99(1 - p_0)$ and $p_{-1} = 0.01(1 - p_0)$.

and with the same probability

$$\begin{aligned} & \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\bar{L}}^-(h, h_{S_1})]}{\mu + 1} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\bar{L}}^-(h, h_{S_2})]}{\mu + 1} \\ & \geq \text{kl}^{-1, -} \left(\frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\bar{L}}^-(h, h_{S_1}, S_2)]}{\mu + 1} + \frac{1}{2} \frac{\mathbb{E}_\rho[\Delta_{\bar{L}}^-(h, h_{S_2}, S_1)]}{\mu + 1}, \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2\sqrt{n/2}}{\delta}}{n/2} \right). \end{aligned} \quad (16)$$

Thus, we can bound Eq. (14) by Eq.(15) and Eq.(16). By replacing Eq.(11) by the upper bound of each term and taking a union bound, we complete the proof. \blacksquare

D Empirical Comparison

We present more results on empirical comparison of the concentration inequalities: the kl, the Empirical Bernstein, the Unexpected Bernstein, and the split-kl. In particular, Section D.1 expands the empirical comparison in Section 2.3 in the body for ternary random variables, and Section D.2 studies the empirical comparison of bounded random variables. The source code for replicating the experiments is available at Github².

D.1 Ternary Random Variables

In this section, we follow the settings and the parameters in Section 2.3, considering n i.i.d. samples taking values in $\{-1, 0, 1\}$. For completeness, Figure 4b and Figure 4c repeats Figures 1a and 1b while we add Figure 4a, where the probability is defined by $p_1 = 0.01(1 - p_0)$ and $p_{-1} = 0.99(1 - p_0)$. In this case, the kl starts well for p_0 close to zero, but similar to the case in Figure 4c falls behind due to its inability of properly handling the values inside the interval. The Unexpected Bernstein and the Empirical Bernstein perform similarly when p_0 is small in Figure 4c since the bounds are cut to 1, while Unexpected Bernstein falls behind Empirical Bernstein when p_0 is small in Figure 4a due to the uncentered second moment. The split-kl matches, and in many cases outperforms, the tightest bounds.

Figure 5 has the same setting with a larger number of samples $n = 1000$. The trends of the bounds are similar to Figure 4. However, the Empirical Bernstein performs better than the Unexpected Bernstein in Figure 5a and Figure 5c when p_0 is less than 0.6. In both cases, split-kl keeps its leading position. When $p_1 = p_{-1} = (1 - p_0)/2$ (Figure 5b), as $p_0 = 0$, the random variable becomes Bernoulli and, as expected, the kl bound performs the best, followed by the split-kl, and then the two Bernstein bounds. As p_0 grows larger, the kl bound falls behind the other three bounds due to inability of properly handling the values inside the interval. The Unexpected Bernstein, the Empirical Bernstein and the split-kl perform similarly well.

²<https://github.com/YiShanAngWu/Split-KL-R/tree/main/simulation>

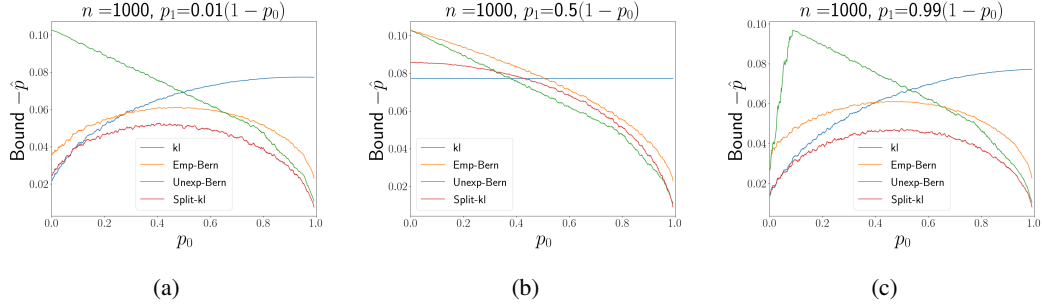


Figure 5: Comparison of the concentration bounds with $n = 1000$, $\delta = 0.05$, and (a) $p_1 = 0.01(1 - p_0)$ and $p_{-1} = 0.99(1 - p_0)$, (b) $p_{-1} = p_1 = 0.5(1 - p_0)$, (c) $p_1 = 0.99(1 - p_0)$ and $p_{-1} = 0.01(1 - p_0)$.

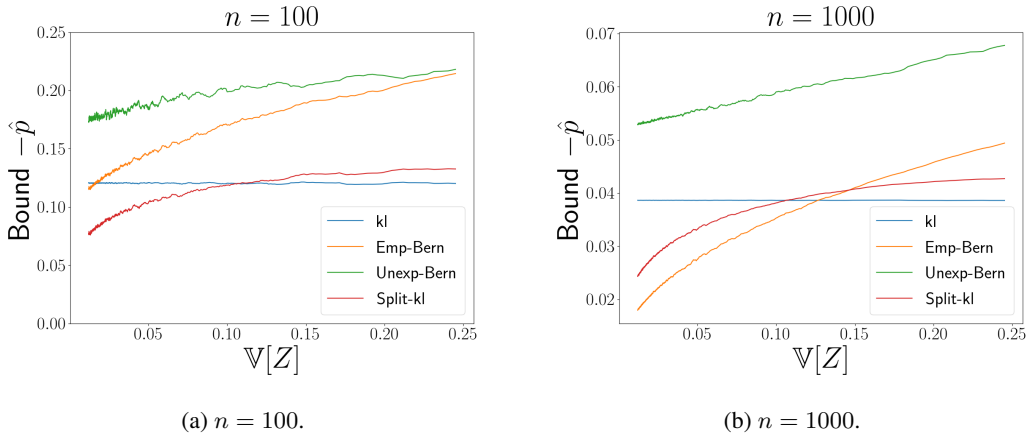


Figure 6: Comparison of the concentration bounds for beta distributions with parameters $\alpha = \beta$ taking values in the $[0.01, 10]$ interval, with $\delta = 0.05$, and with the number of samples $n = 100$ and $n = 1000$, respectively.

D.2 Bounded Random Variables

In this section, we study a more general setting, where the i.i.d. random variables Z_1, \dots, Z_n taking values in $[0, 1]$. Naturally, we consider the random variables following beta distribution with parameters $\alpha > 0$ and $\beta > 0$, where the mean $p = \frac{\alpha}{\alpha + \beta}$ and the variance $\mathbb{V}[Z] = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)}$. For the Empirical Bernstein bound, we take $a = 0$ and $b = 1$. For the Unexpected Bernstein bound we take a grid of $\gamma \in \{1/(2b), \dots, 1/(2^k b)\}$ for $b = 1$, $k = \lceil \log_2(\sqrt{n/\ln(1/\delta)}/2) \rceil$, and a union bound over the grid, as in Section 2.3. For the split-kl bound we take μ to be the middle value 0.5. Again, in the experiments we take $\delta = 0.05$ and cut the bounds to 1.

In Figure 6 we take $\alpha = \beta$ in an interval of $[0.01, 10]$. The mean is a constant $p = 0.5$ throughout the interval and the variance is in an interval of $[0, 0.12, 0.245]$, where a small α and β corresponds to a large variance, and a large α and β corresponds to a small variance. We plot the difference between the values of the bounds and \hat{p} as a function of the variance $\mathbb{V}[Z]$. Since the true mean is a constant, the kl bound is also almost a constant throughout the interval. When the variance large, the kl bound performs the best, followed by split-kl and the two Bernstein bounds. When the variance is small, the Empirical Bernstein bound exploit the low variance and outperform all the others when the number of samples is sufficiently large. The Unexpected Bernstein falls behind due the uncentered second moment. The split-kl bound is comparable to the kl bound when the variance is large and also comparable to the tightest bound when the variance is small.

In Figure 7 we consider another case where the variances stay similar but the means lie across the spectrum in between 0 and 1. We define the distributions being studied by a combination of two sets

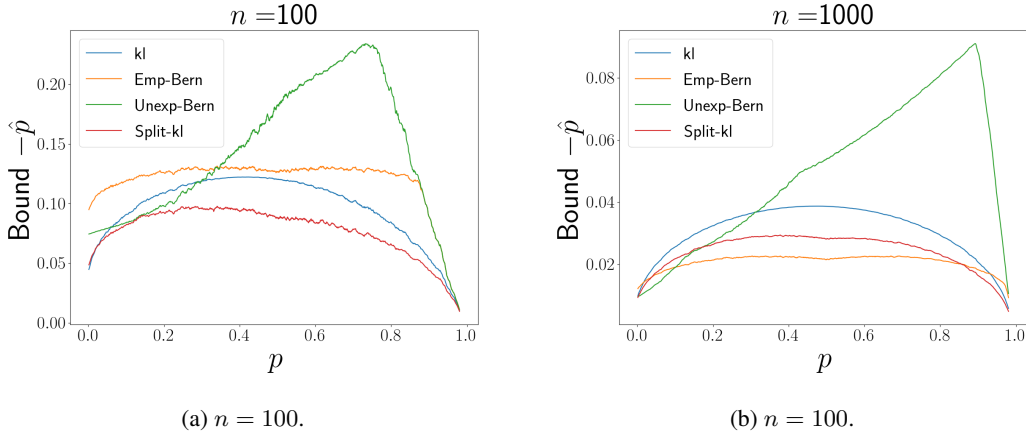


Figure 7: Empirical comparison of concentration bounds for beta distribution with parameters α and β and with the number of samples $n = 100$ and $n = 1000$. For $p \in [0, 0.5]$, we take $\beta = 5$ and $\alpha \in [0.01, 5]$ while for $p \in [0.5, 1]$, we take $\alpha = 5$ and $\beta \in [0.01, 5]$.

of probability distributions. First of all, we take $\beta = 5$ and $\alpha \in [0.01, 5]$, resulting in the mean p in between 0 and 0.5. We define another part by taking $\alpha = 5$ and $\beta \in [0.01, 5]$, resulting in the mean p in between 0.5 and 1. We plot the difference between the values of the bounds and \hat{p} as a function of p . The kl bound is relatively weak around $p = 0.5$ as expected. Since the variances stay similar across the interval, the performance of the Empirical Bernstein stay similar throughout the spectrum, and is tighter than the kl bound when the number of samples is sufficiently large. The split-kl bound is comparable and sometimes outperform the tightest bounds. The Unexpected Bernstein bound again falls behind due to the uncentered second moments.

E Experiments

E.1 Data Sets

As mentioned in Section 4, we consider data sets from UCI and LibSVM repositories [Dua and Graff, 2019, Chang and Lin, 2011], as well as Fashion-MNIST from Zalando Research³. An overview of the data sets is listed in Table 1, where Banknote stands for Banknote Authentication, Breast-C stands for Breast Cancer Wisconsin, Fashion stands for Fashion-MNIST, Haberman stands for Haberman’s Survival, and kr-vs-kp stands for Chess (King-Rook vs. King-Pawn). For data sets with a training and a testing set, we combine the training and the testing sets.

Linear Classifiers. For the linear classifiers experiments, we consider selected data sets with binary class ($c = 2$). We rescale all the real-valued attributes to the $[-1, 1]$ interval and use one-hot encoding to encode categorical variables to $\{-1, 1\}$, which increases the dimension of the attributes for some of the data sets. In particular, the effective dimension of Adult becomes 108, kr-vs-kp becomes 73, and Mushroom becomes 116. We remove rows containing missing features. For each data set, we shuffle the data sets and take four 5-fold train-test split, which gives 20 runs in total.

Weighted Majority Vote. For the weighted majority vote experiments, including the ensemble of multiple heterogeneous classifiers and the random forest, we consider several binary and multiclass ($c > 2$) data sets. We encode the categorical variables into integers and remove rows containing missing features. For each data set we take 10 runs, and for each run we randomly set aside 20% of sample as the test set.

³<https://github.com/zalandoresearch/fashion-mnist>

Table 1: Data set overview. c_{\min} and c_{\max} denote the minimum and maximum class frequency.

Data set	N	d	c	c_{\min}	c_{\max}	Source
Adult	32561	14	2	0.2408	0.7592	LIBSVM (a1a)
Banknote	1372	4	2	0.4447	0.5553	UCI
Breast-C	699	9	2	0.3448	0.6552	UCI
Cod-RNA	59535	8	2	0.3333	0.6667	LIBSVM
Connect-4	67557	126	3	0.0955	0.6583	LIBSVM
Fashion	70000	784	10	0.1000	0.1000	Zalando Research
Haberman	306	3	2	0.2647	0.7353	UCI
kr-vs-kp	3196	36	2	0.48	0.52	UCI
Letter	20000	16	26	0.0367	0.0406	UCI
MNIST	70000	780	10	0.0902	0.1125	LIBSVM
Mushroom	8124	22	2	0.4820	0.5180	LIBSVM
Pendigits	10992	16	10	0.0960	0.1041	LIBSVM
Phishing	11055	68	2	0.4431	0.5569	LIBSVM
Protein	24387	357	3	0.2153	0.4638	LIBSVM
SVMGuide1	3089	4	2	0.3525	0.6475	LIBSVM
SatImage	6435	36	6	0.0973	0.2382	LIBSVM
Sensorless	58509	48	11	0.0909	0.0909	LIBSVM
Shuttle	58000	9	7	0.0002	0.7860	LIBSVM
Spambase	4601	57	2	0.394	0.606	UCI
Splice	3175	60	2	0.4809	0.5191	LIBSVM
TicTacToe	958	9	2	0.347	0.653	UCI
USPS	9298	256	10	0.0761	0.1670	LIBSVM
w1a	49749	300	2	0.0297	0.9703	LIBSVM

E.2 Linear Classifiers

In this section, we describe the details of the experimental setting of the linear classifiers E.2.1, the details of the bounds E.2.2, and the details of optimization E.2.3. The source code for replicating the experiments is available at Github⁴.

E.2.1 Experimental Setting

In this section, we detail the settings and the construction of informed priors and excess losses using linear classifiers with Gaussian posterior. We follow the construction by Mhammedi et al. [2019].

As described in Section 4, the posterior $\rho = \mathcal{N}(w_S, \Sigma_S)$ is a Gaussian distribution centered at w_S , which is learned on S using regularized logistic regression

$$w_S = \arg \min_{w \in \mathbb{R}^d} \frac{\lambda \|w\|^2}{2} + \frac{1}{|S|} \sum_{(X,Y) \in S} -(Y \ln \phi(w^\top X) + (1 - Y) \ln(1 - \phi(w^\top X))), \quad (17)$$

where $\phi(x) := 1/(1 + e^{-x})$ for $x \in \mathbb{R}$ is the sigmoid function. The covariance of the posterior is a diagonal matrix $\sigma^2 I_d$, where the variance σ^2 is learned from the corresponding PAC-Bayes bounds. We use the informed priors in all the PAC-Bayes bounds. The informed priors $\pi_{S_1} = \mathcal{N}(w_{S_1}, \Sigma_{S_1})$ and $\pi_{S_2} = \mathcal{N}(w_{S_2}, \Sigma_{S_2})$ are also chosen to be Gaussian distributions over \mathbb{R}^d , where the centers of the distributions are learned similarly using regularized logistic regression on the corresponding sample S_1 and S_2 . If using excess losses, we take the classifier associated with w_{S_1} as the reference classifier h_{S_1} for the “forward” approach and take the classifier associated with w_{S_2} as the reference classifier h_{S_2} for the “backward” approach. We let the covariance of the informed priors to be also diagonal matrices $\Sigma_{S_1} = \Sigma_{S_2} = \sigma_\pi^2 I_d$, where σ_π^2 is selected from a grid $\mathcal{G} = \{1/2, \dots, 1/2^j\}$ for $j = \lceil \log_2 |S| \rceil$.

For all data sets, we use $\lambda = 0.01$ in equation (17) and solve it using the BFGS algorithm. For all the bounds, we take $\delta = 0.05$. Note that to be able to select the variance of the priors from a grid \mathcal{G} , we

⁴<https://github.com/YiShanAngWu/Split-KL-R>

have to take a union bound over \mathcal{G} . Since the hypothesis space is infinitely large, we approximate the excess risk by drawing 100 classifiers from the posterior ρ and compute the excess losses with respect to the reference classifiers.

E.2.2 Bounds

As mentioned in the body that we used informed priors for all the bounds we applied. The PBSkl_{Ex} bound is presented in Theorem 9, while the PBUB_{Ex} bound and the PBkl bound will be presented in the following. The idea to derive PAC-Bayes bounds with informed priors in general is similar to the technique used in the proof of Theorem 9 in Appendix C.

The key element of the derivations is to bound $\mathbb{E}_{S'}[\mathbb{E}_\pi[e^{f_n(h,S')}]]$ for a given function $f_n : \mathcal{H} \times (\mathcal{X} \times \mathcal{Y})^n \rightarrow \mathbb{R}$ in Lemma 12. Let the prior $\pi = \frac{1}{2}\pi_{S_1} + \frac{1}{2}\pi_{S_2}$, and let S_* be either S_1 or S_2 . If h is sampled from π_{S_*} , we estimate the loss on $\bar{S}_* = S \setminus S_*$. Then,

$$\mathbb{E}_S \mathbb{E}_\pi \left[e^{f_n(h,S)} \right] = \frac{1}{2} \sum_{i=1,2} \mathbb{E}_S \mathbb{E}_{\pi_{S_i}} \left[e^{f_n(h,S)} \right] = \frac{1}{2} \sum_{i=1,2} \mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} \left[e^{f_n(h,S)} \right],$$

where the second equality is due to the fact that π_{S_*} is independent of \bar{S}_* so they are exchangeable. We will then select the function $f_n(h, S)$ later such that $\mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} [e^{f_n(h,S)}]$ is bounded for $i = 1, 2$.

Similarly, we let $\rho = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$. If h is sampled from ρ_* , we estimate the loss on $\bar{S}_* = S \setminus S_*$. Then we have

$$\mathbb{E}_\rho[\tilde{L}(h)] = \frac{1}{2}\mathbb{E}_{\rho_1}[\tilde{L}(h)] + \frac{1}{2}\mathbb{E}_{\rho_2}[\tilde{L}(h)] \quad (18)$$

and

$$\mathbb{E}_\rho[\hat{L}(h, S_*)] = \frac{1}{2}\mathbb{E}_{\rho_1}[\hat{L}(h, S_2)] + \frac{1}{2}\mathbb{E}_{\rho_2}[\hat{L}(h, S_1)] \quad (19)$$

for any loss $\tilde{\ell}$ and the corresponding quantities following the definitions in Section 3.3. We assume that $\rho_1 = \rho_2 = \rho$ in all the bounds. Note that for simpler computation, we replace $\text{kl}(\rho \parallel \pi)$ by its upper bound $\frac{1}{2}\text{kl}(\rho \parallel \pi_{S_1}) + \frac{1}{2}\text{kl}(\rho \parallel \pi_{S_2})$ for all the bounds in the experiments.

PAC-Bayes-kl bound with Informed Priors (PBkl). We take the PAC-Bayes-kl bound with informed priors as the baseline:

$$\mathbb{E}_\rho[L(h)] \leq \text{kl}^{-1,+} \left(\frac{1}{2}\mathbb{E}_\rho[\hat{L}(h, S_1)] + \frac{1}{2}\mathbb{E}_\rho[\hat{L}(h, S_2)], \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{2|\mathcal{G}|\sqrt{n/2}}{\delta}}{n/2} \right),$$

which is obtained by letting $f_n(h, S) = \frac{n}{2} \text{kl}(\hat{L}(h, \bar{S}_*) \parallel L(h))$ and plugging it into Lemma 12. In particular, we have $\mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} [e^{f_n(h,S)}] = \mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} [e^{\frac{n}{2} \text{kl}(\hat{L}(h, \bar{S}_i) \parallel L(h))}] \leq 2\sqrt{n/2}$ for $i = 1, 2$ by Lemma 13. Also, by the convexity of KL, we further have

$$\text{kl} \left(\mathbb{E}_\rho[\hat{L}(h, S_*)] \parallel \mathbb{E}_\rho[L(h)] \right) \leq \mathbb{E}_\rho \left[\text{kl}(\hat{L}(h, S_*) \parallel L(h)) \right].$$

By taking the inverse of kl, applying the relations in Eq. (18) and Eq. (19), and taking a union bound over \mathcal{G} , we obtain the desired formula.

PAC-Bayes-Unexpected Bernstein Bound with Excess Loss and Informed Priors (PBUB_{Ex}). Let $\Delta_{\hat{\Psi}}(h, h^*, S) = \frac{1}{|\mathcal{S}|} \sum_{(X,Y) \in \mathcal{S}} (\Delta_\ell(h(X), h^*(X), Y))^2$ denote the average of the second moment of the excess losses. Then, the PBUB_{Ex} has the form:

$$\begin{aligned} \mathbb{E}_\rho[L(h)] &\leq \frac{1}{2}\mathbb{E}_\rho[\Delta_{\hat{L}}(h, h_{S_1}, S_2)] + \frac{1}{2}\mathbb{E}_\rho[\Delta_{\hat{L}}(h, h_{S_2}, S_1)] \\ &+ \frac{\psi(-\gamma b)}{\gamma b^2} \left(\frac{1}{2}\mathbb{E}_\rho[\Delta_{\hat{\Psi}}(h, h_{S_1}, S_2)] + \frac{1}{2}\mathbb{E}_\rho[\Delta_{\hat{\Psi}}(h, h_{S_2}, S_1)] \right) + \frac{\text{KL}(\rho \parallel \pi) + \ln \frac{3|\mathcal{G}|\Gamma}{\delta}}{\gamma(n/2)} \\ &+ \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_1}, S_2), \frac{\delta}{3|\mathcal{G}|} \right) + \text{Bin}^{-1} \left(\frac{n}{2}, \frac{n}{2} \hat{L}(h_{S_2}, S_1), \frac{\delta}{3|\mathcal{G}|} \right), \end{aligned}$$

where $|\Gamma|$ comes from a union bound over a grid of $\gamma \in \Gamma = \{1/(2b), \dots, 1/(2^k b)\}$ for $k = \lceil \log_2(\sqrt{|S|/\ln(1/\delta)}/2) \rceil$ when applying the PAC-Bayes-Unexpected-Bernstein inequality.

The last line of the bound is by applying Theorem 8 to $L(h_{S_1})$ and $L(h_{S_2})$ as in Theorem 9, while the first two lines of the bound are derived from applying the PAC-Bayes-Unexpected-Bernstein inequality to the first two terms in equation (8). In particular, let $f_n(h, S) = \gamma \frac{n}{2} (\Delta_L(h, h_{S_*}) - \Delta_{\hat{L}}(h, h_{S_*}, \bar{S}_*)) - \frac{\psi(-b\gamma)}{b^2} \frac{n}{2} \Delta_{\hat{\psi}}(h, h_{S_*}, \bar{S}_*)$ and plug it into Lemma 12. Then we have

$$\mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} [e^{f_n(h, S)}] = \mathbb{E}_{S_i} \mathbb{E}_{\pi_{S_i}} \mathbb{E}_{\bar{S}_i} [e^{\gamma \frac{n}{2} (\Delta_L(h, h_{S_i}) - \Delta_{\hat{L}}(h, h_{S_i}, \bar{S}_i)) - \frac{\psi(-b\gamma)}{b^2} \frac{n}{2} \Delta_{\hat{\psi}}(h, h_{S_i}, \bar{S}_i)}] \leq 1$$

for $i = 1, 2$ by Lemma 10. By moving the empirical quantities to the right hand side, applying the relations in Eq. (18) and Eq. (19), and taking the union bounds, we obtain the desired formula.

PAC-Bayes-spli-kl Bound with Excess Loss and Informed Priors (PBSkl_{Ex}). The bound is stated in Theorem 9, except that we replace δ by $\delta/|\mathcal{G}|$ for the union bound of \mathcal{G} . We take $\mu = 0$ for the bound in the experiments.

E.2.3 Optimization

Since the center of the posterior w_S is learned using regularized logistic regression, the only thing remains is to decide the variance of the posterior σ^2 using the PAC-Bayes bounds. In general, the variance can be any non-negative values since the bound holds with high probability for all ρ simultaneously. For simpler computation, we only consider the variance taking the same value as the variance of the priors *i.e.*, taking $\sigma^2 = \sigma_\pi^2 \in \mathcal{G}$. For each PAC-Bayes bounds, we find the optimal σ^2 by iterating over variances $\sigma^2 = \sigma_\pi^2 \in \mathcal{G}$ and return the one corresponds to the tightest bound. We approximate $\mathbb{E}_\rho[\cdot]$ by sampling 100 classifiers from ρ .

The inverse kl in the PBkl and the PBSkl_{Ex} bounds can be computed by binary search. The inverse of the binomial tail distribution in the PBUB_{Ex} and the PBSkl_{Ex} bounds can also be computed by binary search. To optimize the PBUB_{Ex} bound, we also need to iterate over $\gamma \in \Gamma$.

E.3 Ensemble of Multiple Heterogeneous Classifiers

In this section, we describe the details of the experimental setting of the ensemble of multiple heterogeneous classifiers E.3.1, the details of bounds and optimization E.3.2, and lastly, the results E.3.3. The source code for replicating the experiments is available at Github⁵.

E.3.1 Experimental Setting

In this experiment, we follow the setting in Wu et al. [2021]. We take the following standard classifiers available in *scikit-learn* using default parameters to build the ensemble: 1. **Linear Discriminant Analysis** 2. **Decision Tree** 3. **Logistic Regression** 4. **Gaussian Naive Bayes**. We also take three versions of **k-Nearest Neighbors**: 1. $k = 3$ with uniform weights (*i.e.*, all points in each neighborhood are weighted equally) 2. $k = 5$ with uniform weights, and 3. $k = 5$ with the weights of the points are defined by the inverse of their L2 distance. Thus, there are 7 classifiers for ensemble in total.

Ensemble Construction by Bagging. We follow the construction used by Masegosa et al. [2020], Wu et al. [2021]. For each classifier h , we generate a random split of the data set S into a pair of subsets $S = S_h \cup \bar{S}_h$, where $\bar{S}_h = S \setminus S_h$. We generate the split by the standard bagging method, where S_h contains $0.8|S|$ samples randomly subsampled with replacement from S . We train the classifier on S_h , and estimate the expected loss on the out-of-bag (OOB) sample \bar{S}_h to make an unbiased estimation. The resulting set of classifiers produces an ensemble, while the estimates are used for calculating the bounds and deciding the weights of the ensemble. In particular, we estimate the expected loss by $\hat{L}(h, \bar{S}_h)$, and let $n = \min_h |\bar{S}_h|$. In the remaining of the paper, we call the tandem loss with an offset α by the α -tandem loss. Then for a pair of classifiers h and h' , we take the overlap of the OOB sample $\bar{S}_h \cap \bar{S}_{h'}$ to estimate the unbiased tandem loss $\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})$,

⁵<https://github.com/StephanLorenzen/MajorityVoteBounds>

α -tandem loss $\hat{L}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$, the second moment of the α -tandem loss $\hat{V}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$, the variance of the α -tandem loss $\hat{\text{Var}}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$, as well as the splits of the α -tandem loss $\hat{L}_\alpha^+(h, h', \bar{S}_h \cap \bar{S}_{h'})$ and $\hat{L}_\alpha^-(h, h', \bar{S}_h \cap \bar{S}_{h'})$. Let $m = \min_{h, h'} |\bar{S}_h \cap \bar{S}_{h'}|$ be the minimum size of the overlap.

E.3.2 Bounds and Optimization

The bounds we are comparing in this section are the TND, CCTND, CCPBB, CCPBUB, and CCPBSkl bounds. The derivations of the first three bounds are provided in Masegosa et al. [2020], Wu et al. [2021], while we will provide the derivations of the CCPBUB bound and the CCPBSkl bound.

In the experiments, we take $\delta = 0.05$ and take π to be a uniform distribution over the classifiers. For CCPBB, CCPBUB and CCPBSkl bounds, we take a grid of $\alpha \in [-0.5, 0.5]$ since the bounds are not differentiable w.r.t α . Note that we don't need a union bound over α [Wu et al., 2021]. To optimize the weighting ρ , we applied iRProp+ for the gradient based optimization [Igel and Hüsken, 2003, Florescu and Igel, 2018], until the bound did not improve more than 10 for 10 iterations. To find the optimal ρ and the parameters, we start by $\rho = \pi$, and apply alternating minimization until the bound doesn't change for more than 10^{-9} . The details of alternating minimization for each bound are provided below.

We first cite the three existing bounds.

Tandem Bound (TND) [Masegosa et al., 2020] They used the following formula to compute the bound after obtaining the optimal weights ρ :

$$L(\text{MV}_\rho) \leq 4 \text{kl}^{-1,+} \left(\mathbb{E}_{\rho^2}[\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})], \frac{2 \text{KL}(\rho \parallel \pi) + \ln(4\sqrt{m}/\delta)}{m} \right),$$

and used the following relaxation, based on the PAC-Bayes- λ inequality F, for easier optimization:

$$L(\text{MV}_\rho) \leq 4 \left(\frac{\mathbb{E}_{\rho^2}[\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{1 - \lambda/2} + \frac{2 \text{KL}(\rho \parallel \pi) + \ln(2\sqrt{m}/\delta)}{\lambda(1 - \lambda/2)m} \right) \quad (20)$$

for any $\lambda \in (0, 2)$. The bound can be optimized by implementing alternating minimization: Given ρ , starting with $\rho = \pi$, find the corresponding optimal λ (Sec. F). Then given λ , optimize ρ by projected gradient descent.

Chebyshev-Cantelli bound with TND empirical loss estimate bound (CCTND) [Wu et al., 2021] They used the following formula to compute the bound after obtaining the optimal weights ρ :

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\text{kl}^{-1,+} \left(\mathbb{E}_{\rho^2}[\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})], \frac{2 \text{KL}(\rho \parallel \pi) + \ln(4\sqrt{m}/\delta)}{m} \right) - 2\alpha \text{kl}^{-1,\circ} \left(\mathbb{E}_\rho[\hat{L}(h, \bar{S}_h)], \frac{\text{KL}(\rho \parallel \pi) + \ln(4\sqrt{n}/\delta)}{n} \right) + \alpha^2 \right],$$

for $\alpha < 0.5$, where \circ is “-” for $\alpha \geq 0$ and “+” otherwise. On the other hand, they used the following relaxations, based on the PAC-Bayes- λ inequality F, for easier optimization:

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\frac{\mathbb{E}_{\rho^2}[\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{1 - \frac{\lambda}{2}} + \frac{2 \text{KL}(\rho \parallel \pi) + \ln(4\sqrt{m}/\delta)}{\lambda(1 - \frac{\lambda}{2})m} - 2\alpha \left(\left(1 - \frac{\gamma}{2}\right) \mathbb{E}_\rho[\hat{L}(h, \bar{S}_h)] - \frac{\text{KL}(\rho \parallel \pi) + \ln(4\sqrt{n}/\delta)}{\gamma n} \right) + \alpha^2 \right]$$

for $0 \leq \alpha < 0.5$, and

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\frac{\mathbb{E}_{\rho^2}[\hat{L}(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{1 - \frac{\lambda}{2}} + \frac{2 \text{KL}(\rho \parallel \pi) + \ln(4\sqrt{m}/\delta)}{\lambda(1 - \frac{\lambda}{2})m} - 2\alpha \left(\frac{\mathbb{E}_\rho[\hat{L}(h, \bar{S}_h)]}{1 - \frac{\gamma}{2}} + \frac{\text{KL}(\rho \parallel \pi) + \ln(4\sqrt{n}/\delta)}{\gamma(1 - \frac{\gamma}{2})n} \right) + \alpha^2 \right]$$

for $\alpha < 0$. The optimization of the bound can, again, be done by alternating minimization of the following steps: 1. Given α and ρ , where we start with $\alpha = 0$ and $\rho = \pi$, compute the corresponding closed-form minimizer λ and γ (Sec. F). 2. Given ρ , λ and γ , find the closed-form minimizer α . 3. Given parameters α , λ , and γ , optimize over ρ using projected gradient descent.

Chebyshev-Cantelli bound with PAC-Bayes-Bennett loss estimate bound (CCPBB) [Wu et al., 2021] The CCPBB bound has the following formula for both computing the bound and optimization:

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\mathbb{E}_{\rho^2}[\hat{L}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})] + \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{2k_\lambda k_\gamma}{\delta}}{\gamma m} \right. \\ \left. + \frac{\phi(\gamma K_\alpha)}{\gamma K_\alpha^2} \left(\frac{\mathbb{E}_{\rho^2}[\hat{\text{Var}}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{1 - \frac{\lambda m}{2(m-1)}} + \frac{K_\alpha^2 \left(2 \text{KL}(\rho \|\pi) + \ln \frac{2k_\lambda k_\gamma}{\delta} \right)}{n\lambda \left(1 - \frac{\lambda m}{2(m-1)} \right)} \right) \right],$$

where $\phi(x) = e^x - x - 1$ and $K_\alpha = \max\{1 - \alpha, 1 - 2\alpha\}$ is the length of the range of the α -tandem loss. The parameter γ is taken in a grid $\{\gamma_1, \dots, \gamma_{k_\gamma}\}$, where $\gamma_i > 0$ for all i and λ is taken in a grid $\{\lambda_1, \dots, \lambda_{k_\lambda}\}$, where $\lambda_i \in \left(0, \frac{2(n-1)}{n}\right)$ for all i . k_γ and k_λ in the bound come from the union bounds over a grid of γ and a grid of λ .

To optimize the bound, we take a grid of $\alpha \in [-0.5, 0.5]$ and iterate over α in the grid. For a given α , we first compute $\hat{L}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$ and $\hat{\text{Var}}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$ for all h, h' . Then, optimize the bound for a fix α by alternating the following steps: 1. Given ρ , starting with $\rho = \pi$, find the corresponding optimal λ , and then the optimal γ in the grids. 2. Given λ and γ , optimize ρ by projected gradient descent.

Next, we present the two new bounds CCPBUB and CCPBSkl, which are based on the oracle parametric form of the Chebyshev-Cantelli bound [Wu et al., 2021, Theorem 8]: For all ρ and for all $\alpha < 0.5$

$$L(\text{MV}_\rho) \leq \frac{\mathbb{E}_{\rho^2}[L_\alpha(h, h')]}{(0.5 - \alpha)^2}. \quad (21)$$

By applying the PAC-Bayes-Unexpected-Bernstein inequality to the α -tandem loss, we obtain the CCPBUB bound, while by applying the PAC-Bayes-split-kl inequality to the α -tandem loss, we obtain the CCPBSkl bound.

Chebyshev-Cantelli bound with PAC-Bayes-Unexpected-Bernstein loss estimate bound (CCPBUB) By applying the PAC-Bayes-Unexpected-Bernstein inequality to the α -tandem loss in equation (21), with the upper bound of the α -tandem loss $b = (1 - \alpha)^2$ for $\alpha < 0.5$, we obtain the bound:

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\mathbb{E}_{\rho^2}[\hat{L}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})] + \frac{\psi(-\gamma(1 - \alpha)^2)}{\gamma(1 - \alpha)^4} \mathbb{E}_{\rho^2}[\hat{\mathbb{V}}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})] \right. \\ \left. + \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{k_\gamma}{\delta}}{\gamma m} \right],$$

where the 2 in front of KL comes from the fact that $\text{KL}(\rho^2 \|\pi^2) = 2 \text{KL}(\rho \|\pi)$. As in the previous experiments, we take a grid of $\gamma \in \{1/(2(1 - \alpha)^2), \dots, 1/(2^{k_\gamma}(1 - \alpha)^2)\}$ for $k_\gamma = \lceil \log_2(\sqrt{m/\ln(1/\delta)}/2) \rceil$ when applying the PAC-Bayes-Unexpected-Bernstein inequality.

Similar to the optimization of the CCPBB bound, we again take a grid of $\alpha \in [-0.5, 0.5]$ and iterate over α in the grid. For a given α , we first compute $\hat{L}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$ and $\hat{\mathbb{V}}_\alpha(h, h', \bar{S}_h \cap \bar{S}_{h'})$ for all h, h' . Then, optimize the bound for a fix α by alternating minimization of ρ and γ in the grid. We initialize $\rho = \pi$ and optimize it by projected gradient descent.

Chebyshev-Cantelli bound with PAC-Bayes-split-kl loss estimate bound (CCPBSkl) Similarly, by applying the PAC-Bayes-split-kl inequality to the α -tandem loss in equation (21), we obtain the

following formula to compute the bound after obtaining the optimal weights ρ :

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\mu + (b - \mu) \text{kl}^{-1,+} \left(\frac{\mathbb{E}_{\rho^2}[\hat{L}_\alpha^+(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{b - \mu}, \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{4\sqrt{m}}{\delta}}{m} \right) \right. \\ \left. - (\mu - a) \text{kl}^{-1,-} \left(\frac{\mathbb{E}_{\rho^2}[\hat{L}_\alpha^-(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{\mu - a}, \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{4\sqrt{m}}{\delta}}{m} \right) \right].$$

We use the following relaxation formula, which is based on the PAC-Bayes- λ inequality F, for optimization.

$$L(\text{MV}_\rho) \leq \frac{1}{(0.5 - \alpha)^2} \left[\mu + (b - \mu) \left(\frac{\mathbb{E}_{\rho^2}[\hat{L}_\alpha^+(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{(b - \mu)(1 - \lambda/2)} + \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{4\sqrt{m}}{\delta}}{\lambda(1 - \lambda/2)m} \right) \right. \\ \left. - (\mu - a) \left(\left(1 - \frac{\gamma}{2}\right) \frac{\mathbb{E}_{\rho^2}[\hat{L}_\alpha^-(h, h', \bar{S}_h \cap \bar{S}_{h'})]}{\mu - a} - \frac{2 \text{KL}(\rho \|\pi) + \ln \frac{4\sqrt{m}}{\delta}}{\gamma m} \right) \right].$$

$2 \text{KL}(\rho \|\pi)$ in both bounds again comes from $\text{KL}(\rho^2 \|\pi^2) = 2 \text{KL}(\rho \|\pi)$. Recall that the α -tandem loss takes values in $\{(1 - \alpha)^2, -\alpha(1 - \alpha), \alpha^2\}$. For $\alpha < 0.5$, $(1 - \alpha)^2$ has the largest value. Therefore, we take $b = (1 - \alpha)^2$ in the bound. Furthermore, for $\alpha < 0$ we have $\alpha^2 < -\alpha(1 - \alpha)$, and for $\alpha \geq 0$ we have $-\alpha(1 - \alpha) \leq \alpha^2$. Therefore, for $\alpha < 0$, we take $a = \alpha^2$ and μ to be the middle value $-\alpha(1 - \alpha)$, while for $\alpha \geq 0$, we take $a = -\alpha(1 - \alpha)$ and $\mu = \alpha^2$.

To optimize the bound, we again take a grid of $\alpha \in [-0.5, 0.5]$ and iterate over α in the grid. For a given α , we compute the parameters a, b, μ , and then compute the losses $\hat{L}_\alpha^+(h, h', \bar{S}_h \cap \bar{S}_{h'})$ and $\hat{L}_\alpha^-(h, h', \bar{S}_h \cap \bar{S}_{h'})$ for all h, h' . The optimization of the bound for a fixed α can be done by alternating minimization: 1. Given ρ , starting with $\rho = \pi$, compute the corresponding optimal λ and γ (Sec. F). 2. Given λ and γ , optimize over ρ using projected gradient descent.

E.3.3 Results

We presented in the body the results of the selected data sets. Here we show the results on more data sets. We present the results for binary data sets in Figure 8, while we present the results for multiclass data sets in Figure 9. Taking $\alpha = 0$ for CCTND, CCPBB, CCPBUB, and CCPBSkl bounds collapses to the TND bound. Therefore, we take TND as a baseline. In both figures, CCTND performs similar to, and often better than the baseline. The second bound, CCPBB, using the α -tandem loss, lags behind due to nested application of concentration bounds. The two new bounds based the α -tandem loss, CCPBUB and CCPBSkl, clearly improve the shortage of the CCPBB bound and often provide tighter bounds than the rest.

E.4 Random Forest Majority Vote Classifiers

In this section, we describe the details of the experimental setting of random forest E.4.1 and the results of the experiments E.4.2. Since both the ensemble of the heterogeneous classifiers and the random forest are examples of weighted majority vote, the bounds and optimization methods in this experiment are the same as described in Sec. E.3.2. The source code for replicating the experiments is available at Github⁶.

E.4.1 Experimental Setting

In this section, we follow the construction used by Wu et al. [2021]. We construct the ensemble from decision trees, which is available in *scikit-learn*. We take 100 fully grown trees to build the random forest. The ensemble is again construct by bagging as described in E.3.1, where each tree h is trained on a subset of a random split S_h and estimated on \bar{S}_h . To train each tree, we use the Gini criterion for splitting and consider \sqrt{d} features in each split, where d is the number of the attributes in data.

⁶<https://github.com/StephanLorenzen/MajorityVoteBounds>

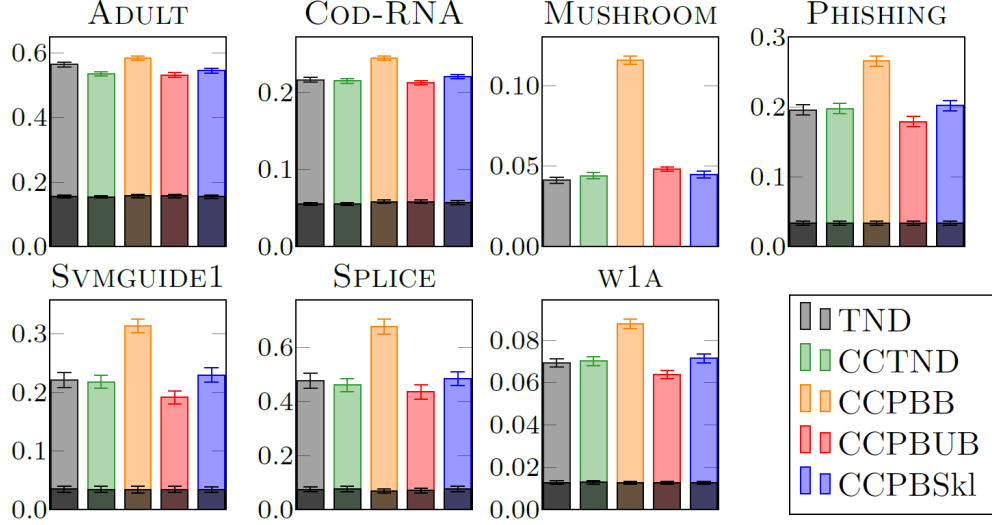


Figure 8: Comparison of the bounds and the test losses of the weighted majority vote on ensembles of heterogeneous classifiers with optimized posterior ρ^* generated by TND, CCTND, CCPBB, CCPBUB, and CCPBSkl. The data sets are binary labeled. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 10 runs of the experiments.

E.4.2 Results

The results of random forest weighted majority vote on binary data sets are shown in Figure 10 while the results on multiclass data sets are shown in Figure 11. Similar to the discussions in Sec. E.3.3 for the ensemble of heterogeneous classifiers, TND serves as a baseline. In both figures, CCTND performs similar to the baseline. The CCPBB, using the α -tandem loss, lags behind due to nested application of concentration bounds. The two new bounds based the α -tandem loss, CCPBUB and CCPBSkl, clearly improve the shortage of the CCPBB bound. The CCPBSkl bound is comparable to the baseline, and the CCPBUB bound often provide tighter bounds than the baseline.

F PAC-Bayes- λ Inequality

Theorem 14 (PAC-Bayes- λ Inequality, Thiemann et al., 2017, Masegosa et al., 2020). *For any loss $\ell \in [0, 1]$, any probability distribution π on \mathcal{H} that is independent of S and any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over a random draw of a sample S , for all distributions ρ on \mathcal{H} and all $\lambda \in (0, 2)$ and $\gamma > 0$ simultaneously:*

$$\mathbb{E}_\rho [L(h)] \leq \frac{\mathbb{E}_\rho[\hat{L}(h, S)]}{1 - \frac{\lambda}{2}} + \frac{\text{KL}(\rho||\pi) + \ln(2\sqrt{n}/\delta)}{\lambda(1 - \frac{\lambda}{2})n}, \quad (22)$$

$$\mathbb{E}_\rho [L(h)] \geq \left(1 - \frac{\gamma}{2}\right) \mathbb{E}_\rho[\hat{L}(h, S)] - \frac{\text{KL}(\rho||\pi) + \ln(2\sqrt{n}/\delta)}{\gamma n}. \quad (23)$$

The upper bound is due to Thiemann et al. [2017] and the lower bound is due to Masegosa et al. [2020], and both hold simultaneously. The PAC-Bayes- λ bound is an optimization friendly relaxation of the PAC-Bayes-kl bound. Tolstikhin and Seldin [2013] has shown that for a given ρ , equation 22 is convex in λ and has the minimizer

$$\lambda_\rho^* = \frac{2}{\sqrt{\frac{2n\mathbb{E}_\rho[\hat{L}(h, S)]}{\text{KL}(\rho||\pi) + \ln \frac{2\sqrt{n}}{\delta}} + 1} + 1}.$$

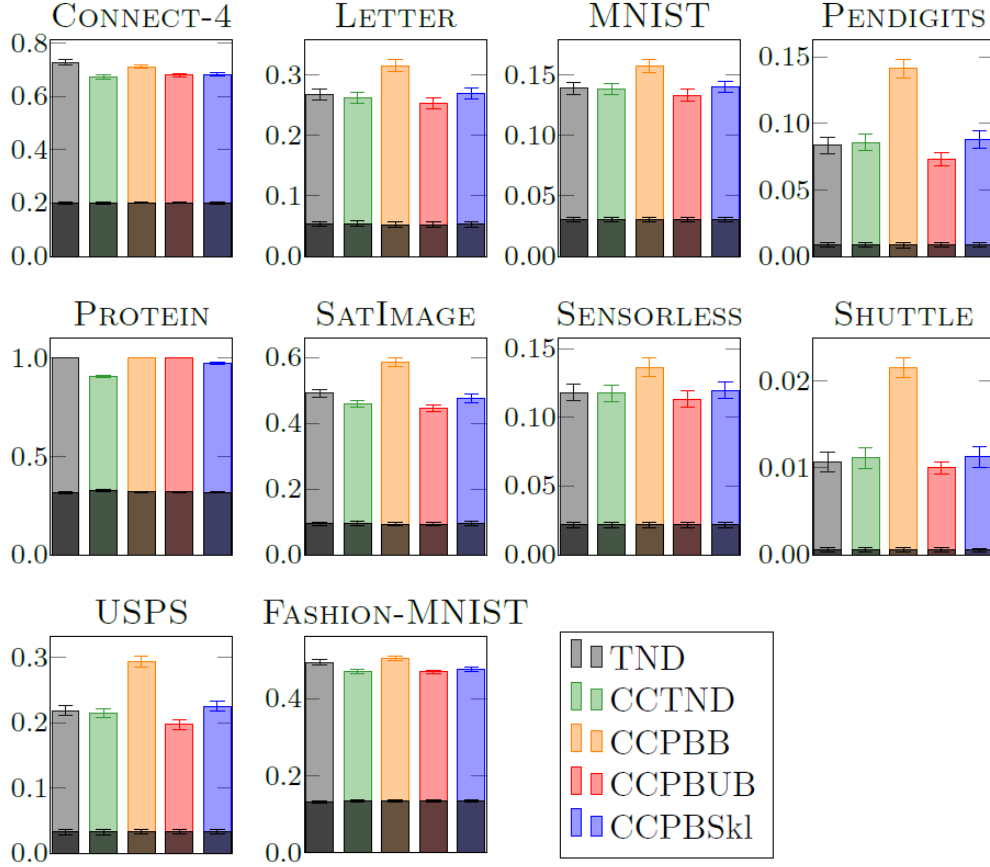


Figure 9: Comparison of the bounds and the test losses of the weighted majority vote on ensembles of heterogeneous classifiers with optimized posterior ρ^* generated by TND, CCTND, CCPBB, CCPBUB, and CCPBSkl. The data sets are multiclass labeled. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 10 runs of the experiments.

On the other hand, Masegosa et al. [2020] has shown that for a given ρ , the optimal γ in equation 23 can be achieved by

$$\gamma_{\rho}^* = \sqrt{\frac{\text{KL}(\rho|\pi) + \ln(2\sqrt{n}/\delta)}{n\mathbb{E}_{\rho}[\hat{L}(h, S)]}}.$$

Furthermore, given λ or γ , the optimal ρ can be achieved by projected gradient descent.

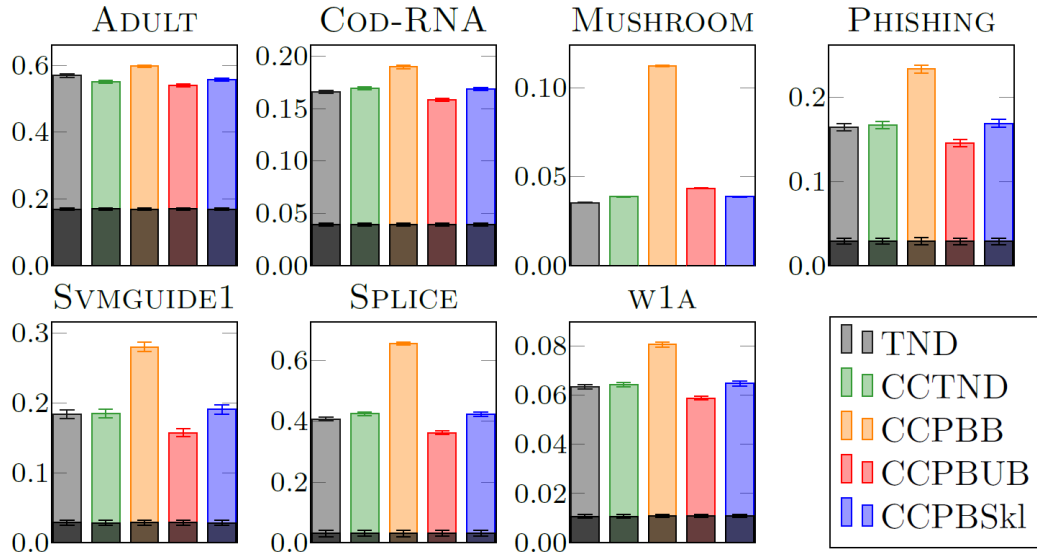


Figure 10: Comparison of the bounds and the test losses of the weighted majority vote on random forest with optimized posterior ρ^* generated by TND, CCTND, CCPBB, CCPBUB, and CCPBSkl. The data sets are binary labeled. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 10 runs of the experiments.

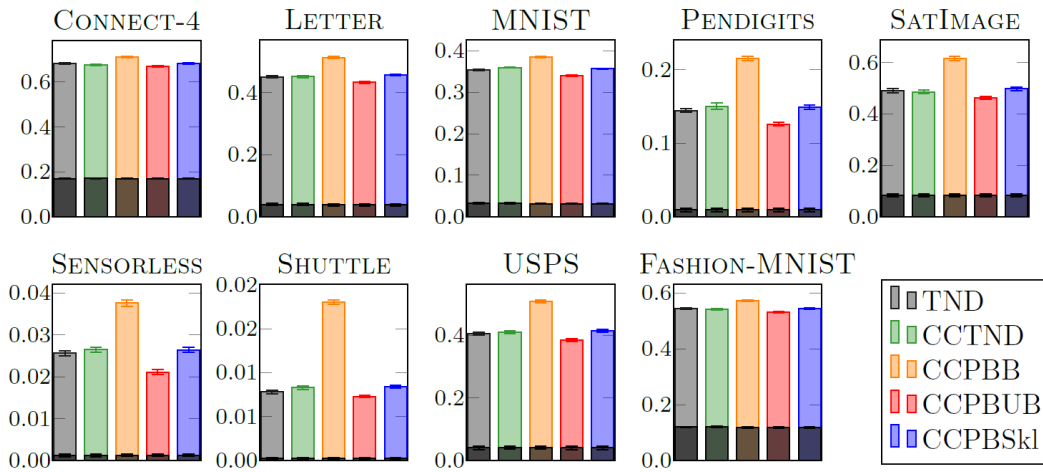


Figure 11: Comparison of the bounds and the test losses of the weighted majority vote on random forest with optimized posterior ρ^* generated by TND, CCTND, CCPBB, CCPBUB, and CCPBSkl. The data sets are multiclass labeled. The test losses of the corresponding bounds are shown in black. We report the mean and the standard deviation over 10 runs of the experiments.