

NONVACUOUS LOSS BOUNDS WITH FAST RATES FOR NEURAL NETWORKS VIA CONDITIONAL INFORMATION MEASURES

Anonymous authors

Paper under double-blind review

ABSTRACT

We present a framework to derive bounds on the test loss of randomized learning algorithms for the case of bounded loss functions. This framework leads to bounds that depend on the conditional information density between the the output hypothesis and the choice of the training set, given a larger set of data samples from which the training set is formed. Furthermore, the bounds pertain to the average test loss as well as to its tail probability, both for the PAC-Bayesian and the single-draw settings. If the conditional information density is bounded uniformly in the size n of the training set, our bounds decay as $1/n$, which is referred to as a fast rate. This is in contrast with the tail bounds involving conditional information measures available in the literature, which have a less benign $1/\sqrt{n}$ dependence. We demonstrate the usefulness of our tail bounds by showing that they lead to estimates of the test loss achievable with several neural network architectures trained on MNIST and Fashion-MNIST that match the state-of-the-art bounds available in the literature.

1 INTRODUCTION

In recent years, there has been a surge of interest in the use of information-theoretic techniques for bounding the loss of learning algorithms. While the first results of this flavor can be traced to the probably approximately correct (PAC)-Bayesian approach (McAllester, 1998; Catoni, 2007) (see also (Guedj, 2019) for a recent review), the connection between loss bounds and classical information-theoretic measures was made explicit in the works of Russo & Zou (2016) and Xu & Raginsky (2017), where bounds on the average population loss were derived in terms of the mutual information between the training data and the output hypothesis. Since then, these average loss bounds have been tightened (Bu et al., 2019; Asadi et al., 2018; Negrea et al., 2019). Furthermore, the information-theoretic framework has also been successfully applied to derive tail probability bounds on the population loss (Bassily et al., 2018; Esposito et al., 2019; Hellström & Durisi, 2020a).

Of particular relevance to the present paper is the *random-subset setting*, introduced by Steinke & Zakyntinou (2020) and further studied in (Hellström & Durisi, 2020b; Haghifam et al., 2020). In this setting, a random vector S is used to select n training samples $Z(S)$ from a larger set \tilde{Z} of $2n$ samples. Then, bounds on the average population loss are derived in terms of the conditional mutual information (CMI) $I(W; S|\tilde{Z})$ between the chosen hypothesis W and the random vector S given the set \tilde{Z} . The bounds obtained by Xu & Raginsky (2017) depend on the mutual information $I(W; Z)$, a quantity that can be unbounded if W reveals too much about the training set Z . In contrast, bounds for the random-subset setting are always finite, since $I(W; S|\tilde{Z})$ is never larger than n bits.

Most information-theoretic population loss bounds mentioned thus far are given by the training loss plus a term with a $\sqrt{\text{IM}(P_{WZ})}/n$ -dependence, where $\text{IM}(P_{WZ})$ denotes an information measure, such as mutual information or maximal leakage (Issa et al., 2020). Assuming that the information measure grows at most polylogarithmically with n , the convergence rate of the population loss to the training loss is $\tilde{O}(1/\sqrt{n})$, where the \tilde{O} -notation hides logarithmic factors. This is sometimes referred to as a *slow rate*. In the context of bounds on the excess risk, defined as the difference between the achieved population loss for a chosen hypothesis w and its infimum over the hypothesis class, it is known that slow rates are optimal for worst-case distributions and hypothesis classes (Talagrand,

1994). However, it is also known that under the assumption of realizability (i.e., the existence of a w in the hypothesis class such that the population loss $L_{P_Z}(w) = 0$) and when the hypothesis class is finite, the dependence on the sample size can be improved to $\tilde{O}(1/n)$ (Vapnik, 1998, Chapter 4). This is referred to as a *fast rate*. Excess risk bounds with fast rates for randomized classifiers have also been derived, under certain additional conditions, for both bounded losses (Van Erven et al., 2015) and unbounded losses (Grünwald & Mehta, 2020).

Notably, Steinke & Zakynthinou (2020, Thm. 2(3)) derive a population loss bound whose dependence on n is $I(W; S|\tilde{Z})/n$. The price for this improved dependence is that the training loss that is added to the n -dependent term is multiplied by a constant larger than 1. Furthermore, (Steinke & Zakynthinou, 2020, Thm. 8) shows that if the Vapnik-Chervonenkis (VC) dimension of the hypothesis class is finite, there exists an empirical risk minimizer (ERM) whose CMI grows at most logarithmically with n . This implies that the CMI approach leads to fast-rate bounds in certain scenarios. However, the result in (Steinke & Zakynthinou, 2020, Thm. 2(3)) pertains only to the average population loss: no tail bounds on the population loss are provided. Throughout the paper, we will, with an abuse of terminology, refer to bounds with an n -dependence of the form $\text{IM}(P_{WZ})/n$ as fast-rate bounds. Such bounds are also known as *linear* bounds (Dziugaite et al., 2020). Note that the n -dependence of the information measure $\text{IM}(P_{WZ})$ has to be at most polylogarithmic for such bounds to actually achieve a fast rate in the usual sense.

An intriguing open problem in statistical learning is to find a theoretical justification for the capability of overparameterized neural networks (NNs) to achieve good generalization performance despite being able to memorize randomly labeled training data sets (Zhang et al., 2017). As a consequence of this behavior, classical population loss bounds that hold uniformly over a given hypothesis class, such as VC bounds, are vacuous when applied to overparameterized NNs. This has stimulated recent efforts aimed at obtaining tighter population loss bounds that are algorithm-dependent or data-dependent.

In the past few years, several studies have shown that promising bounds are attainable by using techniques from the PAC-Bayesian literature (Dziugaite & Roy, 2017; Zhou et al., 2019; Dziugaite et al., 2020). The PAC-Bayesian approach entails using the Kullback-Leibler (KL) divergence to compare the distribution on the weights of the NN induced by training to some reference distribution. These distributions are referred to as the posterior and the prior, respectively. Recently, Dziugaite et al. (2020) used data-dependent priors to obtain state-of-the-art bounds for LeNet-5 trained on MNIST and Fashion-MNIST. In their approach, the available data is used both for training the network and for choosing the prior. This leads to a bound that is tighter than previously available bounds. Furthermore, the bound can be further improved by minimizing the KL divergence between the posterior and the chosen prior during training. One drawback of the PAC-Bayesian approach is that it applies only to stochastic NNs, whose weights are randomly chosen each time the network is used, and not to deterministic NNs with fixed weights.

Information-theoretic bounds have also been derived for iterative, noisy training algorithms such as stochastic gradient Langevin dynamics (SGLD) (Bu et al., 2019). These bounds lead to nonvacuous estimates of the population loss of overparameterized NNs that are trained using SGLD through the use of data-dependent priors (Negrea et al., 2019). However, these bounds do not apply to deterministic NNs, nor to standard stochastic gradient descent (SGD) training. Furthermore, the bounds pertain to the average population loss, and not to its tails. Although the techniques yielding these estimates can be adapted to the PAC-Bayesian setting, as discussed by Negrea et al. (2019, App. I), the resulting bounds are generally loose.

1.1 CONTRIBUTIONS

In this paper, we extend the fast-rate average loss bound by Steinke & Zakynthinou (2020) to the PAC-Bayesian and the single-draw settings. We then use the resulting PAC-Bayesian and single-draw bounds to characterize the test loss of NNs used to classify images from the MNIST and Fashion-MNIST data sets. The single-draw bounds can be applied to deterministic NNs trained through SGD but with Gaussian noise added to the final weights, whereas the PAC-Bayesian bounds apply only to randomized neural networks, whose weights are drawn from a Gaussian distribution each time the network is used. For the same setup, we also evaluate the slow-rate PAC-Bayesian and single-draw bounds from (Hellström & Durisi, 2020b). Our numerical experiments reveal that both the slow-rate

bounds from (Hellström & Durisi, 2020b) and the newly derived fast-rate bounds are nonvacuous. Furthermore, for some settings, the fast-rate bounds presented in this paper are quantitatively stronger than the corresponding slow-rate ones from (Hellström & Durisi, 2020b), and essentially match the best bounds available in the literature for SGD-trained NNs (Dziugaite et al., 2020).

1.2 PRELIMINARIES

We now detail some notation and describe the random-subset setting introduced in (Steinke & Zakynthinou, 2020). Let \mathcal{Z} be the instance space, \mathcal{W} be the hypothesis space, and $\ell : \mathcal{W} \times \mathcal{Z} \rightarrow \mathbb{R}^+$ be the loss function. Throughout the paper, we will assume that the range of $\ell(w, z)$ is restricted to $[0, 1]$ for all $w \in \mathcal{W}$ and all $z \in \mathcal{Z}$. A typical example of such a loss function is the classification error. In this setting, the sample Z consists of an example $X \in \mathcal{X}$ and a corresponding label $Y \in \mathcal{Y}$. Then, the loss is given by $\ell(W, Z) = 1\{f_W(X) \neq Y\}$, where $f_W(\cdot)$ is the map from \mathcal{X} to \mathcal{Y} induced by the hypothesis W . We note that, when applying our bounds to NNs, the function $\ell(\cdot, \cdot)$ used to characterize the performance of the network does not necessarily need to coincide with the loss function used when training the NN. For instance, one could use the (unbounded) cross-entropy loss when training the NN, and apply the bounds for the scenario in which $\ell(\cdot, \cdot)$ is the classification error.

In the random-subset setting, $2n$ training samples $\tilde{Z} = (\tilde{Z}_1, \dots, \tilde{Z}_{2n})$ are available, with all entries of \tilde{Z} being drawn independently from some distribution P_Z on \mathcal{Z} . However, only a randomly selected subset of cardinality n is actually used for training. Following (Steinke & Zakynthinou, 2020), we assume that the training data $Z(S)$ is selected as follows. Let $S = (S_1, \dots, S_n)$ be an n -dimensional random vector, the elements of which are drawn independently from a $\text{Bern}(1/2)$ distribution and are independent of \tilde{Z} . Then, for $i = 1, \dots, n$, the i th training sample in $Z(S)$ is $Z_i(S_i) = \tilde{Z}_{i+S_i n}$. Thus, the binary variable S_i determines whether the training set $Z(S)$ will contain the sample \tilde{Z}_i or the sample \tilde{Z}_{i+n} . The selected training procedure, including the loss function used for training, will determine the conditional distribution $P_{W|Z(S)}$ on the hypothesis class given the training data. For a given $W \sim P_{W|Z(S)}$, we let $L_{Z(S)}(W) = \frac{1}{n} \sum_{i=1}^n \ell(W, Z_i(S_i))$ denote the training loss. Furthermore, we let \bar{S} denote the modulo-2 complement of S . Then $L_{Z(\bar{S})}(W)$ can be interpreted as a test loss, since W is conditionally independent of $Z(\bar{S})$ given $Z(S)$. Finally, we note that the average over (\tilde{Z}, S) of the test loss is the population loss $L_{P_Z}(W) = \mathbb{E}_{P_{\tilde{Z}S}}[L_{Z(\bar{S})}(W)] = \mathbb{E}_{P_Z}[\ell(W, Z)]$.

Our bounds will depend on several different information-theoretic quantities, which we shall introduce next. The information density $\imath(W, Z)$ between W and Z is defined as $\imath(W, Z) = \log \frac{dP_{WZ}}{dP_W P_Z}$, where $\frac{dP_{WZ}}{dP_W P_Z}$ is the Radon-Nikodym derivative of P_{WZ} with respect to $P_W P_Z$. The information density is well-defined if P_{WZ} is absolutely continuous with respect to $P_W P_Z$, denoted by $P_{WZ} \ll P_W P_Z$. The conditional information density $\imath(W, S|\tilde{Z})$ between W and S given \tilde{Z} is defined as $\imath(W, S|\tilde{Z}) = \log \frac{dP_{W\tilde{Z}S}}{dP_{W|\tilde{Z}} P_{\tilde{Z}S}}$, provided that $P_{W\tilde{Z}S} \ll P_{W|\tilde{Z}} P_{\tilde{Z}S}$. The mutual information can be obtained as $I(W; Z) = \mathbb{E}_{P_{WZ}}[\imath(W, Z)]$ and the conditional mutual information as $I(W; S|\tilde{Z}) = \mathbb{E}_{P_{W\tilde{Z}S}}[\imath(W, S|\tilde{Z})]$. We will also need the KL divergences $D(P_{W|Z} \| P_W) = \mathbb{E}_{P_{W|Z}}[\imath(W, Z)]$ and $D(P_{W|\tilde{Z}S} \| P_{W|\tilde{Z}}) = \mathbb{E}_{P_{W|\tilde{Z}S}}[\imath(W, S|\tilde{Z})]$. In practical applications, the marginal distribution P_W is not available, since P_Z is unknown. Furthermore, $P_{W|\tilde{Z}}$ is also difficult to compute, since marginalizing $P_S P_{W|\tilde{Z}S}$ over S involves performing training 2^n times. Hence, bounds depending on $\imath(W, Z)$ or on $\imath(W, S|\tilde{Z})$ cannot typically be evaluated. Therefore, it will be convenient to replace the information density $\imath(W, Z)$ with the proxy $\log \frac{dP_{WZ}}{dQ_W P_Z}$ and $\imath(W, S|\tilde{Z})$ with $\log \frac{dP_{W\tilde{Z}S}}{dQ_{W|\tilde{Z}} P_{\tilde{Z}S}}$. Here, Q_W and $Q_{W|\tilde{Z}}$ are suitably chosen auxiliary distributions (priors) that are used in place of the intractable, true marginals.

2 BACKGROUND

We next review the bounds available in the literature that are relevant for this paper. Then, in Section 3, we will present novel fast-rate bounds. The canonical PAC-Bayesian population-loss bound for a given posterior $P_{W|Z}$ and loss functions bounded between 0 and 1 states that the following holds

with probability at least $1 - \delta$ under P_Z (Guedj & Pujol, 2019, Prop. 3):

$$\mathbb{E}_{P_{W|Z}}[L_{P_Z}(W)] \leq \mathbb{E}_{P_{W|Z}}[L_Z(W)] + \sqrt{\frac{D(P_{W|Z} \parallel Q_W) + \log \frac{1}{\delta}}{2n}}. \quad (1)$$

Here, Q_W is the prior on the hypothesis space \mathcal{W} , which has to be independent of Z . The version of the bound given in (1) slightly improves the dependence on the sample size n as compared to the bound reported in (McAllester, 2003, Thm. 1), at the cost of not holding uniformly for all posteriors. We note that, due to the square root, this is a slow-rate bound. By adapting a proof technique introduced by Catoni (2007, Thm. 1.2.6), McAllester (2013, Eq. (21)) derived the following alternative bound: for all $\gamma \in \mathbb{R}$, and with probability at least $1 - \delta$ under P_Z ,

$$d_\gamma(\mathbb{E}_{P_{W|Z}}[L_Z(W)] \parallel \mathbb{E}_{P_{W|Z}}[L_{P_Z}(W)]) \leq \frac{1}{n} \left(D(P_{W|Z} \parallel Q_W) + \log \frac{1}{\delta} \right). \quad (2)$$

Here, $d_\gamma(q \parallel p) = \gamma q - \log(1 - p + pe^\gamma)$, and one can show that $\sup_\gamma d_\gamma(q \parallel p) = d(q \parallel p)$, where $d(q \parallel p)$ indicates the KL divergence between two Bernoulli distributions with parameters q and p respectively. This bound, with $d_\gamma(q \parallel p)$ replaced by $d(q \parallel p)$, slightly improves the dependence on the sample size n as compared to an earlier bound reported in (Seeger, 2002, Thm. 1), but again at the cost of losing uniformity over posteriors. Let $q = \mathbb{E}_{P_{W|Z}}[L_Z(W)]$ and let c denote the right-hand side of (2). To use the result in (2) to bound the population loss, we need to find

$$p^*(q, c) = \sup\{p : p \in [0, 1], d(q \parallel p) \leq c\}. \quad (3)$$

This is the largest population loss that satisfies the inequality (2). For small q and c , we have $p^*(q, c) \approx c$, which gives us a fast-rate bound. More generally, for any permissible values of q and c , the bound in (2) can be weakened to obtain the following fast-rate bound (McAllester, 2013, Thm. 2): for all $\lambda \in (0, 1)$, and with probability at least $1 - \delta$ under P_Z ,

$$\mathbb{E}_{P_{W|Z}}[L_{P_Z}(W)] \leq \frac{1}{\lambda} \left[\mathbb{E}_{P_{W|Z}}[L_Z(W)] + \frac{D(P_{W|Z} \parallel Q_W) + \log \frac{1}{\delta}}{2(1 - \lambda)n} \right]. \quad (4)$$

Note that the faster decay in n of this bound comes at the price of a multiplication of the training loss and the KL term by a constant that is larger than 1. As a consequence, if the training loss or the KL term are large, this multiplicative constant may make the fast-rate bound in (4) quantitatively worse than the slow-rate bound in (1) for a fixed n . In the so called *interpolating setting*, where the training loss is 0, we can set $\lambda = 1/2$ in (4) and conclude that it is enough for the square-root term in (1) to be smaller than $1/4$ for the fast-rate bound (4) to be tighter than the slow-rate bound (1). [Additional insights on the tightness of these bounds are provided in \(Letarte et al., 2019, Thm. 3\).](#)

We now turn to the random-subset setting introduced by Steinke & Zakynthinou (2020), and described in Section 1.2. In (Steinke & Zakynthinou, 2020, Thm. 2), several bounds on the average population loss are derived for loss functions bounded between 0 and 1, including the following slow-rate and fast-rate bounds:

$$\mathbb{E}_{P_{W\tilde{Z}S}}[L_{P_Z}(W)] \leq \mathbb{E}_{P_{W\tilde{Z}S}}[L_{Z(S)}(W)] + \sqrt{\frac{2I(W; S|\tilde{Z})}{n}} \quad (5)$$

$$\mathbb{E}_{P_{W\tilde{Z}S}}[L_{P_Z}(W)] \leq 2 \mathbb{E}_{P_{W\tilde{Z}S}}[L_{Z(S)}(W)] + \frac{3I(W; S|\tilde{Z})}{n}. \quad (6)$$

Similar to the bound in (4), the price for a fast rate is a multiplicative constant in front of the training loss and the mutual information term. The slow-rate bound in (5) was improved in (Haghifam et al., 2020, Thm. 3.4) by combining the samplewise approach from (Bu et al., 2019) with the disintegration approach in (Negrea et al., 2019), whereby the expectation over \tilde{Z} , which is implicit in the definition of CMI, is pulled outside of the square root. As we detail in the following proposition, the bound in (Haghifam et al., 2020, Thm. 3.4) can be further tightened by also pulling out the expectation with respect to the conditional distribution $P_{W|\tilde{Z}}$ from the square root. The proof of the resulting bound, which is novel, is deferred to Appendix A.1.

Proposition 1. *Consider the random-subset setting described in Section 1.2. Then,*

$$\mathbb{E}_{P_{W\tilde{Z}S}}[L_{P_Z}(W)] \leq \mathbb{E}_{P_{W\tilde{Z}S}}[L_{Z(S)}(W)] + \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{W\tilde{Z}}} \left[\sqrt{2D(P_{S_i|W\tilde{Z}} \parallel P_{S_i})} \right]. \quad (7)$$

We recover (Haghifam et al., 2020, Thm. 3.4) by applying Jensen’s inequality to move the expectation with respect to $P_{W|\tilde{Z}}$ inside the square root. Furthermore, we recover (5) by applying Jensen’s inequality once more to move the remaining expectation over $P_{\tilde{Z}}$ and the empirical average over i inside the square root and by upper-bounding the resulting sum of samplewise CMLs by $I(W; S|\tilde{Z})$.

In (Hellström & Durisi, 2020b), the slow-rate average loss bound (5) was extended to the PAC-Bayesian setting and the single-draw setting through the use of an exponential inequality. Specifically, the following bounds on the test loss $L_{Z(\tilde{S})}(W)$ are derived: with probability at least $1 - \delta$ under $P_{\tilde{Z}S}$,

$$\mathbb{E}_{P_{W|\tilde{Z}S}}[L_{Z(\tilde{S})}(W)] \leq \mathbb{E}_{P_{W|\tilde{Z}S}}[L_{Z(S)}(W)] + \sqrt{\frac{2}{n} \left(D(P_{W|\tilde{Z}S} \| P_{W|\tilde{Z}}) + \log \frac{1}{\delta} \right)}. \quad (8)$$

Furthermore, with probability at least $1 - \delta$ under $P_{W\tilde{Z}S}$,

$$L_{Z(\tilde{S})}(W) \leq L_{Z(S)}(W) + \sqrt{\frac{2}{n} \left(i(W, S|\tilde{Z}) + \log \frac{1}{\delta} \right)}. \quad (9)$$

While the bounds in (8) and (9) pertain to the test loss instead of the population loss, one can obtain population loss bounds by adding a penalty term to (8) and (9), as discussed in (Hellström & Durisi, 2020b, Thm. 2). However, when comparing the bounds to the empirical performance of learning algorithms, the population loss is unknown. Thus, in practice, one has to resort to evaluating a test loss. In Section 3 below, we will derive fast-rate analogues of the tail bounds (8) and (9), again at the price of multiplicative constants.

3 FAST-RATE RANDOM-SUBSET BOUNDS

We now present an exponential inequality from which several test loss bounds can be derived, in a similar manner as was done in (Hellström & Durisi, 2020b). The derivation, which echos part of the proof of (Steinke & Zakynthinou, 2020, Thm. 2.(3)), is provided in Appendix A.2. [This result and its proof illustrate how to combine the exponential-inequality approach from \(Hellström & Durisi, 2020b\) with fast-rate derivations, like those presented in \(Steinke & Zakynthinou, 2020, Thm. 2.\(3\)\) and \(McAllester, 2013, Thm. 2\).](#)

Theorem 1. *Consider the random-subset setting introduced in Section 1.2. Let $W \in \mathcal{W}$ be distributed according to $P_{W|Z(S)} = P_{W|\tilde{Z}S}$. Also, assume that the joint distribution $P_{W\tilde{Z}S} = P_{W|\tilde{Z}S}P_{\tilde{Z}}P_S$ is absolutely continuous with respect to $Q_{W|\tilde{Z}}P_{\tilde{Z}}P_S$ for some conditional prior $Q_{W|\tilde{Z}}$. Then, the following holds:*

$$\mathbb{E}_{P_{W\tilde{Z}S}} \left[\exp \left(\frac{n}{3} L_{Z(S)}(W) - \frac{2n}{3} L_{Z(\tilde{S})}(W) - \log \frac{dP_{W\tilde{Z}S}}{dQ_{W|\tilde{Z}}P_{\tilde{Z}}P_S} \right) \right] \leq 1. \quad (10)$$

Note that the exponential function in (10) depends linearly on the population loss. In contrast, the exponential inequality derived in (Hellström & Durisi, 2020b, Thm. 4) to establish slow-rate generalization bounds for the random-subset setting depends quadratically on the population loss (after the parameter λ therein is suitably optimized). This difference explains why Theorem 1 allows for the derivation of fast-rate bounds, whereas (Hellström & Durisi, 2020b, Thm. 4) unavoidably leads to slow-rate bounds. Also note that, since in the random-subset setting W and \tilde{Z} are dependent both before and after any change of measure argument, the proof technique used in (McAllester, 2013, App. A) and, previously in (Seeger, 2002, Thm. 1), to derive (2) cannot be used in the random-subset setting.

By simple applications of Jensen’s inequality and Markov’s inequality, the exponential inequality (10) can be used to derive bounds on the population loss or test loss. In particular, as detailed in the proof of Corollary 1 below (see Appendix A.3), it can be used to recover (6), but also to establish novel PAC-Bayesian and single-draw versions of (6).

Corollary 1. *Consider the setting of Theorem 1. Then, the average population loss is bounded by*

$$\mathbb{E}_{P_{W\tilde{Z}S}}[L_{P_Z}(W)] \leq 2 \mathbb{E}_{P_{W\tilde{Z}S}}[L_{Z(S)}(W)] + \frac{3 \mathbb{E}_{P_{\tilde{Z}S}}[D(P_{W|\tilde{Z}S} \| Q_{W|\tilde{Z}})]}{n}. \quad (11)$$

Furthermore, with probability at least $1 - \delta$ over $P_{\tilde{\mathbf{Z}}\mathbf{S}}$, the PAC-Bayesian test loss is bounded by

$$\mathbb{E}_{P_{W|\tilde{\mathbf{Z}}\mathbf{S}}} [L_{\mathbf{Z}(\tilde{\mathbf{S}})}(W)] \leq 2\mathbb{E}_{P_{W|\tilde{\mathbf{Z}}\mathbf{S}}} [L_{\mathbf{Z}(\mathbf{S})}(W)] + \frac{3 \left(D(P_{W|\tilde{\mathbf{Z}}\mathbf{S}} \| Q_{W|\tilde{\mathbf{Z}}}) + \log \frac{1}{\delta} \right)}{n}. \quad (12)$$

Finally, with probability at least $1 - \delta$ over $P_{W\tilde{\mathbf{Z}}\mathbf{S}}$, the single-draw test loss is bounded by

$$L_{\mathbf{Z}(\tilde{\mathbf{S}})}(W) \leq 2L_{\mathbf{Z}(\mathbf{S})}(W) + \frac{3 \left(\log \frac{dP_{W\tilde{\mathbf{Z}}\mathbf{S}}}{dQ_{W|\tilde{\mathbf{Z}}}P_{\tilde{\mathbf{Z}}\mathbf{S}}} + \log \frac{1}{\delta} \right)}{n}. \quad (13)$$

Setting $Q_{W|\tilde{\mathbf{Z}}} = P_{W|\tilde{\mathbf{Z}}}$ in (11), we recover the CMI bound in (Steinke & Zakyntinou, 2020) since $\mathbb{E}_{P_{\tilde{\mathbf{Z}}\mathbf{S}}} [D(P_{W|\tilde{\mathbf{Z}}\mathbf{S}} \| P_{W|\tilde{\mathbf{Z}}})] = I(W; \mathbf{S} | \tilde{\mathbf{Z}})$. As illustrated in Corollary 2 below, the bound on the average population loss in (11) can be tightened by replacing the CMI with a sum of samplewise CMIs. The proof of this result, which involves the same argument used to establish Proposition 1, is presented in Appendix A.4.

Corollary 2. *Consider the setting of Theorem 1. Then, the average population loss is bounded by*

$$\mathbb{E}_{P_{W\tilde{\mathbf{Z}}\mathbf{S}}} [L_{P_{\mathbf{Z}}}(W)] \leq 2\mathbb{E}_{P_{W\tilde{\mathbf{Z}}\mathbf{S}}} [L_{\mathbf{Z}(\mathbf{S})}(W)] + \sum_{i=1}^n \frac{3I(W; S_i | \tilde{\mathbf{Z}})}{n}. \quad (14)$$

The bounds in (12) and (13) are data-dependent, i.e., they depend on the specific instances of $\tilde{\mathbf{Z}}$ and \mathbf{S} . They can be turned into data-independent bounds that are functions of the average of the information measures appearing in (12) and (13), at the cost of a less benign polynomial dependence on the confidence parameter δ . Alternatively, one can obtain bounds that have a more benign dependence on δ if one allows the bounds to depend on sufficiently high moments of the information measures appearing in (12) and (13), or if one replaces these measures by quantities such as conditional maximal leakage or conditional α -divergence. See (Hellström & Durisi, 2020b) for further discussion.

We conclude by noting that for the interpolating case where $L_{\mathbf{Z}(\mathbf{S})}(W) = 0$, and under the additional assumption that $Q_{W|\tilde{\mathbf{Z}}} = P_{W|\tilde{\mathbf{Z}}}$, one can obtain a different exponential inequality than the one reported in Theorem 1, which leads to tighter bounds than the ones reported in Corollary 1. Specifically, in these alternative bounds, the factor 3 is replaced with a factor of $1/\log(2) \approx 1.44$. These bounds are presented in Appendix B.

4 EXPERIMENTS

To assess the ability of the bounds just discussed to predict the performance of overparameterized NNs, we next present the result of several numerical experiments for different NN architectures. Specifically, we consider fully connected NNs (FCNNs) and convolutional NNs (CNNs). The performance of the networks is evaluated on the MNIST and Fashion-MNIST data sets. The bounds that we will consider are (8), (9), (12) and (13), where we set the loss function to be the classification error defined in Section 1.2.

The following procedure is used to evaluate the bounds: from the $2n$ available training samples $\tilde{\mathbf{Z}}$ (from MNIST or Fashion-MNIST), the training set $\mathbf{Z}(\mathbf{S})$ is constructed by selecting n training samples uniformly at random. A network is then trained on this data set using a standard SGD procedure, which is described in more detail in Appendix C.2. Let $\boldsymbol{\mu}_1$ be the vector containing the weights of the network after training. The posterior distribution $P_{W|\tilde{\mathbf{Z}}\mathbf{S}}$ is then chosen to be a Gaussian distribution with mean $\boldsymbol{\mu}_1$ and covariance matrix equal to $\sigma_1^2 \mathbf{I}_d$, where d is the number of parameters in the network. The standard deviation σ_1 is chosen as the largest real number, determined to some finite precision (see Appendix C.2), for which the absolute value of the difference between the training loss of the network with weights $\boldsymbol{\mu}_1$ and the empirical average of the training loss achieved by 5 NNs with weights randomly drawn from $\mathcal{N}(W | \boldsymbol{\mu}_1, \sigma_1^2 \mathbf{I}_d)$ is less than some specified threshold. Unless otherwise specified, we use a threshold of 0.05 for MNIST and 0.10 for Fashion-MNIST. Note that this procedure is performed for a fixed $\mathbf{Z}(\mathbf{S})$. Consequently, σ_1 depends on $\tilde{\mathbf{Z}}$ and \mathbf{S} .

To select the prior $Q_{W|\tilde{Z}}$, we proceed as follows. We form 10 subsets of \tilde{Z} , each of size n . The first subset contains the first n samples in \tilde{Z} , the last contains the last n samples in \tilde{Z} , and the remaining subsets contain the linearly spaced sequences in between. We then train one NN on each subset and denote the average of the final weights of these networks by μ_2 . Finally, we choose $Q_{W|\tilde{Z}}$ as a Gaussian distribution with mean μ_2 and covariance matrix $\sigma_2^2 \mathbf{I}_d$. To select σ_2 , we proceed as follows. First, we determine the largest real number $\tilde{\sigma}_2$, again to some finite precision, for which the absolute value of the difference between the training loss of a NN with weights μ_2 and the empirical average of the training loss of 5 NNs with weights drawn from $\mathcal{N}(W | \mu_2, \tilde{\sigma}_2^2 \mathbf{I}_d)$ is below the selected threshold. Note that this time the training loss is evaluated over the entire data set \tilde{Z} , so that there is no dependence on S . We then use $\tilde{\sigma}_2$ to form a set of 27 candidate values for σ_2 , from which we pick the one that results in the tightest bound on the test loss. This procedure, the details of which are given in Appendix C.2, typically results in $\sigma_2 = \sigma_1$. Note that the prior and the posterior distribution satisfy the assumptions needed for the bounds (8), (9), (12) and (13) to hold. Indeed, (μ_1, σ_1) depend on \tilde{Z} only through $Z(S)$, while (μ_2, σ_2) are independent of S but depend on \tilde{Z} .

Equipped with these Gaussian distributions, we evaluate the bounds by noting that, for the chosen prior and posterior, the Radon-Nikodym derivatives in (9) and (13) reduce to likelihood ratios, and the KL divergences in (8) and (12) can be evaluated as

$$D(P_{W|\tilde{Z}S} \| Q_{W|\tilde{Z}}) = \frac{\|\mu_1 - \mu_2\|_2^2}{\sigma_2^2} + d \left(\frac{\sigma_1^2}{\sigma_2^2} + \log \frac{\sigma_2^2}{\sigma_1^2} - 1 \right). \quad (15)$$

Since the MNIST and Fashion-MNIST data sets are fixed and we are unable to draw several data sets from some underlying data distribution, we evaluate our bounds for these particular instances of \tilde{Z} . We do however have control over S , so we run experiments for 10 instances of S and present the resulting mean as well as standard deviation. Note that since we pick the training set uniformly at random from \tilde{Z} , we implicitly randomize over the ordering of the elements of \tilde{Z} .

Our results are obtained by setting $\delta \approx 0.001$ as the confidence parameter. However, since the bounds are optimized over the choice of σ_2 , we need to use a union bound argument (Dziugaite & Roy, 2017; Dziugaite et al., 2020) to guarantee that the final slow-rate and fast-rate bounds hold for all of these candidates simultaneously. As a consequence, the bounds depicted in the figure hold with probability at least 95%. The test loss and training loss are computed empirically by averaging the performance of 5 NNs whose weights are sampled from $\mathcal{N}(W | \mu_1, \sigma_1^2 \mathbf{I}_d)$. For the FCNNs, we use the notation W^L to denote an architecture consisting of L hidden layers with width W . For the case of CNNs, we consider the modified LeNet-5 architecture used in (Zhou et al., 2019) and (Dziugaite et al., 2020). Detailed descriptions of these architectures are provided in Appendix C.1.

In Figure 1, we study the dependence of the bounds on the number of training epochs. The shaded areas around the curves correspond to two standard deviations. The differences between the PAC-Bayesian and the single-draw bounds turn out to be negligible, so we include only the PAC-Bayesian bounds in the figure. The networks are optimized using SGD either with or without momentum. Specifically, in Figures 1a–d, we use SGD without momentum, while in Figures 1e–f, we use SGD with momentum. In Figures 1a–d, we look at the early training epochs, while in Figures 1e–f, we train the networks until a small training loss (on the order of 0.001) is achieved. More details about the training procedures are given in Appendix C.2.

As seen in Figures 1a–d, where SGD without momentum is used, the bounds on the test loss are fairly accurate for both architectures, and both MNIST and Fashion-MNIST. As previously discussed, the relative ordering of the slow-rate bounds and the fast-rate bounds from a quantitative standpoint depends on the details of the learning setup. In particular, higher values for the training loss and information measures tend to make the slow-rate bounds tighter due to their smaller constant factors. As a consequence, the fast-rate bounds are superior for the MNIST data set, for which low training loss and information measures are achieved, while the slow-rate bounds are tighter for the more challenging Fashion-MNIST data set. Note that, due to the training procedure used, the underlying deterministic NNs upon which Figures 1a–d are based never reach training errors below a few percent.

To shed light on the relationship between the results presented in Fig. 1a–d and previously obtained bounds, we compare our bounds on the test loss with those reported in (Dziugaite et al., 2020), which established the best available PAC-Bayesian bounds for the settings we consider. The approach used

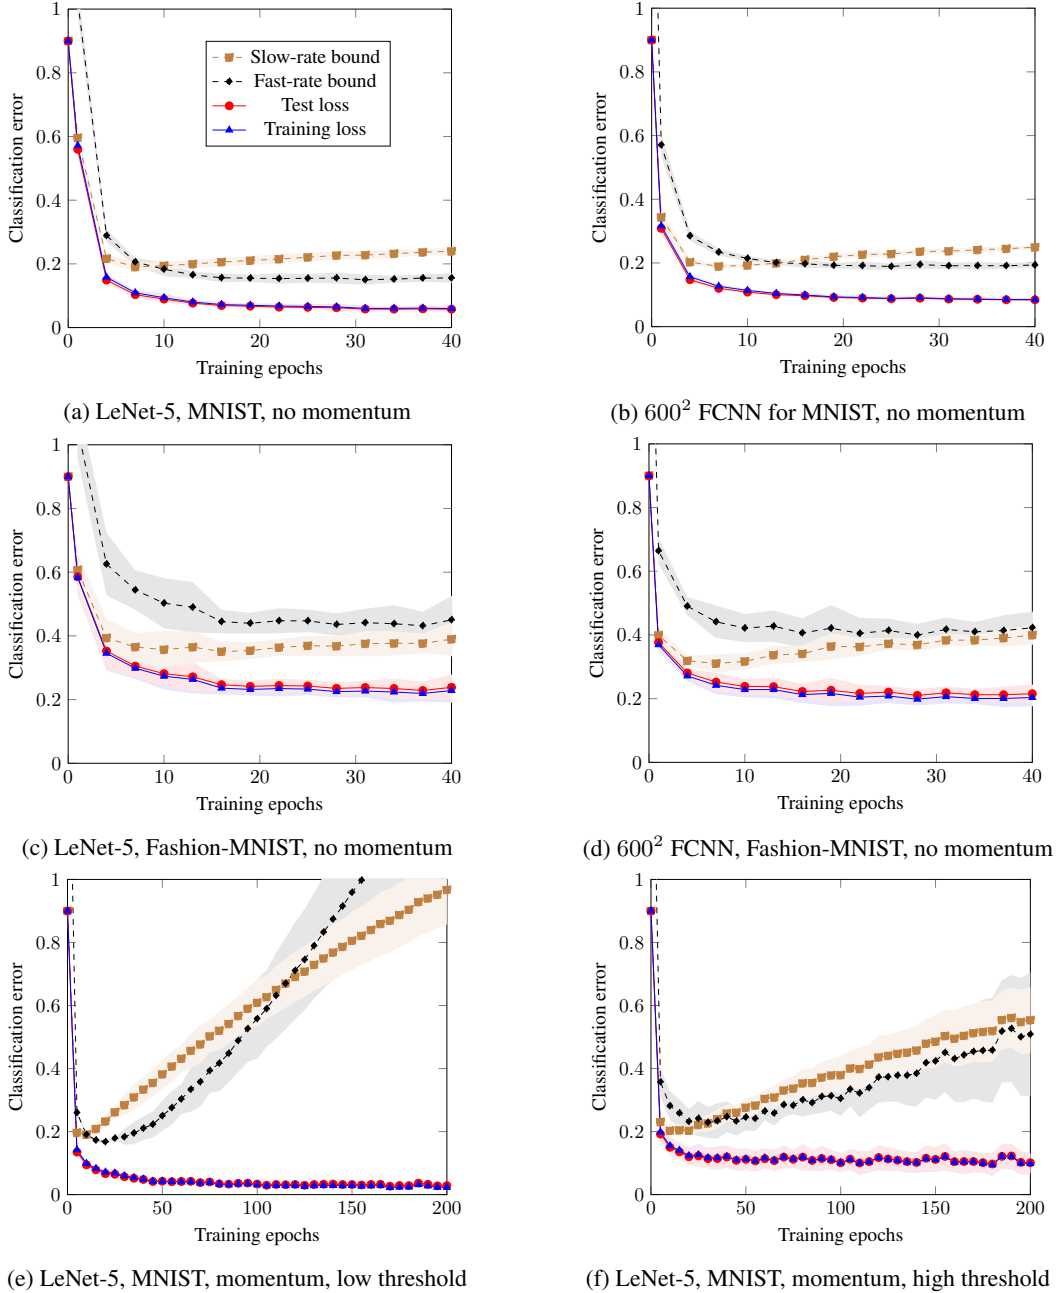


Figure 1: The estimated training losses and test losses as well as the slow-rate (8) and fast-rate (12) PAC-Bayesian bounds on the test loss for two NNs trained on MNIST or Fashion-MNIST. The shaded regions correspond to two standard deviations. In (a)–(d), we perform training using SGD without momentum with a decaying learning rate. In (e)–(f), we use SGD with momentum and a fixed learning rate. Further details on the experimental setup are given in Appendix C.

therein is similar to the random-subset setting considered in this paper, in that the authors make use of a data-dependent prior. The key differences with respect to the framework considered in this paper is that the posterior in Dziugaite et al. (2020) is allowed to depend on the entire data set \tilde{Z} , whereas the training loss and prior depend on randomly selected disjoint subsets of \tilde{Z} . In contrast, in the random-subset setting considered in this paper, the prior is allowed to depend on the entire data set \tilde{Z} , whereas the training loss and posterior depend only on a randomly selected portion $Z(S)$ of \tilde{Z} .

For the case in which training is performed using SGD, the minimum test loss bounds (averaged over 50 runs) for MNIST reported in (Dziugaite et al., 2020, Fig. 4) are approximately 0.13 for LeNet-5 and 0.18 for the 600^2 FCNN. These values are similar to our best bounds, which are 0.15 for LeNet-5 and 0.19 for the 600^2 FCNN. For LeNet-5 trained on Fashion-MNIST, our tightest bound on the test loss is 0.35, whereas the corresponding one in (Dziugaite et al., 2020, Fig. 4) is approximately 0.36. Taking error bars into account, our bounds are not clearly distinguishable from those reported in (Dziugaite et al., 2020, Fig. 4). It is important to mention that significantly tighter bounds are reported in (Dziugaite et al., 2020, Fig. 5) for the case in which the PAC-Bayesian bound considered therein is used as a regularizer during the training process. Such a direct optimization of the bound does not appear to be feasible for the random-subset setting considered in this paper.

Next, we discuss the results presented in Figures 1e–f. As shown in the figure, while our bounds become tighter in the initial phase of training, they lose tightness as training progresses when momentum is used and smaller training errors (on the order of 0.001) are reached for the deterministic NNs. This is similar to what is noted by Dziugaite et al. (2020, p. 12). Specifically, when the underlying deterministic NN therein is trained to achieve very low errors (or equivalently, is trained for many epochs), the PAC-Bayesian bound they consider becomes loose, and the corresponding stochastic NN has a significantly higher test error than the underlying deterministic NN.

Finally, the difference in behavior of our bounds in Figure 1e and Figure 1f illustrates the role played by the variances σ_1 and σ_2 . In Figure 1e, we set the threshold used to determine σ_1 and σ_2 to 0.05, which leads to small values for σ_1 and σ_2 . In Figure 1f, we use a threshold of 0.15 instead, which allows for larger variances. The results illustrate the intuitive fact that larger variances yield better generalization bounds at the cost of a higher true test error.

Further numerical experiments, in which we study how the bounds evolve as a function of the size of the training set, and how they are affected by randomized labels, are reported in Appendix D.

5 CONCLUSION

We have studied information-theoretic bounds on the test loss in the random-subset setting, in which the posterior and the training loss depend on a randomly selected subset of the available data set, and the prior is allowed to depend on the entire data set. In particular, we derived new fast-rate bounds for the PAC-Bayesian and single-draw settings. Provided that the information measures appearing in the bounds scale sublinearly with n , these fast-rate bounds have a better asymptotic dependence on n than the slow-rate PAC-Bayesian and single-draw bounds previously reported in (Hellström & Durisi, 2020b), at the price of larger multiplicative constants. We also presented improvements on previously presented bounds on the average loss by using samplewise information measures and disintegration.

Through numerical experiments, we show that our novel fast-rate PAC-Bayesian bound, as well as its slow-rate counterpart, result in test-loss bounds for some overparameterized NNs trained through SGD that essentially match the best available bounds in the literature (Dziugaite et al., 2020). Furthermore, the single-draw counterparts of these bounds, which are as tight as the PAC-Bayesian bounds, are applicable also to deterministic NNs trained through SGD and with Gaussian noise added to the final weights. On the negative side, as illustrated in Fig. 1e, the bounds turn out to be loose when applied to NNs trained to achieve very small training errors. Moreover, the additional experiments described in Appendix D reveal that the bounds overestimate the number of training samples needed to guarantee generalization, and that they become vacuous when randomized labels are introduced.

Still, the results demonstrate the value of the random-subset approach in studying the generalization capabilities of NNs, and show that fast-rate versions of the available information-theoretic bounds can be beneficial in this setting. In particular, the random-subset setting provides a natural way to select data-dependent priors, namely by marginalizing the learning algorithm $P_{W|\tilde{Z}_S}$ over \mathcal{S} , either exactly or approximately. Such data-dependent priors are a key element in obtaining tight information-theoretic generalization bounds (Dziugaite et al., 2020).

REFERENCES

- A. R. Asadi, E. Abbe, and S. Verdú. Chaining mutual information and tightening generalization bounds. In *Proc. Conf. Neural Inf. Process. Syst. (NeurIPS)*, Montreal, Canada, Dec. 2018.
- R. Bassily, S. Moran, I. Nachum, J. Shafer, and A. Yehudayoff. Learners that use little information. *J. of Mach. Learn. Res.*, 83:25–55, Apr. 2018.
- Y. Bu, S. Zou, and V. V. Veeravalli. Tightening mutual information based bounds on generalization error. In *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Paris, France, July 2019.
- O. Catoni. *PAC-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning*, volume 56. IMS Lecture Notes Monogr. Ser., 2007.
- G.K. Dziugaite and D.M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. In *Proc. Conf. Uncertainty in Artif. Intell. (UAI)*, Sydney, Australia, Aug. 2017.
- G.K. Dziugaite, K. Hsu, W. Gharbieh, and D.M. Roy. On the role of data in PAC-Bayes bounds, June 2020. URL <https://arxiv.org/abs/2006.10929>.
- A.R. Esposito, M. Gastpar, and I. Issa. Generalization error bounds via R nyi f -divergences and maximal leakage. *arXiv*, Dec. 2019. URL <http://arxiv.org/abs/1912.01439>.
- P.D. Gr nwald and N.A. Mehta. Fast rates for general unbounded loss functions: from ERM to generalized Bayes. *J. of Mach. Learn. Res.*, 83:1–80, Mar. 2020.
- B. Guedj. A primer on PAC-Bayesian learning. *arXiv*, Jan. 2019. URL <http://arxiv.org/abs/1901.05353>.
- B. Guedj and L. Pujol. Still no free lunches: the price to pay for tighter PAC-Bayes bounds. *arXiv*, Oct. 2019. URL <http://arxiv.org/abs/1910.04460>.
- M. Haghifam, J. Negrea, A. Khisti, D.M. Roy, and G.K. Dziugaite. Sharpened generalization bounds based on conditional mutual information and an application to noisy, iterative algorithms. *arXiv*, Apr. 2020. URL <http://arxiv.org/abs/2004.12983>.
- F. Hellstr m and G. Durisi. Generalization error bounds via m th central moments of the information density. In *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Los Angeles, CA, June 2020a.
- F. Hellstr m and G. Durisi. Generalization bounds via information density and conditional information density. (June 2020), 2020b. URL <http://arxiv.org/abs/2005.08044>.
- I. Issa, S. Kamath, and A. B. Wagner. An operational approach to information leakage. *IEEE Trans. Inf. Theory*, 66(3):1625–1657, Mar. 2020.
- Ga l Letarte, Pascal Germain, Benjamin Guedj, and Francois Laviolette. Dichotomize and generalize: PAC-Bayesian binary activated deep neural networks. In *Proc. Conf. Neural Inf. Process. Syst. (NeurIPS)*, Vancouver, Canada, Dec 2019.
- D.A. McAllester. Some PAC-Bayesian theorems. In *Proc. Conf. Learn. Theory (COLT)*, pp. 230–234, Madison, WI, July 1998.
- D.A. McAllester. PAC-Bayesian stochastic model selection. *Mach. Learn.*, 51:5–21, Apr. 2003.
- D.A. McAllester. A PAC-Bayesian tutorial with a dropout bound. July 2013. URL <http://arxiv.org/abs/1307.2118>.
- J. Negrea, M. Haghifam, G.K. Dziugaite, A. Khisti, and D.M. Roy. Information-theoretic generalization bounds for SGLD via data-dependent estimates. In *Proc. Conf. Neural Inf. Process. Syst. (NeurIPS)*, Vancouver, Canada, Dec. 2019.
- Y. Polyanskiy and Y. Wu. *Lecture Notes On Information Theory*. 2019. URL <http://www.stat.yale.edu/%7Eyw562/teaching/itlectures.pdf>.

- D. Russo and J. Zou. Controlling bias in adaptive data analysis using information theory. In *Proc. Artif. Intell. Statist. (AISTATS)*, Cadiz, Spain, May 2016.
- M. Seeger. PAC-Bayesian generalisation error bounds for Gaussian process classification. *J. of Mach. Learn. Res.*, 3:233–269, Oct. 2002.
- T. Steinke and L. Zakynthinou. Reasoning about generalization via conditional mutual information. *Conf. Learn Theory (COLT)*, 125:1–16, July 2020.
- M. Talagrand. Sharper bounds for Gaussian and empirical processes. *Ann. Probab.*, 22(1):28–76, Jan. 1994.
- T. Van Erven, P.D. Grünwald, N.A. Mehta, M.D. Reid, and R.C. Williamson. Fast rates in statistical and online learning. *J. of Mach. Learn. Res.*, 16:1793–1861, Sep. 2015.
- V. Vapnik. *Statistical Learning Theory*. John Wiley & Sons, New York, 1998.
- A. Xu and M. Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Proc. Conf. Neural Inf. Process. Syst. (NeurIPS)*, Long Beach, CA, Dec. 2017.
- C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals. Understanding deep learning requires rethinking generalization. In *Proc. Int. Conf. Learn. Representations (ICLR)*, Toulon, France, Apr. 2017.
- W. Zhou, V. Veitch, M. Austern, R.P. Adams, and P. Orbanz. Non-vacuous generalization bounds at the ImageNet scale: a PAC-Bayesian compression approach. In *Proc. Int. Conf. Learn. Representations (ICLR)*, New Orleans, LA, May 2019.

A PROOFS

A.1 PROOF OF PROPOSITION 1

Consider a fixed hypothesis $w \in \mathcal{W}$ and a supersample $\tilde{z} \in \mathcal{Z}^{2n}$. Due to the boundedness of $\ell(\cdot, \cdot)$, the random variable $\widehat{\text{gen}}_i(w, \tilde{z}, S_i) = \ell(w, Z_i(\tilde{S}_i)) - \ell(w, Z_i(S_i))$ is bounded to $[-1, 1]$ for $i = 1, \dots, n$, and it has zero mean. Subgaussianity then implies that the following holds for all $\lambda > 0$:

$$\mathbb{E}_{P_{S_i}}[\exp(\lambda \widehat{\text{gen}}_i(w, \tilde{z}, S_i))] \leq \exp\left(\frac{\lambda^2}{2}\right). \quad (16)$$

Now, let $\mathcal{E}(w, \tilde{z}) \equiv \text{supp}(P_{S_i|w\tilde{z}})$ denote the support of $P_{S_i|w\tilde{z}}$, where $P_{S_i|w\tilde{z}}$ is shorthand for the distribution $P_{S_i|W=w, \tilde{Z}=\tilde{z}}$. Then, with $1_{\mathcal{E}(w, \tilde{z})}$ denoting the indicator function of $\mathcal{E}(w, \tilde{z})$,

$$\mathbb{E}_{P_{S_i}}[1_{\mathcal{E}(w, \tilde{z})} \cdot \exp(\lambda \widehat{\text{gen}}_i(w, \tilde{z}, S_i))] \leq \exp\left(\frac{\lambda^2}{2}\right). \quad (17)$$

Through a change of measure from P_{S_i} to $P_{S_i|w\tilde{z}}$ (Polyanskiy & Wu, 2019, Prop. 17.1), we get (after reorganizing terms)

$$\mathbb{E}_{P_{S_i|w\tilde{z}}} \left[\exp \left(\lambda \widehat{\text{gen}}_i(w, \tilde{z}, S_i) - \frac{\lambda^2}{2} - \log \frac{dP_{S_i|w\tilde{z}}}{dP_{S_i}} \right) \right] \leq 1. \quad (18)$$

We now have a disintegrated, samplewise exponential inequality. Next, we use Jensen’s inequality and then minimize over λ to find that

$$\mathbb{E}_{P_{S_i|w\tilde{z}}}[\widehat{\text{gen}}_i(w, \tilde{z}, S_i)] \leq \min_{\lambda > 0} \frac{\frac{\lambda^2}{2} + \mathbb{E}_{P_{S_i|w\tilde{z}}} \left[\log \frac{dP_{S_i|w\tilde{z}}}{dP_{S_i}} \right]}{\lambda} = \sqrt{2 \mathbb{E}_{P_{S_i|w\tilde{z}}} \left[\log \frac{dP_{S_i|w\tilde{z}}}{dP_{S_i}} \right]}. \quad (19)$$

We now use that $\mathbb{E}_{P_{S_i|w\tilde{z}}} \left[\log \frac{dP_{S_i|w\tilde{z}}}{dP_{S_i}} \right] = D(P_{S_i|w\tilde{z}} \| P_{S_i})$ and then take the expectation with respect to $P_{W\tilde{Z}}$ to find that

$$\mathbb{E}_{P_{W\tilde{Z}}}[\widehat{\text{gen}}_i(W, \tilde{Z}, S_i)] \leq \mathbb{E}_{P_{W\tilde{Z}}} \left[\sqrt{2D(P_{S_i|W\tilde{Z}} \| P_{S_i})} \right]. \quad (20)$$

The desired bound then follows because

$$\mathbb{E}_{P_{W\tilde{Z}S}} [L_{P_Z}(W) - L_{Z(S)}(W)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{W\tilde{Z}S_i}} [\widehat{\text{gen}}_i(W, \tilde{Z}, S_i)] \quad (21)$$

$$\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{W\tilde{Z}}} [\sqrt{2D(P_{S_i|W\tilde{Z}} \| P_{S_i})}]. \quad (22)$$

A.2 PROOF OF THEOREM 1

The proof essentially mimics parts of the derivation of (Steinke & Zakynthinou, 2020, Thm. 2.(3)). For convenience, we begin by proving an exponential inequality for a binary random variable X satisfying $P(X = a) = P(X = b) = 1/2$ where $a, b \in [0, 1]$. Also, let $\bar{X} = b$ if $X = a$ and $\bar{X} = a$ if $X = b$. Finally, let $\lambda, \gamma > 0$ and $c = e^\lambda - 1 - \lambda$. Then,

$$\mathbb{E} [e^{\lambda(X - \gamma\bar{X})}] \leq \mathbb{E} [1 + \lambda(X - \gamma\bar{X}) + c(X - \gamma\bar{X})^2] \quad (23)$$

$$= 1 + \frac{\lambda(1 - \gamma)}{2} (a + b) + \frac{c}{2} (a - \gamma b)^2 + \frac{c}{2} (b - \gamma a)^2. \quad (24)$$

Here, the first inequality follows because $e^y \leq 1 + y + cy^2/\lambda^2$ for all $y \leq \lambda$. Expanding the squares and removing negative terms, we find that

$$\mathbb{E} [e^{\lambda(X - \gamma\bar{X})}] \leq 1 + \frac{\lambda(1 - \gamma)}{2} (a + b) + \frac{c(1 + \gamma^2)}{2} (a^2 + b^2) \quad (25)$$

$$\leq 1 + \lambda(1 - \gamma) + (e^\lambda - 1 - \lambda)(1 + \gamma^2). \quad (26)$$

In view of (10), we are interested in values of λ and γ such that $\lambda(1 - \gamma) + (e^\lambda - 1 - \lambda)(1 + \gamma^2) \leq 0$, so that the left-hand side of (26) is no larger than 1. Furthermore, it will turn out convenient to select pairs (λ, γ) so that λ is as large as possible and γ is as small as possible. A possible choice is $\lambda = 1/3$ and $\gamma = 2$.¹ Thus, we conclude that

$$\mathbb{E} [e^{\frac{1}{3}(X - 2\bar{X})}] \leq 1. \quad (27)$$

Next, we will apply (27) with $X = \ell(w, Z_i(\tilde{S}_i))$ and $\bar{X} = \ell(w, Z_i(S_i))$ for fixed w and \tilde{z} . Note that these random variables satisfy the required assumptions on X and \bar{X} , since the loss function is supported on $[0, 1]$ and the random variables S_i are Bernoulli distributed. Let $Q_{W\tilde{Z}} = Q_{W|\tilde{Z}}P_{\tilde{Z}}$. It then follows that

$$\mathbb{E}_{Q_{W\tilde{Z}}P_S} \left[e^{\frac{n}{3}(L_{Z(\tilde{S})}(W) - 2L_{Z(S)}(W))} \right] = \mathbb{E}_{Q_{W\tilde{Z}}} \left[\prod_{i=1}^n \mathbb{E}_{P_{S_i}} \left[e^{\frac{1}{3}(\ell(W, Z_i(\tilde{S}_i)) - 2\ell(W, Z_i(S_i)))} \right] \right] \leq 1. \quad (28)$$

Now let $\mathcal{E} = \text{supp}(P_{W\tilde{Z}S})$. Then,

$$\mathbb{E}_{Q_{W\tilde{Z}}P_S} \left[1_{\mathcal{E}} \cdot e^{\frac{n}{3}(L_{Z(\tilde{S})}(W) - 2L_{Z(S)}(W))} \right] \leq 1. \quad (29)$$

The desired result follows after a change of measure to $P_{W\tilde{Z}S}$ (Polyanskiy & Wu, 2019, Prop. 17.1).

A.3 PROOF OF COROLLARY 1

We begin by applying Jensen's inequality to (10) to move the expectation inside the exponential. We then obtain (11) by simply taking the logarithm of both sides and reorganizing terms.

To derive (12), we first apply Jensen's inequality in (10), this time only with respect only $P_{W|\tilde{Z}S}$, to get

$$\mathbb{E}_{P_{\tilde{Z}S}} \left[\exp \left(\mathbb{E}_{P_{W|\tilde{Z}S}} \left[\frac{n}{3} L_{Z(S)}(W) - \frac{2n}{3} L_{Z(\tilde{S})}(W) \right] - D(P_{W|\tilde{Z}S} \| Q_{W|\tilde{Z}}) \right) \right] \leq 1. \quad (30)$$

¹Another permissible choice is $\lambda = 1/2.98$ and $\gamma = 1.795$. It turns out that this choice leads to tighter bounds for the setup considered in Section 4. Hence, it will be used in that section.

We now use Markov's inequality in the following form. Let $U \sim P_U$ be a nonnegative random variable satisfying $\mathbb{E}[U] \leq 1$. Then,

$$P_U[U \leq 1/\delta] \geq 1 - \mathbb{E}[U] \delta \geq 1 - \delta. \quad (31)$$

Applying (31) to (30) we find that, with probability at least $1 - \delta$ under $P_{\tilde{\mathbf{z}}\mathbf{s}}$,

$$\exp\left(\mathbb{E}_{P_{W|\tilde{\mathbf{z}}\mathbf{s}}}\left[\frac{n}{3}L_{\mathbf{Z}(\tilde{\mathbf{s}})}(W) - \frac{2n}{3}L_{\mathbf{Z}(\mathbf{s})}(W)\right] - D(P_{W|\tilde{\mathbf{z}}\mathbf{s}} \| Q_{W|\tilde{\mathbf{z}}})\right) \leq \frac{1}{\delta}. \quad (32)$$

Taking the logarithm and reorganizing terms, we obtain (12).

Finally, to derive (13), we apply (31) to (10) immediately to conclude that, with probability at least $1 - \delta$ under $P_{W\tilde{\mathbf{z}}\mathbf{s}}$,

$$\exp\left(\frac{n}{3}L_{\mathbf{Z}(\tilde{\mathbf{s}})}(W) - \frac{2n}{3}L_{\mathbf{Z}(\mathbf{s})}(W) - \log \frac{dP_{W\tilde{\mathbf{z}}\mathbf{s}}}{dQ_{W|\tilde{\mathbf{z}}\mathbf{s}}}\right) \leq \frac{1}{\delta}. \quad (33)$$

The desired bound (13) follows after taking the logarithm and reorganizing terms.

A.4 PROOF OF COROLLARY 2

Consider a fixed $w \in \mathcal{W}$ and $\tilde{\mathbf{z}} \in \mathcal{Z}^{2n}$. As shown in Appendix A.2,

$$\mathbb{E}_{P_{S_i}}\left[e^{\frac{1}{3}(\ell(w, Z_i(\tilde{S}_i)) - 2\ell(w, Z_i(S_i)))}\right] \leq 1. \quad (34)$$

Let $\mathcal{E} = \text{supp}(P_{S_i|w\tilde{\mathbf{z}}})$, where $P_{S_i|w\tilde{\mathbf{z}}}$ is short for $P_{S_i|W=w, \tilde{\mathbf{Z}}=\tilde{\mathbf{z}}}$. By changing measure we get

$$\mathbb{E}_{P_{S_i}}\left[1_{\mathcal{E}} \cdot e^{\frac{1}{3}(\ell(w, Z_i(\tilde{S}_i)) - 2\ell(w, Z_i(S_i)))}\right] = \mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}\left[e^{\frac{1}{3}(\ell(w, Z_i(\tilde{S}_i)) - 2\ell(w, Z_i(S_i))) - \log \frac{dP_{S_i|w\tilde{\mathbf{z}}}}{dP_{S_i}}}\right] \leq 1. \quad (35)$$

Moving the expectation inside the exponential through the use of Jensen's inequality and taking the logarithm, we get

$$\mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}[\ell(w, Z_i(\tilde{S}_i))] \leq 2\mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}[\ell(w, Z_i(S_i))] + 3\mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}\left[\log \frac{dP_{S_i|w\tilde{\mathbf{z}}}}{dP_{S_i}}\right]. \quad (36)$$

$$= 2\mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}[\ell(w, Z_i(S_i))] + 3I(W; S_i | \tilde{\mathbf{Z}}). \quad (37)$$

The desired result now follows by noting that

$$\mathbb{E}_{P_{W\tilde{\mathbf{z}}\mathbf{s}}}[L_{P_{\mathbf{Z}}}(W)] = \mathbb{E}_{P_{W\tilde{\mathbf{z}}}}\left[\frac{1}{n} \sum_{i=1}^n \mathbb{E}_{P_{S_i|w\tilde{\mathbf{z}}}}[\ell(w, Z_i(\tilde{S}_i))]\right] \quad (38)$$

and applying (37) to each term in the sum in (38).

B FAST-RATE BOUNDS FOR THE INTERPOLATING CASE

In this section, we discuss how to tighten the bound in Corollary 1 under the additional assumption that the training loss $L_{\mathbf{Z}(\mathbf{s})}(W)$ is 0 for all $W \sim P_{W|\mathbf{Z}(\mathbf{s})}$ (interpolating assumption), and for the special case $Q_{W|\tilde{\mathbf{z}}} = P_{W|\tilde{\mathbf{z}}}$.

We begin by proving the following exponential inequality, the derivation of which is similar to part of the proof of the realizable fast-rate bound from (Steinke & Zakyntinou, 2020).

Proposition 2. *Consider the setting of Theorem 1, with the additional assumption that $L_{\mathbf{Z}(\mathbf{s})}(W) = 0$ for all $W \sim P_{W|\mathbf{Z}(\mathbf{s})}$ and that $Q_{W|\tilde{\mathbf{z}}} = P_{W|\tilde{\mathbf{z}}}$. Then,*

$$\mathbb{E}_{P_{W\tilde{\mathbf{z}}\mathbf{s}}}\left[\exp\left(n \log 2 \cdot L_{\mathbf{Z}(\tilde{\mathbf{s}})}(W) - \iota(W, \mathbf{S} | \tilde{\mathbf{Z}})\right)\right] \leq 1. \quad (39)$$

Proof. Let $\lambda, \gamma > 0$. Furthermore, let S' be independent of W , \tilde{Z} , and S , and distributed as S . Then,

$$\mathbb{E}_{P_{W\tilde{Z}S}} \left[\prod_{i=1}^n \left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right] \quad (40)$$

$$= \mathbb{E}_{P_{W\tilde{Z}S} P_{S'}} \left[\prod_{i=1}^n e^{\lambda \ell(W, Z_i(\tilde{S}'_i)) - \gamma \ell(W, Z_i(S'_i))} \right] = \mathbb{E}_{P_{W\tilde{Z}} P_S} \left[\prod_{i=1}^n e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} \right]. \quad (41)$$

Let $\mathcal{E} = \text{supp}(P_{W\tilde{Z}S})$. It now follows from (41) that

$$\mathbb{E}_{P_{W\tilde{Z}} P_S} \left[1_{\mathcal{E}} \cdot e^{n(\lambda L_{Z(S)}(W) - \gamma L_{Z(S)}(W))} \right] \leq \mathbb{E}_{P_{W\tilde{Z}} P_S} \left[\prod_{i=1}^n e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} \right] \quad (42)$$

$$= \mathbb{E}_{P_{W\tilde{Z}S}} \left[\prod_{i=1}^n \left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right]. \quad (43)$$

We now change measure to $P_{W\tilde{Z}S}$ to conclude that

$$\begin{aligned} & \mathbb{E}_{P_{W\tilde{Z}S}} \left[e^{n(\lambda L_{Z(S)}(W) - \gamma L_{Z(S)}(W)) - \iota(W, S | \tilde{Z})} \right] \\ & \leq \mathbb{E}_{P_{W\tilde{Z}S}} \left[\prod_{i=1}^n \left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right]. \end{aligned} \quad (44)$$

We now use the interpolating assumption, set $\lambda = \log 2$, and let $\gamma \rightarrow \infty$. These steps, together with the assumption that $\ell(W, Z_i(\tilde{S}_i)) \in [0, 1]$, imply that the right-hand side of (44) is no larger than 1. From this, the desired result follows. \square

Using Proposition 2, we can derive bounds that are analogous to those of Corollary 1. We present these bounds below without proof, since they can be established following steps similar to the ones used to prove Corollary 1.

Corollary 3. *Consider the setting of Proposition 2. Then, the average population loss is bounded by*

$$\mathbb{E}_{P_{W\tilde{Z}S}} [L_{P_Z}(W)] \leq \frac{I(W; S | \tilde{Z})}{n \log 2}. \quad (45)$$

Furthermore, with probability at least $1 - \delta$ over $P_{\tilde{Z}S}$, the PAC-Bayesian population loss is bounded by

$$\mathbb{E}_{P_{W|\tilde{Z}S}} [L_{Z(\tilde{S})}(W)] \leq \frac{D(P_{W|\tilde{Z}S} \| P_{W|\tilde{Z}}) + \log \frac{1}{\delta}}{n \log 2}. \quad (46)$$

Finally, with probability at least $1 - \delta$ over $P_{W\tilde{Z}S}$, the single-draw population loss is bounded by

$$L_{Z(\tilde{S})}(W) \leq \frac{\iota(W, S | \tilde{Z}) + \log \frac{1}{\delta}}{n \log 2}. \quad (47)$$

Finally, we present a samplewise bound that tightens Corollary 2 under the interpolating assumption. Its derivation is inspired by the techniques used to establish Proposition 1 and Proposition 2.

Corollary 4. *Consider the setting of Proposition 2. Then, the average population loss is bounded by*

$$\mathbb{E}_{P_{W\tilde{Z}S}} [L_{P_Z}(W)] \leq \sum_{i=1}^n \frac{I(W; S_i | \tilde{Z})}{n \log 2}. \quad (48)$$

Proof. Let $\lambda, \gamma > 0$ and let S'_i be independent of W , \tilde{Z} , and S_i , and distributed as S_i . Then, for all i ,

$$\mathbb{E}_{P_{W\tilde{Z}S}} \left[\left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right] \quad (49)$$

$$= \mathbb{E}_{P_{W\tilde{Z}S} P_{S'_i}} \left[e^{\lambda \ell(W, Z_i(\tilde{S}'_i)) - \gamma \ell(W, Z_i(S'_i))} \right] = \mathbb{E}_{P_{W\tilde{Z}} P_{S_i}} \left[e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} \right]. \quad (50)$$

We now let $\mathcal{E} = \text{supp}(P_{W\tilde{Z}S_i})$. It follows from (49)–(50) that

$$\begin{aligned} & \mathbb{E}_{P_{W\tilde{Z}S_i}} \left[1_{\mathcal{E}} \cdot e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} \right] \\ & \leq \mathbb{E}_{P_{W\tilde{Z}S}} \left[\left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right]. \end{aligned} \quad (51)$$

By performing a change of measure from $P_{W\tilde{Z}S_i}$ to $P_{W\tilde{Z}S}$ we conclude that

$$\begin{aligned} & \mathbb{E}_{P_{W\tilde{Z}S_i}} \left[e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i)) - \iota(W, S_i|\tilde{Z})} \right] \\ & \leq \mathbb{E}_{P_{W\tilde{Z}S}} \left[\left(\frac{1}{2} e^{\lambda \ell(W, Z_i(\tilde{S}_i)) - \gamma \ell(W, Z_i(S_i))} + \frac{1}{2} e^{\lambda \ell(W, Z_i(S_i)) - \gamma \ell(W, Z_i(\tilde{S}_i))} \right) \right]. \end{aligned} \quad (52)$$

Here, $\iota(W, S_i|\tilde{Z}) = \log \frac{dP_{W\tilde{Z}S_i}}{dP_{W\tilde{Z}S}}$. We now use the interpolating assumption, set $\lambda = \log 2$, and let $\gamma \rightarrow \infty$. These steps, together with the assumption that $\ell(\cdot, \cdot) \in [0, 1]$, imply that the right-hand side of (52) is no larger than 1. Thus,

$$\mathbb{E}_{P_{W\tilde{Z}S_i}} \left[e^{\log 2 \cdot \ell(W, Z_i(\tilde{S}_i)) - \iota(W, S_i|\tilde{Z})} \right] \leq 1. \quad (53)$$

Next, we use Jensen’s inequality to move the expectation in (53) inside the exponential. Taking the logarithm and reorganizing terms, we get

$$\mathbb{E}_{P_{W\tilde{Z}S_i}} [\ell(W, Z_i(\tilde{S}_i))] \leq \mathbb{E}_{P_{W\tilde{Z}S_i}} \left[\frac{\iota(W, S_i|\tilde{Z})}{\log 2} \right] = \frac{I(W; S_i|\tilde{Z})}{\log 2}. \quad (54)$$

The result now follows because

$$\mathbb{E}_{P_{W\tilde{Z}S}} [L_{P_Z}(W)] = \mathbb{E}_{P_{W\tilde{Z}S}} \left[\frac{1}{n} \sum_{i=1}^n \ell(W, Z_i(\tilde{S}_i)) \right] \leq \sum_{i=1}^n \frac{I(W; S_i|\tilde{Z})}{n \log 2}. \quad (55)$$

□

C EXPERIMENT DETAILS

Here, we provide a detailed description of the network architectures and training procedures considered in this paper. We also note that, when evaluating the fast-rate bounds (12) and (13), we use the constants 1.975 and 2.98 in place of 2 and 3, respectively. This choice leads to valid bounds, as pointed out in Appendix A.2.

C.1 NETWORK ARCHITECTURES

The LeNet-5 architecture used in the numerical results is described in Table 1. This is different from most standard implementations of LeNet-5, but coincides with the architecture used by Zhou et al. (2019) and Dziugaite et al. (2020). It has 431 080 parameters. The fully connected neural network denoted by 600^2 consists of an input layer with 784 units, 2 fully connected layers with 600 units and ReLU activations, followed by an output layer with 10 units and softmax activations. It has 837 610 parameters.

C.2 TRAINING PROCEDURES

We now provide additional details on the training procedures described in Section 4. The initial weights of all the networks used for each instance of $\mathbf{Z}(\mathbf{S})$ were set to the same randomly selected values drawn from a zero-mean normal distribution with standard deviation 0.01. All networks were trained using the cross-entropy loss, optimized using either SGD with momentum and a fixed learning rate or SGD without momentum and a decaying learning rate. First, we describe the details of SGD with momentum. For MNIST, we used a learning rate of 0.001, and for Fashion-MNIST, we used 0.003. In all experiments, the momentum parameter is set to 0.9. We used a batch size of 512.

Table 1: The LeNet-5 architecture used in Section 4.

Convolutional layer, 20 units, 5×5 size, linear activation, 1×1 stride, valid padding
Max pooling layer, 2×2 size, 2×2 stride
Convolutional layer, 50 units, 5×5 size, linear activation, 1×1 stride, valid padding
Max pooling layer, 2×2 size, 2×2 stride
Flattening layer
Fully connected layer, 500 units, ReLU activation
Fully connected layer, 10 units, softmax activation

For SGD without momentum, we used a decaying learning rate schedule, where the learning rate α for a given epoch E is given by

$$\alpha(E) = \frac{\alpha_0}{1 + \gamma \cdot \lfloor E/E_0 \rfloor}. \quad (56)$$

Here, α_0 is the initial learning rate, γ is the decay rate, and E_0 is the number of epochs between each decay. In all experiments, we used $\alpha_0 = 0.01$, $\gamma = 2$, and $E_0 = 20$. Again, we used a batch size of 512.

To choose σ_1 , we pick the largest value with one significant digit (i.e., of the form $a \cdot 10^{-b}$ with $a \in [1 : 9]$ and $b \in \mathbb{Z}$) such that the absolute value of the difference between the training loss on $\mathbf{Z}(\mathcal{S})$ of the deterministic network with weights $\boldsymbol{\mu}_1$ and empirical average of the training loss of 5 NNs with weights drawn independently from $\mathcal{N}(\mathbf{W} \mid \boldsymbol{\mu}_1, \sigma_1^2 \mathbf{I}_d)$ was no larger than some specified threshold.

When producing the results reported in Figure 2 and Figures 1a–d, we used a threshold of 0.05 for MNIST, while for Fashion-MNIST, we used a threshold of 0.10. In Appendix D, we perform additional experiments with other thresholds. Specifically, for Figure 1e, we use a threshold of 0.05, while we use a threshold of 0.15 for Figure 1f. For the randomized label experiment in Table 2, we use a threshold of 0.10.

To find σ_2 , we use as starting point the same procedure as for determining σ_1 , but with $\boldsymbol{\mu}_2$ in place of $\boldsymbol{\mu}_1$ and the training loss evaluated on all of $\tilde{\mathbf{Z}}$. Let us call the value found by this procedure $\sigma'_2 = a' \cdot 10^{-b'}$. Then, among the values of the form $a \cdot 10^{-b}$ with $a \in [1 : 9]$ and $b \in \{b' - 1, b', b' + 1\}$, we choose σ_2 to be the one that minimizes the bound on the test loss. In all our experiments, this procedure resulted in $\sigma_2 = \sigma_1$. To guarantee that the final bound holds with a given confidence level, all 27 bounds resulting from all possible choices of a and b need to hold with the same confidence level. Since we consider both slow-rate and fast-rate bounds, a total of 54 bounds need to hold simultaneously. We ensure that this is the case via the union bound. Thus, if each individual bound holds with probability at least $1 - \delta$, the optimized bounds hold with probability at least $1 - 54\delta$. We compute the bounds with $\delta = 0.05/54$, so the optimized bounds hold with 95% confidence.

D ADDITIONAL EXPERIMENTS

D.1 DEPENDENCE ON THE SIZE OF THE TRAINING SET

In this section, we study the dependence of the bounds on the size n of the training set. We perform experiments for different values of n by restricting $\tilde{\mathbf{Z}}$ to be a $2n$ -dimensional randomly chosen subset of the set of $6 \cdot 10^4$ training samples available in MNIST and Fashion-MNIST. The training set $\mathbf{Z}(\mathcal{S})$ is then formed by selecting n of these samples at random. We then train a network, either LeNet-5 or the 600² FCNN, on this restricted training set, until the training error is lower than some target error. For MNIST, we use a target training error of 0.05, while we use 0.15 for Fashion-MNIST. The results are shown in Figure 2.

As seen in Figure 2, the bounds on the test loss for large values of n are fairly accurate for both of these architectures, especially so for MNIST. However, they are loose for smaller values of n . As previously discussed, the relative ordering of the slow-rate bounds and the fast-rate bounds from a quantitative standpoint depends on the details of the learning setup. In particular, higher values for the training loss and information measures tend to make the slow-rate bounds tighter due to their smaller constant factors.

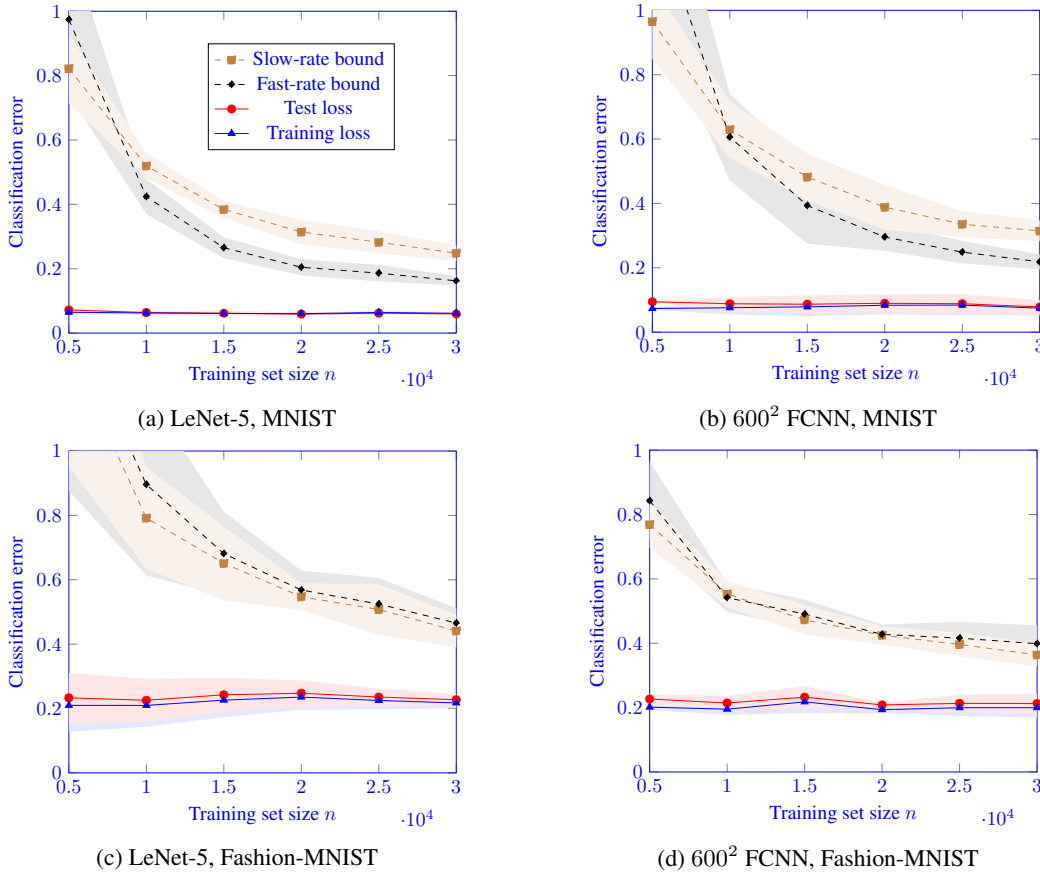


Figure 2: The estimated training losses and test losses as well as the slow-rate (8) and fast-rate (12) PAC-Bayesian bounds on the test loss for two NNs trained on MNIST or Fashion-MNIST. These slow-rate and fast-rate PAC-Bayesian bounds essentially coincide with the single-draw bounds in (9) and (13) respectively. The shaded regions correspond to two standard deviations.

Table 2: The estimated training errors, test errors, and the corresponding slow-rate (8) and fast-rate (12) PAC-Bayesian generalization bounds for LeNet-5 trained on binarized MNIST with partially corrupted labels.

Randomized labels	25%	50%	75%	100%
Training error	0.106	0.088	0.090	0.081
Test error	0.216	0.364	0.461	0.494
Slow-rate bound	5.561	9.811	10.45	11.67
Fast-rate bound	44.52	141.0	160.1	200.4

D.2 RANDOMIZED LABELS

In order to examine the behavior of our bounds in an overfitting scenario, we consider data sets with partially randomized labels. Specifically, we set the labels of a fixed proportion of both the training and test sets of MNIST uniformly at random, and then perform training using SGD with momentum as described in Appendix C. In order to simplify training with randomized labels, we consider a binarized version of MNIST where the digits $0, \dots, 4$ are merged into one class and $5, \dots, 9$ into another. The results are shown in Table 2. The slow-rate bound is computed using (8), while the fast-rate bound is based on (12). As usual, the quantitative difference between these bounds and the corresponding single-draw bounds in (9) and (13) is negligible.

As shown in Table 2, our bounds become vacuous when randomized labels are used. The fast-rate bound is significantly worse than its slow-rate counterpart, which is to be expected: when the prior and posterior are selected using randomized labels, a larger discrepancy between them arises. This increases the value of the KL divergence in (8) and (12), which, as previously discussed, penalizes the fast-rate bound more. We note, though, that the qualitative behavior of the bounds is in agreement with the empirically evaluated test error: an increased proportion of randomized labels, and thus an increased test error, is accompanied by an increase in the values of our bounds. Furthermore, the slow-rate bound consistently overestimates the test error by a factor of approximately 25.

To the best of our knowledge, all bounds available in the literature for overfitting situations such as the one considered in this section are vacuous. The best result can be found in (Dziugaite & Roy, 2017, Tab. 1), where an FCNN with one hidden layer is trained on a binarized version of MNIST with fully randomized labels. Despite directly optimizing the evaluated PAC-Bayesian bound as part of the training procedure, the obtained test error bound of 1.365 is vacuous.