
Sharpened Generalization Bounds based on Conditional Mutual Information and an Application to Noisy, Iterative Algorithms

Mahdi Haghifam
University of Toronto,
Vector Institute

Jeffrey Negrea
University of Toronto,
Vector Institute

Ashish Khisti
University of Toronto

Daniel M. Roy
University of Toronto,
Vector Institute

Gintare Karolina Dziugaite
Element AI

Abstract

The information-theoretic framework of Russo and Zou (2016) and Xu and Raginsky (2017) provides bounds on the generalization error of a learning algorithm in terms of the mutual information between the algorithm’s output and the training sample. In this work, we study the proposal by Steinke and Zakynthinou (2020) to reason about the generalization error of a learning algorithm by introducing a super sample that contains the training sample as a random subset and computing mutual information conditional on the super sample. We first show that these new bounds based on the conditional mutual information are tighter than those based on the unconditional mutual information. We then introduce yet tighter bounds, building on the “individual sample” idea of Bu et al. (2019) and the “data dependent” ideas of Negrea et al. (2019), using disintegrated mutual information. Finally, we apply these bounds to the study of the Langevin dynamics algorithm, showing that conditioning on the super sample allows us to exploit information in the optimization trajectory to obtain tighter bounds based on hypothesis tests.

1 Introduction

Let \mathcal{D} be an unknown distribution on a space \mathcal{Z} , and let \mathcal{W} be a set of parameters that index a set of predictors, $\ell : \mathcal{Z} \times \mathcal{W} \rightarrow [0, 1]$ be a bounded loss function. Consider a (randomized) learning algorithm \mathcal{A} that selects an element W in \mathcal{W} , based on an IID sample $S = (Z_1, \dots, Z_n) \sim \mathcal{D}^{\otimes n}$. For $w \in \mathcal{W}$, let $R_{\mathcal{D}}(w) = \mathbb{E}\ell(Z, w)$ denote the risk of predictor w , and $\hat{R}_S(w) = \frac{1}{n} \sum_{i=1}^n \ell(Z_i, w)$ denote the empirical risk. Our interest in this paper is the (expected) generalization error of \mathcal{A} with respect to \mathcal{D} ,

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) = \mathbb{E}[R_{\mathcal{D}}(W) - \hat{R}_S(W)].$$

In this work, we study bounds on generalization error in terms of information-theoretic measures of dependence between the data and the output of the learning algorithm. This approach was initiated by Russo and Zou [18, 19] and has since been extended [2, 3, 6, 9, 17, 24]. The basic result in this line of work is that the generalization error can be bounded in terms of the mutual information $I(W; S)$ between the data and the learned parameter, a quantity that has been called the *information usage* or *input–output mutual information of \mathcal{A} with respect to \mathcal{D}* , which we denote by $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$. The following result is due to Russo and Zou [18] and Xu and Raginsky [24].

Theorem 1.1. $\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \sqrt{\frac{\text{IOMI}_{\mathcal{D}}(\mathcal{A})}{2n}}.$

Theorem 1.1 formalizes the intuition that a learning algorithm without heavy dependence on the training set will generalize well. This result has been extended in many directions: Raginsky et al. [17] connect variants of $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ to different notions of stability. Asadi et al. [3] establish refined bounds using chaining techniques for subgaussian processes. Bu et al. [6] obtain a tighter bound by replacing $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ with the mutual information between W and a single training data point. Negrea et al. [15] propose variants that allow for data-dependent estimates. See also [1, 2, 4, 9, 10].

Our focus in this paper is on a new class of information-theoretic bounds on generalization error, proposed by Steinke and Zakynthinou [21]. Fix $k \geq 2$, let $[k] = \{1, \dots, k\}$, let $U^{(k)} = (U_1, \dots, U_n) \sim \text{Unif}([k]^n)$, and let $\tilde{Z}^{(k)} \sim \mathcal{D}^{\otimes(k \times n)}$ be a $k \times n$ array of IID random elements in \mathcal{Z} , independent from $U^{(k)}$. Let $S = (Z_{U_{1,1}}, \dots, Z_{U_{n,n}})$ and let W be a random element in \mathcal{W} such that conditional on $S, U^{(k)}$, and $\tilde{Z}^{(k)}$, W has distribution $\mathcal{A}(S)$. It follows that, conditional on S , W is independent from $U^{(k)}$ and $\tilde{Z}^{(k)}$. By construction, the data set S is hidden inside the super sample; the indices $U^{(k)}$ specify where. Steinke and Zakynthinou [21] use these additional structures to define:

Definition 1.2. The conditional mutual information of \mathcal{A} w.r.t. \mathcal{D} is $\text{CMI}_{\mathcal{D}}^k(\mathcal{A}) = I(W; U^{(k)} | \tilde{Z}^{(k)})$.

Intuitively, $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ captures how well we can recognize which samples from the given super-sample $\tilde{Z}^{(k)}$ were in the training set, given the learned parameters. This intuition and the connection of $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ with the membership attack [20] can be formalized using Fano’s inequality, showing that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ can be used to lower bound the error of any estimator of $U^{(k)}$ given W and $\tilde{Z}^{(k)}$. (See Appendix A.) Steinke and Zakynthinou [21] connect $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ with well-known notions in learning theory such as distributional stability, differential privacy, and VC dimension, and establish the following bound [21, Thm. 5.1] in the case $k = 2$, the extension to $k \geq 2$ being straightforward:

Theorem 1.3. $\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \sqrt{\frac{2\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{n}}$.

This paper improves our understanding of the framework introduced by Steinke and Zakynthinou [21], identifies tighter bounds, and applies these techniques to the analysis of a real algorithm. In Section 2, we present several formal connections between the two aforementioned information-theoretic approaches for studying generalization. Our first result bridges $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ and $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$, showing that for any learning algorithm, any data distribution, and any k , $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ is less than $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$. We also show that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ converges to $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ as $k \rightarrow \infty$ when $|\mathcal{W}|$ is finite. In Section 3, we establish two novel bounds on generalization error using the random index and super sample structure of Steinke and Zakynthinou, and show that both our bounds are tighter than those based on $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$. Finally, in Section 4, we show how to construct generalization error bounds for noisy, iterative algorithms using the generalization bound proposed in Section 3. Using the Langevin dynamics algorithm as our example, we introduce a new type of prior for iterative algorithms that “learns” from the past trajectory, using a form of *hypothesis testing*, in order to not “pay” again for information obtained at previous iterations. Experiments show that our new bound is tighter than [14, 15], especially in the late stages of training, where the hypothesis test component of the bound discounts the contributions of new gradients. Our new bounds are non-vacuous for a great deal more epochs than related work, and do not diverge or exceed 1 even when severe overfitting occurs.

1.1 Contributions

1. We characterize the connections between the $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ and $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$. We show that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ is always less than the $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ for any data distribution, learning algorithms and k . Further, we prove that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ converges to $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ when k goes to infinity for finite parameter spaces.
2. We provide novel generalization bounds that relate generalization to the mutual information between learned parameters and a random subset of the random indices U_1, \dots, U_n .
3. We apply our generalization bounds to the Langevin dynamics algorithm by constructing a specific *generalized prior and posterior*. We employ a generalized prior that learns about the values of the indices U from the optimization trajectory. To our knowledge, this is the first generalized prior that learns about the dataset from the iterates of the learning algorithm.
4. We show empirically that our bound on the expected generalization error of Langevin dynamics algorithm is tighter than other existing bounds in the literature.

1.2 Definitions from Probability and Information Theory

Let \mathcal{S}, \mathcal{T} be measurable spaces, let $\mathcal{M}_1(\mathcal{S})$ be the space of probability measures on \mathcal{S} , and define a probability kernel from \mathcal{S} to \mathcal{T} to be a measurable map from \mathcal{S} to $\mathcal{M}_1(\mathcal{T})$. For random elements X in \mathcal{S} and Y in \mathcal{T} , write $\mathbb{P}[X] \in \mathcal{M}_1(\mathcal{S})$ for the distribution of X and write $\mathbb{P}^Y[X]$ for (a regular version of) the conditional distribution of X given Y , viewed as a $\sigma(Y)$ -measurable random element in $\mathcal{M}_1(\mathcal{S})$. Recall that $\mathbb{P}^Y[X]$ is a regular version if, for some probability kernel κ from \mathcal{T} to \mathcal{S} , we have $\mathbb{P}^Y[X] = \kappa(Y)$ a.s. . If Y is $\sigma(X)$ -measurable then Y is a function of X . If random measure, P , is $\sigma(X)$ -measurable then the measure P is determined by X , but a random element Y with $\mathbb{P}^X[Y] = P$ is not X measurable unless it is degenerate. If X is a random variable, write $\mathbb{E}X$ for the expectation of X and write $\mathbb{E}^Y X$ or $\mathbb{E}[X|Y]$ for (an arbitrary version of) the conditional expectation of X given Y , which is Y -measurable. For a random element X on \mathcal{S} and a probability kernel P from \mathcal{S} to \mathcal{T} , the composition $P(X) := P \circ X$ is a $\sigma(X)$ -measurable random measure of a random element taking values in \mathcal{T} . We occasionally use this notation to refer to a kernel P implicitly by the way it acts on X .

Let P, Q be probability measures on a measurable space \mathcal{S} . For a P -integrable or nonnegative measurable function f , let $P[f] = \int f dP$. When Q is absolutely continuous with respect to P , denoted $Q \ll P$, we write $\frac{dQ}{dP}$ for the Radon–Nikodym derivative of Q with respect to P . We rely on several notions from information theory: The *KL divergence of Q with respect to P* , denoted $\text{KL}(Q \| P)$, is $Q[\log \frac{dQ}{dP}]$ when $Q \ll P$ and ∞ otherwise. Let X, Y , and Z be random elements, and let \otimes form product measures. The *mutual information between X and Y* is $I(X; Y) = \text{KL}(\mathbb{P}[(X, Y)] \| \mathbb{P}[X] \otimes \mathbb{P}[Y])$. The *disintegrated mutual information between X and Y given Z* , is¹

$$I^Z(X; Y) = \text{KL}(\mathbb{P}^Z[(X, Y)] \| \mathbb{P}^Z[X] \otimes \mathbb{P}^Z[Y]).$$

The *conditional mutual information* of X and Y given Z is $I(X; Y|Z) = \mathbb{E}I^Z(X, Y)$.

2 Connections between $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ and $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$

In this section, we compare approaches for the information-theoretic analysis of generalization error, and we aim to unify the two main information-theoretic approaches for studying generalization. In Theorems 2.1 and 2.2 we will show that for any learning algorithm and any data distribution, $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ provides a tighter measure of dependence than $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$, and that one can recover $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ -based bounds from $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ for finite parameter spaces.

A fundamental difference between $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ and $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ is that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ is bounded by $n \log k$ [21], while $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ can be infinite even for learning algorithms that provably generalize [6]. One of the motivations of Steinke and Zakynthinou was that proper empirical risk minimization algorithms over threshold functions on \mathbb{R} have large $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ [4]. In contrast, some such algorithms have small $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$. Our first result shows that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ is never larger than $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$.

Theorem 2.1. *For every $k \geq 2$, $I(W; S) = I(W; \tilde{Z}^{(k)}) + I(W; U^{(k)} | \tilde{Z}^{(k)})$ and*

$$\text{CMI}_{\mathcal{D}}^k(\mathcal{A}) \leq \text{IOMI}_{\mathcal{D}}(\mathcal{A}).$$

Next, we address the role of the size of the super-sample in CMI. In [21], CMI is defined using a super-sample of size $2n$ ($k = 2$) only. Our next result demonstrates that $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ agree $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ in the limit as $k \rightarrow \infty$ when the parameter space is finite.

Theorem 2.2. *If the output of \mathcal{A} takes value in a finite set then*

$$\lim_{k \rightarrow \infty} \text{CMI}_{\mathcal{D}}^k(\mathcal{A}) = \text{IOMI}_{\mathcal{D}}(\mathcal{A}).$$

Combining Theorems 1.3 and 2.2, we obtain

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \lim_{k \rightarrow \infty} \sqrt{\frac{2\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{n}} = \sqrt{\frac{2\text{IOMI}_{\mathcal{D}}(\mathcal{A})}{n}}, \quad (1)$$

when the parameter space is finite. Comparing Eq. (1) with Theorem 1.1 we observe that Eq. (1) is twice as large. In Theorem B.1, we present a refined bound based on $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ which asymptotically match Theorem 1.1. The proofs of the results of this section appear in Appendix C.

¹ Letting ϕ satisfy $\phi(Z) = I^Z(X; Y)$ a.s., define $I(X, Y|Z = z) = \phi(z)$. This notation is necessarily well defined only up to a null set under the marginal distribution of Z .

3 Sharpened Bounds based on Individual Samples

We now present two novel generalization bounds and show they provide a tighter characterization of the generalization error than Theorem 1.3. The results are inspired by the improvements on $\text{IOMI}_{\mathcal{D}}(\mathcal{A})$ due to Bu et al. [6]. In particular, Theorem 3.1 bounds the expected generalization error in terms of the mutual information between the output parameter and a random subsequence of the indices $U^{(2)}$, given the super-sample. Theorem 3.4 provides a generalization bound in terms of the disintegrated mutual information between each individual element of $U^{(2)}$ and the output of the learning algorithm, W . The bound in Theorem 3.4 is an analogue of [6, Prop. 1] for Theorem 1.3. In this section as in Steinke and Zakynthinou [21], we only consider $\tilde{Z}^{(k)}$ and $U^{(k)}$ with $k = 2$, so we will drop the superscript from $U^{(k)}$. Let $U = (U_1, \dots, U_n)$. The proofs for the results of this section appear in Appendix D.

Theorem 3.1. *Fix $m \in [n]$ and let $J = (J_1, \dots, J_m)$ be a random subset of $[n]$, distributed uniformly among all subsets of size m and independent from W , $\tilde{Z}^{(2)}$, and U . Then*

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \mathbb{E} \sqrt{\frac{2I^{\tilde{Z}^{(2)}}(W; U_J | J)}{m}}. \quad (2)$$

By applying Jensen’s inequality to Theorem 3.1, we obtain

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \sqrt{\frac{2I(W; U_J | \tilde{Z}^{(2)}, J)}{m}}. \quad (3)$$

Our next results in Theorem 3.2 let us compare Eq. (3) for different values of $m = |J|$.

Theorem 3.2. *Let $m_1 < m_2 \in [n]$, and let $J^{(m_1)}, J^{(m_2)}$ be random subsets of $[n]$, distributed uniformly among all subsets of size m_1 and m_2 , respectively, and independent from W , $\tilde{Z}^{(2)}$, and U . Then*

$$\frac{I(W; U_{J^{(m_1)}} | \tilde{Z}^{(2)}, J^{(m_1)})}{m_1} \leq \frac{I(W; U_{J^{(m_2)}} | \tilde{Z}^{(2)}, J^{(m_2)})}{m_2}. \quad (4)$$

Consequently, taking $m_2 = n$, for all $1 \leq m_1 \leq n$

$$\mathbb{E} \sqrt{\frac{2I^{\tilde{Z}^{(2)}}(W; U_{J^{(m_1)}} | J^{(m_1)})}{m_1}} \leq \sqrt{\frac{2I(W; U | \tilde{Z}^{(2)})}{n}}. \quad (5)$$

Corollary 3.3. $\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \sqrt{2I(W; U_J | \tilde{Z}^{(2)}, J)/m}$. The case $m = |J| = n$ is equivalent to Theorem 1.3. The bound is increasing in $m \in [n]$, and, the tightest bound is achieved when $m = |J| = 1$. Also, Eq. (5) shows our bound in Theorem 3.1 is tighter than Theorem 1.3 for $k = 2$.

To further tighten Theorem 3.2 when $m = 1$, we show that we can pull the expectation over both $\tilde{Z}^{(2)}$ and J outside the concave square-root function.

Theorem 3.4. *Let $J \sim \text{Unif}([n])$ (i.e., $m = 1$ above) be independent from W , $\tilde{Z}^{(2)}$, and U . Then*

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \mathbb{E} \sqrt{2I^{\tilde{Z}^{(2)}, J}(W; U_J)} = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \sqrt{2I^{\tilde{Z}^{(2)}}(W; U_i)}. \quad (6)$$

Remark 3.5. Theorem 3.4 is tighter than Theorem 1.3 since

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} \sqrt{2I^{\tilde{Z}^{(2)}}(W; U_i)} \leq \sqrt{\sum_{i=1}^n \frac{2}{n} I(W; U_i | \tilde{Z}^{(2)})} \leq \sqrt{\frac{2}{n} I(W; U | \tilde{Z}^{(2)})} \quad (7)$$

The first inequality is Jensen’s, while the second follows from the independence of indices U_i . \triangleleft

3.1 Controlling CMI bounds using KL Divergence

It is often difficult to compute MI directly. One standard approach in the literature is to bound MI by the expectation of the KL divergence of the conditional distribution of the parameters given the data (the “posterior”) with respect to a “prior”. The statement below is adapted from Negrea et al. [15].

Lemma 3.6. *Let X , Y , and Z be random elements. For all $\sigma(Z)$ -measurable random probability measures P on the space of Y ,*

$$I^Z(X; Y) \leq \mathbb{E}^Z[\text{KL}(\mathbb{P}^{X,Z}[Y] \parallel P)] \text{ a.s.,} \quad \text{with a.s. equality for } P = \mathbb{E}^Z[\mathbb{P}^{X,Z}[Y]] = \mathbb{P}^Z[Y].$$

We refer to the conditional law of W given S as the “posterior” of W given S , which we denote $Q = \mathbb{P}^S[W] = \mathbb{P}^{\tilde{Z}^{(2)}, U}[W]$, and to P as the prior. This can be used in combination with, for example, Lemma 3.6 and Theorem 1.3 to obtain that for any $\tilde{Z}^{(2)}$ -measurable random prior $P(\tilde{Z}^{(2)})$

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \sqrt{\frac{2 I(W; U | \tilde{Z}^{(2)})}{n}} \leq \sqrt{\frac{2 \mathbb{E}[\text{KL}(Q \parallel P(\tilde{Z}^{(2)}))]}{n}}. \quad (8)$$

Note that the prior only has access to $\tilde{Z}^{(2)}$, therefore from its perspective the training set can take 2^n different values. Alternatively, combining Lemma 3.6 and Theorem 3.1 yields

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \mathbb{E} \sqrt{\frac{2 \mathbb{E}^{\tilde{Z}^{(2)}} I^{\tilde{Z}^{(2)}}(W; U_J | U_{J^c}, J)}{m}} \leq \mathbb{E} \sqrt{\frac{2 \mathbb{E}^{\tilde{Z}^{(2)}} [\text{KL}(Q \parallel P(\tilde{Z}^{(2)}, U_{J^c}, J))]}{m}}. \quad (9)$$

In Eq. (9) the prior has access to $n - m$ samples in the training set, S_{J^c} , because $\tilde{Z}_{U_{J^c}}^{(2)} = S_{J^c}$. However, since $\tilde{Z}^{(2)}$ is known to the prior, the training set can take only 2^m distinct values from the point of view of the prior in Eq. (9). This is a significant reduction in the amount of information that can be carried by the indexes in U_J about the output hypothesis. Consequently, priors can be designed to better exploit the dependence of the output hypothesis and the index set.

3.2 Tighter Generalization bound for the case $m = 1$

Since the strategy above controls MI-based expressions via KL divergences, one may ask whether a bound derived with similar tools, but directly in terms of KL, can be tighter than the combination Lemma 3.6 and Theorem 3.1. The following result shows that for $m = 1$ a tighter bound can be derived by pulling the expectation over both U_{J^c} and J outside the concave square-root function.

Theorem 3.7. *Let $J \sim \text{Unif}([n])$ be independent from W , U , and $\tilde{Z}^{(2)}$. Let $Q = \mathbb{P}^{\tilde{Z}^{(2)}, U}[W]$ and P be a $\sigma(\tilde{Z}^{(2)}, U_{J^c}, J)$ -measurable random probability measure. Then*

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \mathbb{E} \sqrt{2 \text{KL}(Q \parallel P)}. \quad (10)$$

Here, the KL divergence is between two $\sigma(\tilde{Z}^{(2)}, J, U)$ -measurable random measures, so is random.

4 Generalization bounds for noisy, iterative algorithms

We apply this new class of generalization bounds to non-convex learning. We analyze the Langevin dynamics (LD) algorithm [8], following the analysis pioneered by Pensia et al. [16]. The example we set here is a blueprint for building bounds for other iterative algorithms. Our approach is similar to the recent advances by Li et al. [14] and Negrea et al. [15], employing data-dependent estimates to obtain easily simulated bounds. We find our new results allow us to exploit past iterates to obtain tighter bounds. The influence of past iterates is seen to take the form of a hypothesis test.

4.1 Bounding Generalization Error via Hypothesis Testing

The chain rule for KL divergence is a key ingredient of information-theoretic generalization error bounds for iterative algorithms [6, 14, 15, 16]. $\mathcal{W}^{\{0, \dots, T\}}$ denotes the space of parameters generated by an iterative algorithm in T iterations. For any measure, ν , on $\mathcal{W}^{\{0, \dots, T\}}$, and $W \sim \nu$, let ν_0 denote the marginal law of W_0 , and $\nu_{t|}$ denote the conditional law of W_t given $W_0 \dots W_{t-1}$.

Lemma 4.1 (Chain Rule for KL). *Let Q, P be probability measures on $\mathcal{W}^{\{0, \dots, T\}}$ with $Q_0 = P_0$. The following lemma bounds the KL divergence involving the posterior for the terminal parameter with one involving the sum of the KL divergences over each individual step of the trajectory. Then*

$$\text{KL}(Q_T \parallel P_T) \leq \text{KL}(Q \parallel P) = \sum_{t=1}^T Q_{0:(t-1)}[\text{KL}(Q_t \parallel P_{t|})]$$

The benefit of using the chain rule to analyze the iterative algorithm are two-fold: first, we gain analytical tractability; many bounds that appear in the literature implicitly require this form of incrementation [6, 14, 15, 16]. Second, and novel to the present work, the *information in the optimization trajectory can be exploited* to identify U from the history of W .

In order to understand how the prior may take advantage of information from the optimization trajectory, consider applying Lemma 4.1 to the KL term in Eq. (9). We have

$$\text{KL}(Q_T \| P_T(\tilde{Z}^{(2)}, U_{J^c}, J)) \leq \sum_{t=1}^T \mathbb{E}^{\tilde{Z}^{(2)}, U_{J^c}, J} [\text{KL}(Q_{t|} \| P_{t|}(\tilde{Z}^{(2)}, U_{J^c}, J))].$$

Here $P_{t|}(\tilde{Z}^{(2)}, U_{J^c}, J)$ is a $\sigma(\tilde{Z}^{(2)}, U_{J^c}, J, W_{0:t-1})$ -measurable random probability measure. The prior may use U_{J^c} , $\tilde{Z}^{(2)}$, and J to reduce the number of possible values that U can take to $2^{|J|}$. Moreover, since U_J is constant during optimization, $W_0, W_1, W_2, \dots, W_{t-1}$ may leak some information about U_J , and the prior can use this information to tighten the bound by choosing a $P_{t|}$ that achieves small $\text{KL}(Q_{t|} \| P_{t|})$. In the special case where the prior can perfectly estimate U_J from $W_0, W_1, W_2, \dots, W_{t-1}$, we can set $P_{t|} = Q_{t|}$ and $\text{KL}(Q_{t|} \| P_{t|})$ will be zero. As will be seen in the next subsection, we can explicitly design a prior that uses the information in the optimization trajectory for the LD algorithm.

The process by which the prior can learn from the trajectory can be viewed as an *online hypothesis test*, or binary decision problem, where the prior at time t allocates belief between 2^m possible explanations, given by the possible values of U_J , based on the evidence provided by W_0, \dots, W_t . If the prior is able to identify U_J based on the W s then the bound stops accumulating, even if the gradients taken by subsequent training steps are large. This means that penalties for information obtained later in training are *discounted* based on the information obtained earlier in training.

4.2 Example: Langevin Dynamics Algorithm for Non-Convex Learning

We apply these results to obtain generalization bounds for a gradient-based iterative noisy algorithm, the Langevin Dynamics (LD) algorithm. For classification with continuous parameters, the 0-1 loss does not provide useful gradients. Typically we optimize a surrogate objective, based on a *surrogate* loss, such as cross entropy. Write $\tilde{\ell}: \mathcal{Z} \times \mathcal{W} \rightarrow \mathbb{R}$ for the surrogate loss and let $\tilde{R}_S(w) = \frac{1}{n} \sum_{i=1}^n \tilde{\ell}(Z_i, w)$ be the empirical surrogate risk. Let η_t be the learning rate at time t , β_t the inverse temperature at time t and let ε_t be sampled i.i.d. from $\mathcal{N}(0, \mathbb{I}_d)$. Then the LD algorithm iterates are given by

$$W_{t+1} = W_t - \eta_t \nabla \tilde{R}_S(W_t) + \sqrt{\frac{2\eta_t}{\beta_t}} \varepsilon_t. \quad (11)$$

The prior We will take $m = 1$, and construct a bespoke $\sigma(\tilde{Z}^{(2)}, U_{J^c}, J)$ -measurable prior for this problem in order to apply Theorem 3.7. The prior is based on a *decision function*, $\theta: \mathbb{R} \rightarrow [0, 1]$, which at each time $t + 1$ takes in a $\sigma(W_0 \dots W_t)$ -measurable *test statistic*, ΔY_t , and returns a *degree of belief* in favor of the hypothesis $U_J = 1$ over $U_J = 2$. The prior predicts an LD step by replacing the unknown (to the prior) contribution to the gradient of the data point at index J with a $\hat{\theta}_t = \theta(\Delta Y_t)$ -weighted average of the gradients due to each candidate $\{Z_{1,J}, Z_{2,J}\}$. The conditional law of the t th iterate under the prior is a $\sigma(\tilde{Z}^{(2)}, U_{J^c}, J, W_0, \dots, W_t)$ -measurable random measure, as required. The exact value of the test statistic is $\Delta Y_t = Y_{t,2} - Y_{t,1}$, here the $Y_{0,1} = Y_{0,2} = 0$ and $Y_{t,u}$ are defined by the formula in Eq. (13). The conditional law of the t th iterate under the prior is described by

$$W_{t+1} = W_t - \frac{\eta_t}{n} \left(\sum_{i=1}^n \nabla \tilde{\ell}(Z_i, W_t) + \hat{\theta}_t \nabla \tilde{\ell}(Z_{1,J}, W_t) + (1 - \hat{\theta}_t) \nabla \tilde{\ell}(Z_{2,J}, W_t) \right) + \sqrt{\frac{2\eta_t}{\beta_t}} \varepsilon_t. \quad (12)$$

The test statistic chosen is based on the log-likelihood-ratio test statistic for the independent mean 0 Gaussian random vectors $(\varepsilon_s)_{s=1}^t$, which is well known to be *uniformly most powerful* for the binary discrimination of means. Natural choices for θ are symmetric CDFs, since they treat possible values of U symmetrically, and are monotone in the test statistic.

We define the *two-sample incoherence* at time t by $\zeta_t = \nabla \tilde{\ell}(Z_{1,J}, W_t) - \nabla \tilde{\ell}(Z_{2,J}, W_t)$. Θ denotes the set of measurable $\theta: \mathbb{R} \rightarrow [0, 1]$. $Y_{0,1} = Y_{0,2} = 0$, and for $t \geq 1$, $Y_{t,1}$ and $Y_{t,2}$ are given by (for $u \in \{1, 2\}$)

$$Y_{t,u} \triangleq \sum_{i=1}^t \frac{\beta_{i-1}}{4\eta_{i-1}} \|W_i - W_{i-1} + \eta_{i-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{i-1}) + \frac{\eta_{i-1}}{n} \nabla \tilde{\ell}(Z_{u,J}, W_{i-1})\|^2. \quad (13)$$

Theorem 4.2 (Generalization bound for LD algorithm). *Let $\{W_t\}_{t \in [T]}$ denote the iterates of the LD algorithm. If $\ell(Z, w)$ is $[0, 1]$ -bounded then*

$$\mathbb{E} [R_D(W_T) - \hat{R}_S(W_T)] \leq \frac{1}{n\sqrt{2}} \inf_{\theta \in \Theta} \mathbb{E} \sqrt{\sum_{t=0}^{T-1} \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \beta_t \eta_t \|\zeta_t\|^2 \left(\mathbb{1}\{U_J = 1\} - \theta(Y_{t,2} - Y_{t,1}) \right)^2}. \quad (14)$$

Remark 4.3. For $\theta \in \Theta$ with $1 - \theta(x) = \theta(-x)$, Eq. (14) simplifies to

$$\mathbb{E} [R_D(W_T) - \hat{R}_S(W_T)] \leq \frac{1}{n\sqrt{2}} \mathbb{E} \sqrt{\sum_{t=0}^{T-1} \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \beta_t \eta_t \|\zeta_t\|^2 \theta^2(-1^{U_J} (Y_{t,2} - Y_{t,1}))}. \quad (15)$$

For instance $\theta(x) = \frac{1}{2} + \frac{1}{2} \tanh(x)$ and $\theta(x) = \frac{1}{2} + \frac{1}{2} \text{sign}(x)$ satisfy $1 - \theta(x) = \theta(-x)$. \triangleleft

Remark 4.4. By the law of total expectation, for any $\theta \in \Theta$, $\text{EGE}_D(\mathcal{A}) \leq \frac{1}{2\sqrt{2}n} \mathbb{E} [V_1 + V_2]$, where

$$V_u \triangleq \sqrt{\sum_{t=0}^{T-1} \mathbb{E}^{\tilde{Z}^{(2)}, U_{J^c}, J, U_J=u} \beta_t \eta_t \|\zeta_t\|^2 (\mathbb{1}\{u = 1\} - \theta(Y_{t,2} - Y_{t,1}))^2}, \quad u \in \{1, 2\}. \quad (16)$$

To estimate V_u ($u \in \{1, 2\}$) for fixed J , the training set is $S_u = \{Z_1, \dots, Z_{J-1}, \tilde{Z}_{u,J}, Z_{J+1}, \dots, Z_n\}$. Hence V_1, V_2 can be simulated from just $n+1$ data points $(Z_1, \dots, Z_{J-1}, Z_{J+1}, \dots, Z_n, \tilde{Z}_{1,J}, \tilde{Z}_{2,J}) \sim \mathcal{D}^{\otimes(n+1)}$. \triangleleft

The generalization bound in Eq. (14) does not place any restrictions on the learning rate or Lipschitz continuity of the loss or its gradient. In the next corollary we study the asymptotic properties of the bound in Eq. (14) when $\tilde{\ell}$ is L -Lipschitz. Then, we draw a comparison between the bound in this paper and some of the existing bounds in the literature.

Corollary 4.5. *Under the assumption that $\tilde{\ell}$ is L -Lipschitz, we have $\|\zeta_t\| \leq 2L$. Then, the generalization bound in Eq. (14) can be upper-bounded as*

$$\mathbb{E}(R_D(W_T) - R_S(W_T)) \leq \frac{\sqrt{2}L}{n} \inf_{\theta \in \Theta} \mathbb{E} \sqrt{\sum_{t=0}^{T-1} \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \beta_t \eta_t (\mathbb{1}\{U_J = 1\} - \theta(Y_{t,2} - Y_{t,1}))^2}. \quad (17)$$

Remark 4.6. Under an L -Lipschitz assumption, for the LD algorithm, Li et al. [14, Thm. 9] have

$$\mathbb{E} [R_D(W_T) - R_S(W_T)] \leq \frac{\sqrt{2}L}{n} \sqrt{\sum_{t=0}^{T-1} \beta_t \eta_t}. \quad (18)$$

We immediately see that Eq. (17) provides a constant factor improvement over Eq. (18) by naïvely using $\theta : x \mapsto 1/2$. Our bound has order-wise improvement with respect to n over that of Bu et al. [6] and Pensia et al. [16] under the L -Lipschitz assumption. Negrea et al. [15, App. E.1] obtain

$$\mathbb{E} [R_D(W_T) - R_S(W_T)] \leq \frac{L}{2(n-1)} \sqrt{\sum_{t=0}^{T-1} \beta_t \eta_t}. \quad (19)$$

which is a constant factor better than our bound for the choice $\theta : x \mapsto 1/2$. However, this θ essentially corresponds to no hypothesis test, yielding the same prior as in [15]. For more sophisticated choices of decision function (θ), even under a Lipschitz-surrogate loss assumption, it is difficult to compare our bound with related work because the exact impact of θ -discounting is difficult to quantify analytically. \triangleleft

Remark 4.7. A prevailing method for analyzing the generalization error in [6, 14, 15, 16] for iterative algorithms is via the chain rule for KL, using priors for the joint distribution of weight vectors that are Markov, i.e., given the t th weight, the $(t+1)$ th weight is conditionally independent from the trajectory so far. Existing results using this approach accumulate a "penalty" for each step. In [6, 14, 15], the penalty terms are, respectively, the squared Lipschitz constant, the squared norm of the gradients, and the trace of the minibatch gradient covariance. The penalty term in our paper is the squared norm of "two-sample incoherence", defined in Theorem 4.2 as the squared norm of the difference between the gradient of a randomly selected training point and the held-out point. However, the use of chain rule along with existing "Markovian" priors introduces a source of looseness, i.e., the accumulating penalty may diverge to $+\infty$ yielding vacuous bounds (as seen in Fig. 1). *The distinguishing feature of our data-dependent CMI analysis is that the penalty terms get "filtered" by the online hypothesis test via our non-Markovian prior, i.e., our prediction for $t+1$ depends on whole trajectory.* When the true index can be inferred from the previous weights, then the penalty essentially stops accumulating. \triangleleft

4.2.1 Empirical Results

In order to better understand the effect of discounting and the degree of improvement due to our new bounds and more sophisticated prior, we turn to simulation studies. We present and compare the empirical evaluations of the generalization bound in Theorem 4.2 with the data-dependent generalization bounds in Li et al. [14] and Negrea et al. [15]. For brevity, many of the details behind our implementation are deferred to Appendix G. The functional form of our bounds and [14, 15] are nearly identical as all of them use the chain rule for KL divergence. Nevertheless, the summands appearing in the bounds are different. The bound in [14] depends on the squared surrogate loss gradients norm, and the generalization bound in Negrea et al. [15] depends on the squared norm of *training set incoherence* defined as $\|\nabla \tilde{\ell}(Z_J, W_t) - \frac{1}{n-1} \sum_{i \in [n], i \neq J} \nabla \tilde{\ell}(Z_i, W_t)\|^2$ where the training set is $\{Z_1, \dots, Z_n\}$ and $J \sim \text{Unif}([n])$. The first key difference between our bound and others is that the summand in our bound consists of two terms: squared norm of the two-sample incoherence, i.e., $\|\zeta_t\|^2$, and the squared error probability of a hypothesis test at time t , given by the term $(\mathbb{1}\{U_J = 1\} - \theta(\sum_{i=0}^t (Y_{i,2} - Y_{i,1})))^2$ in our bound. A consequence of this, and the second fundamental difference between our bound and existing bounds, is that our bound exhibits a trade-off in $\|\zeta_t\|^2$ because large $\|\zeta_t\|^2$ will make the error of the hypothesis test small on future iterations, whereas the bounds in [14, 15] are uniformly increasing with respect to the squared norm of surrogate loss gradients and the training set incoherence, respectively. In this section we empirically evaluate and compare our bound with related work across various neural network architectures and datasets.

Using Monte Carlo (MC) simulation, we compared estimates of our expected generalization error bounds with estimates of the bound from [14, 15] for the MNIST [13], CIFAR10 [12], and Fashion-MNIST [23] datasets in Fig. 1 and Table 1. For all the plots we consider $\theta(x) = \frac{1}{2}(1 + \text{erf}(x))$ for our bound. Also, in the last row of Table 1, we report the *unbiased estimate* of our bound optimized over the choice of θ function. We plot the squared norm of the two sample incoherence and training set incoherence, as well as the squared error probability of the hypothesis test. Fig. 1 and Table 1 show that our bound is tighter, and remain non-vacuous after many more iterations. We also observe that the variances for MC estimates of our bound are smaller than those of Negrea et al. [15], and it is also smaller than Li et al. [14] for CIFAR10 and MNIST-CNN experiments. Moreover, we observe that the error probability of the hypothesis test decays with the number of iterations, which matches the intuition that, as one observes more noisy increments of the process, one is more able to determine which point is contributing to the gradient. For CIFAR10, $\|\zeta_t\|^2$ is large because the generalization gap is large. However, as mentioned in the beginning of this section, large $\|\zeta_t\|^2$ makes the hypothesis testing easier on subsequent iterations. For instance, after iteration 600 the error is vanishingly small for CIFAR10 experiments which results in a plateau region in the bound. We can also observe the same phenomenon for the Fashion-MNIST experiment. This property distinguishes our bound from those in [14, 15].

Results for MNIST with CNN demonstrate that $\|\zeta_t\|^2$ and training set incoherence are close to each other. The reason behind this observations is that the generalization gap is small. Also, for this experiment the performance of the hypothesis testing is only slightly better than random guessing since the generalization gap is small, and it is difficult to distinguish the training samples from the test samples. This observation explains why our generalization bound is close to that of [15]. Nevertheless, the hypothesis testing performance improves with more training iterations, leading the two bounds to diverge, with our new bound performing better at later iterations. Finally, the scaling of our bound with respect to the number of iteration is tighter than in the bounds in [14, 15] as can be seen in Fig. 1.

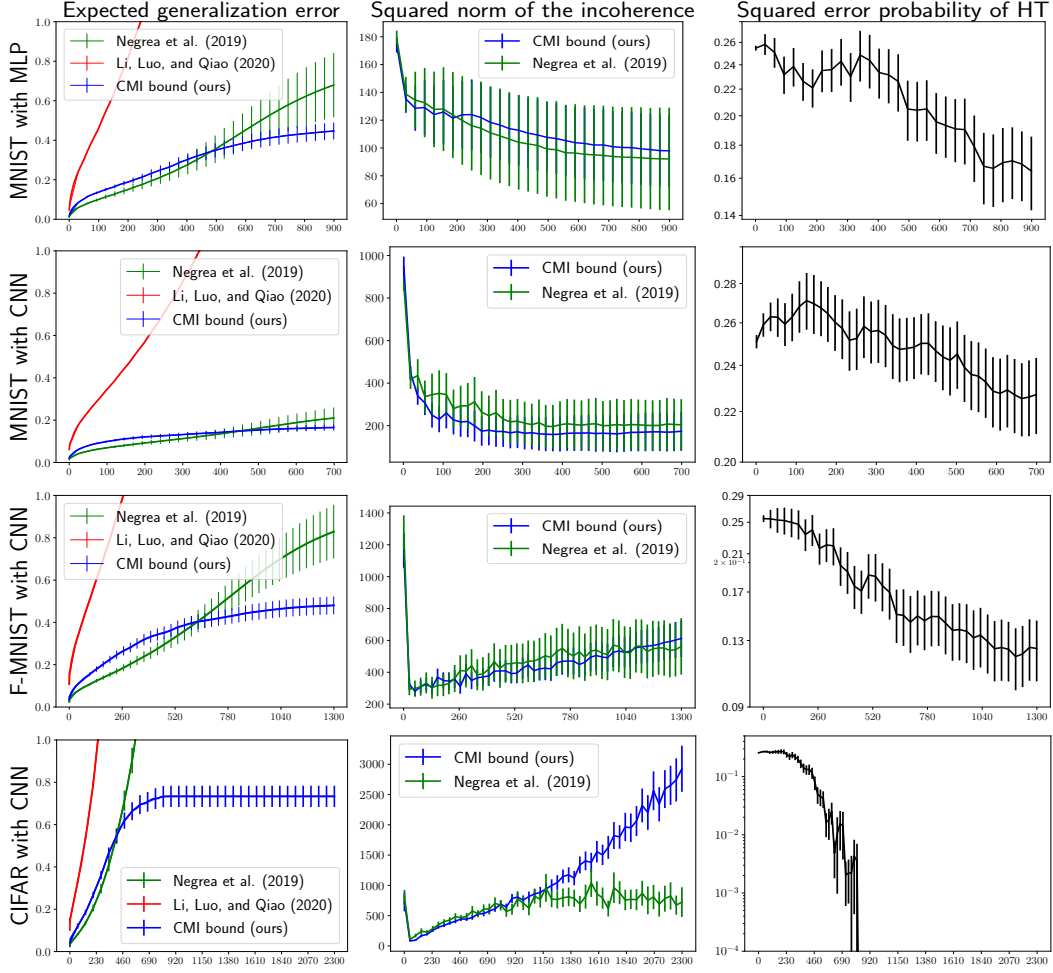


Figure 1: Numerical results for various datasets and architectures. All the x-axes represent the training iteration. The plots in the first column depict a Monte Carlo estimate of our bounds with that of Li et al. [14] and Negrea et al. [15]. The plots in the second column compare the mean of the *training set incoherence* in [15] with the two-sample incoherence in our bound. Finally, the plots in the third column show the mean of the squared error probability of the hypothesis testing performed by the proposed prior in our bound.

	MNIST-MLP	MNIST-CNN	CIFAR10-CNN	FMNIST-CNN
Training error	$4.33 \pm 0.01\%$	$2.59 \pm 0.01\%$	$9.39 \pm 0.36\%$	$7.96 \pm 0.03\%$
Generalization error	$0.88 \pm 0.01\%$	$0.55 \pm 0.01\%$	$32.89 \pm 0.44\%$	$3.71 \pm 0.03\%$
Negrea et al. [15]	$67.93 \pm 16.25\%$	$20.98 \pm 5.01\%$	$4112.63 \pm 567.08\%$	$82.89 \pm 12.64\%$
Li et al. [14]	$600.29 \pm 1.99\%$	$245.03 \pm 2.37\%$	$20754.32 \pm 75.95\%$	$598.62 \pm 3.21\%$
CMI (Ours)	$44.65 \pm 4.27\%$	$16.51 \pm 1.41\%$	$71.76 \pm 4.82\%$	$48.01 \pm 4.22\%$
CMI-OPT(Ours)	$39.06 \pm 5.52\%$	$13.24 \pm 1.53\%$	$63.00 \pm 5.97\%$	$41.17 \pm 5.85\%$

Table 1: Summary of the results. The generalization bounds are reported at the end of training.

Acknowledgments

The authors would like to thank Blair Bilodeau and Yasaman Mahdaviyeh for feedback on drafts of this work, and Shiva Ketabi for helpful discussions on the implementation of the bounds.

Funding

MH is supported by the Vector Institute. JN is supported by an NSERC Vanier Canada Graduate Scholarship, and by the Vector Institute. DMR is supported by an NSERC Discovery Grant and an Ontario Early Researcher Award. This research was carried out in part while MH, JN, GKD, and DMR were visiting the Institute for Advanced Study. JN’s visit to the Institute was funded by an NSERC Michael Smith Foreign Study Supplement. Resources used in preparing this research were provided, in part, by the Province of Ontario, the Government of Canada through CIFAR, and companies sponsoring the Vector Institute www.vectorinstitute.ai/partners.

Broader Impact This work builds upon the community’s understanding of generalization error for machine learning methods. This has a positive impact on the scientific advancement of the field, and may lead to further improvements in our understanding, methodologies and applications of machine learning and AI. While there are not obvious direct societal implications of the present work, the indirect and longer term impact on society may be positive, negative or both depending on how, where and for what machine learning method that will have benefited from our research are used in the future.

References

- [1] A. Lopez and V. Jog. “Generalization error bounds using Wasserstein distances”. In: *IEEE Information Theory Workshop*. 2018.
- [2] A. R. Asadi and E. Abbe. *Chaining Meets Chain Rule: Multilevel Entropic Regularization and Training of Neural Nets*. 2019. arXiv: 1906.11148.
- [3] A. Asadi, E. Abbe, and S. Verdú. “Chaining mutual information and tightening generalization bounds”. In: *Advances in Neural Information Processing Systems*. 2018, pp. 7234–7243.
- [4] R. Bassily, S. Moran, I. Nachum, J. Shafer, and A. Yehudayoff. “Learners that Use Little Information”. In: *Algorithmic Learning Theory*. 2018, pp. 25–55.
- [5] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- [6] Y. Bu, S. Zou, and V. V. Veeravalli. “Tightening mutual information based bounds on generalization error”. In: *2019 IEEE International Symposium on Information Theory (ISIT)*. IEEE. 2019, pp. 587–591.
- [7] R. Durrett. *Probability: theory and examples*. Vol. 49. Cambridge university press, 2019.
- [8] S. B. Gelfand and S. K. Mitter. “Recursive stochastic algorithms for global optimization in \mathbb{R}^d ”. *SIAM Journal on Control and Optimization* 29.5 (1991), pp. 999–1018.
- [9] I. Issa, A. R. Esposito, and M. Gastpar. “Strengthened Information-theoretic Bounds on the Generalization Error”. In: *2019 IEEE International Symposium on Information Theory (ISIT)*. 2019, pp. 582–586.
- [10] J. Jiao, Y. Han, and T. Weissman. “Dependence measures bounding the exploration bias for general measurements”. In: *IEEE International Symposium on Information Theory*. 2017.
- [11] O. Kallenberg. *Foundations of modern probability*. Springer Science & Business Media, 2006.
- [12] A. Krizhevsky. *Learning Multiple Layers of Features from Tiny Images*. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>. 2009.
- [13] Y. LeCun, C. Cortes, and C. J. C. Burges. *MNIST handwritten digit database*. <http://yann.lecun.com/exdb/mnist/>. 2010.
- [14] J. Li, X. Luo, and M. Qiao. “On Generalization Error Bounds of Noisy Gradient Methods for Non-Convex Learning”. In: *International Conference on Learning Representations*. 2020.
- [15] J. Negrea, M. Haghifam, G. K. Dziugaite, A. Khisti, and D. M. Roy. “Information-Theoretic Generalization Bounds for SGLD via Data-Dependent Estimates”. In: *Advances in Neural Information Processing Systems*. 2019, pp. 11013–11023.

- [16] A. Pensia, V. Jog, and P.-L. Loh. “Generalization error bounds for noisy, iterative algorithms”. In: *2018 IEEE International Symposium on Information Theory (ISIT)*. 2018, pp. 546–550.
- [17] M. Raginsky, A. Rakhlin, M. Tsao, Y. Wu, and A. Xu. “Information-theoretic analysis of stability and bias of learning algorithms”. In: *2016 IEEE Information Theory Workshop (ITW)*. IEEE. 2016, pp. 26–30.
- [18] D. Russo and J. Zou. “Controlling Bias in Adaptive Data Analysis Using Information Theory”. In: *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*. Ed. by A. Gretton and C. C. Robert. Vol. 51. Proceedings of Machine Learning Research. Cadiz, Spain: PMLR, May 2016, pp. 1232–1240.
- [19] D. Russo and J. Zou. *How much does your data exploration overfit? Controlling bias via information usage*. 2015. arXiv: 1511.05219.
- [20] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. “Membership inference attacks against machine learning models”. In: *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2017, pp. 3–18.
- [21] T. Steinke and L. Zakynthinou. “Reasoning About Generalization via Conditional Mutual Information” (2020). arXiv: 2001.09122.
- [22] H. Te Sun. “Nonnegative entropy measures of multivariate symmetric correlations”. *Information and Control* 36 (1978), pp. 133–156.
- [23] H. Xiao, K. Rasul, and R. Vollgraf. *Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms*. Aug. 28, 2017. arXiv: cs.LG/1708.07747 [cs.LG].
- [24] A. Xu and M. Raginsky. “Information-theoretic analysis of generalization capability of learning algorithms”. In: *Advances in Neural Information Processing Systems*. 2017, pp. 2524–2533.

A CMI, Membership Attack, and Fano's Inequality

Let $\tilde{Z}^{(k)}$, $U^{(k)}$, and S as in Definition 1.2. Consider the following hypothesis testing problem. Assume a decision maker observes W and wishes to recover $U^{(k)}$ by having access to the super-sample $\tilde{Z}^{(k)}$. For any estimate $\hat{U} = \Psi(W, \tilde{Z}^{(k)})$, we have the Markov chain

$$U^{(k)} \rightarrow S \rightarrow W \rightarrow \widehat{U^{(k)}}.$$

and so, combined with the fact that $U^{(k)}$ is uniformly distributed over a set of size k^n , we can invoke Fano's inequality to bound the error probability of the decision maker. In particular,

$$\inf_{\Psi} \mathbb{P} \left[\Psi(W, \tilde{Z}^{(k)}) \neq U^{(k)} \right] \geq 1 - \frac{I(W; U^{(k)} | \tilde{Z}^{(k)}) + \log 2}{n \log k}.$$

Hence, $I(W; U^{(k)} | \tilde{Z}^{(k)})$ provides a lower bound on the hardness of the hypothesis testing problem, where one wants to identify the training sample given access to $\tilde{Z}^{(k)}$ and W .

Some interpretation of our result is helpful. Consider an adversary who has access to the supersample $\tilde{Z}^{(k)}$ and wishes to identify the training set that was used for the training after observing the output of a learning algorithm W . Our result here showed that the CMI upperbounds the success probability of *every* adversary. Also, recall that the CMI upper bounds the expected generalization error. In the literature of data privacy in machine learning, this problem is known as Membership Attack [20], and it is empirically observed that a machine learning model leaks information about its training set when the generalization error is large [20]. Our result in this section provides a formal connection between generalization and this specific membership attack problem.

B Matching the leading coefficient of Theorem 1.1 with $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$

Theorem B.1. *Let $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$ as defined in the introduction. Then, for $k > 2$*

$$\text{EGE}_{\mathcal{D}}(\mathcal{A}) \leq \frac{\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{\lambda^*} + \frac{\exp(\frac{\lambda^*}{n}) - \frac{\lambda^*}{n} - 1}{\frac{\lambda^*}{n}} \frac{k^3 + 7k^2 - 8k - 16}{4(k^3 - 2k^2)},$$

where

$$\lambda^* = n\mathfrak{W} \left(\left(\frac{4(k^3 - 2k^2)\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{n(k^3 + 7k^2 - 8k - 16)} - 1 \right) \exp(-1) \right) + n,$$

and \mathfrak{W} is the Lambert W function.

The proof is deferred to Appendix C. Here, we quantitatively compare Theorem 1.1, Theorem 1.3, and Theorem B.1, we consider the case that the output of \mathcal{A} takes value in a finite set and $k \rightarrow \infty$. In this case, Theorem B.1 can be rephrased as

$$\begin{aligned} \text{EGE}_{\mathcal{D}}(\mathcal{A}) &\leq \lim_{k \rightarrow \infty} \frac{\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{\lambda^*} + \frac{\exp(\frac{\lambda^*}{n}) - \frac{\lambda^*}{n} - 1}{\frac{\lambda^*}{n}} \frac{k^3 + 7k^2 - 8k - 16}{4(k^3 - 2k^2)} \\ &= \frac{\text{IOMI}_{\mathcal{D}}(\mathcal{A})}{\lambda_{\infty}^*} + \frac{\exp(\frac{\lambda_{\infty}^*}{n}) - \frac{\lambda_{\infty}^*}{n} - 1}{4 \frac{\lambda_{\infty}^*}{n}} \end{aligned} \quad (20)$$

where $\lambda_{\infty}^* = n\mathfrak{W} \left(\left(\frac{4\text{IOMI}_{\mathcal{D}}(\mathcal{A})}{n} - 1 \right) \exp(-1) \right) + n$. In the next plot, we compare the values of the bounds in Theorem 1.1, Theorem 1.3, and Theorem B.1 assuming $\text{IOMI}_{\mathcal{D}}(\mathcal{A}) = 1$. As seen, the bound in Theorem B.1 provides much tighter constant compared with Theorem 1.3, and it matches with Theorem 1.1.




Figure 2: Comparison between constants of Theorem 1.1, Theorem 1.3, and Theorem B.1 for the case $k \rightarrow \infty$.

C Proofs of Section 2

Proof of Theorem 2.1. By the chain rule for the mutual information, we have

$$I(W; U^{(k)}, \tilde{Z}^{(k)}) = I(W; \tilde{Z}^{(k)}) + I(W; U^{(k)} | \tilde{Z}^{(k)}). \quad (21)$$

Since S is $\sigma(\tilde{Z}^{(k)}, U^{(k)})$ -measurable, $I(W; U^{(k)}, \tilde{Z}^{(k)}) = I(W; S, U^{(k)}, \tilde{Z}^{(k)})$. But then W is independent of $\tilde{Z}^{(k)}, U^{(k)}$ given S , hence $I(W; S, U^{(k)}, \tilde{Z}^{(k)}) = I(W; S)$. The result follows from the nonnegativity of mutual information. \square

Proof of Theorem 2.2. By Theorem 2.1, $I(W; U^{(k)} | \tilde{Z}^{(k)}) = I(W; S) - I(W; \tilde{Z}^{(k)})$. Therefore, in order to prove the claim, we need to show $\lim_{k \rightarrow \infty} I(W; \tilde{Z}^{(k)}) = 0$ when $I(W; S)$ is finite.

Recall that \mathcal{A} is a probability kernel from the space of tuples in \mathcal{Z} to \mathcal{W} . Assume $\mathcal{W} = \{w_1, \dots, w_m\}$. For each $l \in [m]$, let $\kappa_l(S) = \mathbb{P}^S[W = w_l]$ and $f_l : \mathcal{Z}^{kn} \rightarrow [0, 1]$ be a measurable function defined as

$$f_l(\tilde{Z}^{(k)}) = \frac{1}{k^n} \sum_{u \in [k]^n} \kappa_l(\tilde{Z}_u^{(k)}).$$

Letting $z, z' \in \mathcal{Z}^{kn}$ be two super-samples that only differ in one element, it is straightforward to see that

$$|f_l(z) - f_l(z')| \leq \frac{1}{k}.$$

Therefore, we can invoke McDiarmid's inequality to obtain

$$\mathbb{P}[|f_l(\tilde{Z}^{(k)}) - \mathbb{E}[f_l(\tilde{Z}^{(k)})]| \geq \varepsilon] \leq \exp\left(-\frac{2k\varepsilon^2}{n}\right). \quad (22)$$

Also, we have $\mathbb{E}[f_l(\tilde{Z}^{(k)})] = \mathbb{P}[W = w_l]$ as each element of $\tilde{Z}^{(k)}$ is IID. Hence, $f_l(\tilde{Z}^{(k)}) \rightarrow \mathbb{P}[W = w_l]$ in probability as k diverges.

By the definition of mutual information and KL divergence,

$$\begin{aligned} I(W; \tilde{Z}^{(k)}) &= \mathbb{E}[\text{KL}(\mathbb{P}^{\tilde{Z}^{(k)}}[W] \parallel \mathbb{P}[W])] \\ &= \mathbb{E}\left[\text{KL}\left(\frac{1}{k^n} \sum_{u \in [k]^n} \mathbb{P}^{\tilde{Z}_u^{(k)}}[W] \parallel \mathbb{P}[W]\right)\right] \\ &= \mathbb{E}\left[\sum_{l=1}^m \frac{1}{k^n} \sum_{u \in [k]^n} \kappa_l(\tilde{Z}_u^{(k)}) \log \frac{\frac{1}{k^n} \sum_{u \in [k]^n} \kappa_l(\tilde{Z}_u^{(k)})}{\mathbb{P}[W = w_l]}\right] \\ &= \sum_{l=1}^m \mathbb{E}\left[f_l(\tilde{Z}^{(k)}) \log \frac{f_l(\tilde{Z}^{(k)})}{\mathbb{P}[W = w_l]}\right]. \end{aligned} \quad (23)$$

Defining $\phi_l : [0, 1] \rightarrow \mathbb{R}$ as $\phi_l(x) = x \log \frac{x}{\mathbb{P}[W=w_l]}$, we have established

$$I(W; \tilde{Z}^{(k)}) = \sum_{l=1}^m \mathbb{E}[\phi_l(f_l(\tilde{Z}^{(k)}))]. \quad (24)$$

Note that ϕ_l is a continuous and bounded function. By a standard result [7, Thm. 2.3.4], $f_l(\tilde{Z}^{(k)}) \rightarrow \mathbb{P}[W = w_l]$ in probability implies that

$$\mathbb{E}[\phi_l(f_l(\tilde{Z}^{(k)}))] \rightarrow \mathbb{E}[\phi_l(\mathbb{P}[W = w_l])] = 0,$$

as k goes to infinity. Using this, we conclude that $I(W; \tilde{Z}^{(k)}) \rightarrow 0$ as k diverges, as was to be shown. \square

Proof of Theorem B.1. For any $k \in \mathbb{N}$ define

$$\rho_i^{(k)}(m) = \begin{cases} -1 & \text{if } m = i, \\ \frac{1}{k-1} & \text{otherwise} \end{cases}.$$

where m and $i \in [k]$. Consider random variables $\tilde{Z}^{(k)}$, U , and W as in the definition of $\text{CMI}_{\mathcal{D}}^k(\mathcal{A})$. Also, let $\hat{U} \stackrel{d}{=} U$ and $\hat{U} \perp\!\!\!\perp (\tilde{Z}^{(k)}, W, U)$. Let $f : \mathcal{Z}^{kn} \times [k]^n \times \mathcal{W} \rightarrow [-1, 1]$ be given by

$$f(\tilde{z}^{(k)}, u, w) = \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^k \rho_i^{(k)}(u_j) \ell(w, z_{i,j}).$$

Then, by the Donsker–Varadhan variational formula [5, Prop. 4.15] of the KL divergence we obtain

$$\begin{aligned} \text{CMI}_{\mathcal{D}}^k(\mathcal{A}) &= \mathbb{E}[\text{KL}(\mathbb{P}^{\tilde{Z}^{(k)}, W}[U] \parallel \mathbb{P}[U])] \\ &\geq \lambda \mathbb{E}[f(\tilde{Z}^{(k)}, U, W)] - \mathbb{E} \log \mathbb{E}^{\tilde{Z}^{(k)}, W}[\exp(\lambda f(\tilde{Z}^{(k)}, \hat{U}, W))] \end{aligned} \quad (25)$$

It is straightforward to see that the first term in the RHS of Eq. (25) is $\mathbb{E}[f(\tilde{Z}^{(k)}, U, W)] = \text{EGE}_{\mathcal{D}}(\mathcal{A})$. In what follows we analyze the second term in the RHS of Eq. (25). We begin with

$$\begin{aligned} \mathbb{E} \log \mathbb{E}^{\tilde{Z}^{(k)}, W}[\exp(\lambda f(\tilde{Z}^{(k)}, \hat{U}, W))] &= \mathbb{E} \log \mathbb{E}^{\tilde{Z}^{(k)}, W} \left[\prod_{j=1}^n \exp\left(\frac{\lambda}{n} \sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j) \ell(W, Z_{i,j})\right) \right] \\ &= \mathbb{E} \log \prod_{j=1}^n \mathbb{E}^{\tilde{Z}^{(k)}, W} \exp\left(\frac{\lambda}{n} \sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j) \ell(W, Z_{i,j})\right) \\ &\leq \mathbb{E} \sum_{j=1}^n \left(\exp\left(\frac{\lambda}{n}\right) - \frac{\lambda}{n} - 1 \right) \mathbb{E}^{\tilde{Z}^{(k)}, W} \left[\sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j) \ell(W, Z_{i,j}) \right]^2 \end{aligned} \quad (26)$$

The first step follows from the independence of the elements in \hat{U} . The last inequality is obtained by the Bennet’s inequality [5, Thm. 2.9] on the moment generating function and the fact that the elements of \hat{U} are independent of $(\tilde{Z}^{(k)}, W)$. Also, we have used $\mathbb{E}^{\tilde{Z}^{(k)}, W} \sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j) \ell(W, Z_{i,j}) = 0$

since $\mathbb{E}\rho_i^{(k)}(\hat{U}_j) = 0$ and $|\sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j)\ell(W, Z_{i,j})| \leq 1$ a.s.. For a fixed j , from Eq. (27) we obtain

$$\mathbb{E}[\sum_{i=1}^k \rho_i^{(k)}(\hat{U}_j)\ell(W, Z_{i,j})]^2 = \frac{1}{k} \mathbb{E} \sum_{\tilde{i}=[k]} \left[\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) - \ell(W, Z_{\tilde{i},j}) \right]^2 \quad (28)$$

$$= \frac{1}{k^2} \mathbb{E} \left[\sum_{u_j \in [k]} \mathbb{E}^{\tilde{Z}^{(k)}, U_j=u_j} \sum_{\tilde{i}=[k]} \left[\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) - \ell(W, Z_{\tilde{i},j}) \right]^2 \right] \quad (29)$$

$$= \frac{1}{k^2} \mathbb{E} \left[\sum_{u_j \in [k]} \mathbb{E}^{\tilde{Z}^{(k)}, U_j=u_j} \left[\sum_{\tilde{i}=[k], \tilde{i} \neq u_j} \left[\frac{1}{k-1} \ell(W, Z_{u_j,j}) + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) - \ell(W, Z_{\tilde{i},j}) \right]^2 \right. \right. \\ \left. \left. + \left[\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) - \ell(W, Z_{u_j,j}) \right]^2 \right] \right] \quad (30)$$

$$\leq \frac{1}{k^2} \mathbb{E} \left[\sum_{u_j \in [k]} \mathbb{E}^{\tilde{Z}^{(k)}, U_j=u_j} \left[\sum_{\tilde{i}=[k], \tilde{i} \neq u_j} \left[\left(\frac{1}{k-1} + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right)^2 + \ell(W, Z_{\tilde{i},j})^2 \right. \right. \right. \\ \left. \left. - 2\ell(W, Z_{\tilde{i},j}) \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right] + \left[\left(\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) \right)^2 + 1 \right] \right] \right] \quad (31)$$

$$= \frac{1}{k^2} \sum_{u_j \in [k]} \mathbb{E} \left[\sum_{\tilde{i}=[k], \tilde{i} \neq u_j} \mathbb{E}^W \left[\left(\frac{1}{k-1} + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right)^2 + \ell(W, Z_{\tilde{i},j})^2 \right. \right. \\ \left. \left. - 2\ell(W, Z_{\tilde{i},j}) \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right] + \mathbb{E}^W \left[\left(\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) \right)^2 + 1 \right] \right] \right]. \quad (32)$$

Here, Eq. (28) is obtained by taking the expectation over \hat{U}_j , the definition of $\rho_i^{(k)}(\hat{U}_j)$, and $\hat{U} \perp (\tilde{Z}^{(k)}, W)$. Then, Eq. (29) is by the law of iterated expectations. Specifically, we condition on U_j , and recall that based on the Definition 1.2 the j -th training sample is $Z_{U_j,j}$. Step Eq. (30) is by some manipulations. Eq. (31) is obtained by

$$\left[\frac{1}{k-1} \ell(W, Z_{u_j,j}) + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) - \ell(W, Z_{\tilde{i},j}) \right]^2 \leq \left(\frac{1}{k-1} + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right)^2 \\ + \ell(W, Z_{\tilde{i},j})^2 - \ell(W, Z_{\tilde{i},j}) \frac{2}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}),$$

and

$$\left[\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) - \ell(W, Z_{u_j,j}) \right]^2 \leq 1 + \left(\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) \right)^2.$$

Finally last step is obtained by changing the order of the expectation over W and $\tilde{Z}^{(k)}$. Then, we can simplify Eq. (32) by considering

$$\mathbb{E}^W \left[\left(\frac{1}{k-1} + \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right)^2 + \ell(W, Z_{\tilde{i},j})^2 - 2\ell(W, Z_{\tilde{i},j}) \frac{1}{k-1} \sum_{i \in [k], i \neq \{\tilde{i}, u_j\}} \ell(W, Z_{i,j}) \right] \\ \leq R_{\mathcal{D}}(W)^2 \frac{-k^2 + k + 2}{(k-1)^2} + R_{\mathcal{D}}(W) \frac{k^2 + k - 5}{(k-1)^2} + \frac{1}{(k-1)^2} \triangleq A_1(k, R_{\mathcal{D}}(W)) \quad (33)$$

$$\mathbb{E}^W \left[\left(\frac{1}{k-1} \sum_{i \in [k], i \neq \tilde{i}} \ell(W, Z_{i,j}) \right)^2 + 1 \right] \leq R_{\mathcal{D}}(W)^2 \frac{k-2}{k-1} + R_{\mathcal{D}}(W) \frac{1}{k-1} + 1 \triangleq A_2(k, R_{\mathcal{D}}(W)). \quad (34)$$

Note that in Eq. (33) and Eq. (34), W and $Z_{i,j}$ s are independent, therefore $\mathbb{E}^W[\ell(W, Z_{i,j})] = R_{\mathcal{D}}(W)$. Also, we have $\text{Var}^W[\ell(W, Z_{i,j})] \leq R_{\mathcal{D}}(W)(1 - R_{\mathcal{D}}(W))$ because Bernoulli random variable has the

largest variance among the $[0, 1]$ -bounded random variables with the same mean. Plugging Eq. (33) and Eq. (34) into Eq. (32) we obtain

$$\mathbb{E}[\sum_{i=1}^k \rho_i^{(k)}(U_j) \ell(W, Z_{i,j})]^2 \leq \frac{1}{k} \mathbb{E}[(k-1)A_1(k, R_{\mathcal{D}}(W)) + A_2(k, R_{\mathcal{D}}(W))]. \quad (35)$$

We can upper bound the LHS of Eq. (35) by maximizing it over $R_{\mathcal{D}}(W)$ to obtain

$$\frac{\partial[(k-1)A_1(k, R) + A_2(k, R)]}{\partial R} = 0 \Rightarrow R^* = \frac{k^2 + k - 4}{2k^2 - 4k}.$$

We can plug the expression for R^* into Eq. (35) to get

$$\begin{aligned} \mathbb{E}[\sum_{i=1}^k \rho_i^{(k)}(U_j) \ell(W, Z_{i,j})]^2 &\leq \frac{1}{k} \mathbb{E}[(k-1)A_1(k, R^*) + A_2(k, R^*)] \\ &= \frac{k^3 + 7k^2 - 8k - 16}{4(k^3 - 2k^2)}. \end{aligned} \quad (36)$$

Then, plugging Eq. (36) into Eq. (25) yields

$$\inf_{\lambda \geq 0} \frac{\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{\lambda} + \frac{\exp(\frac{\lambda}{n}) - \frac{\lambda}{n} - 1}{\frac{\lambda}{n}} \frac{k^3 + 7k^2 - 8k - 16}{4(k^3 - 2k^2)} \geq \text{EGE}_{\mathcal{D}}(\mathcal{A}). \quad (37)$$

Finally, the closed form solution of Eq. (37) is given by

$$\begin{aligned} \frac{\partial[\frac{\text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{\lambda} + \frac{\exp(\frac{\lambda}{n}) - \frac{\lambda}{n} - 1}{\frac{\lambda}{n}} \frac{k^3 + 7k^2 - 8k - 16}{4(k^3 - 2k^2)}]}{\partial \lambda} &= 0 \Rightarrow \\ \lambda^* &= n \mathfrak{W} \left(\left(\frac{4(k^3 - 2k^2) \text{CMI}_{\mathcal{D}}^k(\mathcal{A})}{n(k^3 + 7k^2 - 8k - 16)} - 1 \right) \exp(-1) \right) + n, \end{aligned}$$

which is the desired result. \square

D Proofs of Section 3

Proof of Theorem 3.1. With $k = 2$, recall from the introduction

$$\tilde{Z}^{(2)} = \begin{pmatrix} Z_{1,1} & \cdots & Z_{1,n} \\ Z_{2,1} & \cdots & Z_{2,n} \end{pmatrix} \sim \mathcal{D}^{\otimes 2n},$$

and $U = (U_1, \dots, U_n) \in \{1, 2\}^n$ where U_i s are IID, and the marginal distribution follows $U_i \sim \text{Bern}(\frac{1}{2})$ for $i \in [n]$. Furthermore, recall $S = \{Z_{U_1,1}, \dots, Z_{U_n,n}\}$. The expected generalization error can be written as

$$\mathbb{E}[R_{\mathcal{D}}(W) - \hat{R}_S(W)] = \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n (-1)^{U_i} (\ell(Z_{1,i}, W) - \ell(Z_{2,i}, W))\right] \quad (38)$$

$$= \mathbb{E}\left[\frac{1}{m} \sum_{i=1}^m (-1)^{U_{J_i}} (\ell(Z_{1,J_i}, W) - \ell(Z_{2,J_i}, W))\right], \quad (39)$$

where the last equality follows because J is independent of U_J , $\tilde{Z}^{(2)}$, and W .

Define \tilde{W} , \tilde{U}_J , and \tilde{J} such that $(W, U_J, J, \tilde{Z}^{(2)}) \stackrel{d}{=} (\tilde{W}, \tilde{U}_J, \tilde{J}, \tilde{Z}^{(2)})$, and \tilde{W} , \tilde{U}_J , and \tilde{J} are independent given $\tilde{Z}^{(2)}$. By the Donsker–Varadhan variational formula [5, Prop. 4.15] and the disintegration theorem [11, Thm. 6.4], for all measurable functions g in \mathcal{G} , i.e., the class of functions g such that $(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}]) (\exp g) < \infty$, with probability one we have

$$I^{\tilde{Z}^{(2)}}(W, J; U_J) = \text{KL}(\mathbb{P}^{\tilde{Z}^{(2)}}[W, J, U_J] \parallel \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J]) \quad (40)$$

$$= \sup_{g \in \mathcal{G}} \mathbb{P}^{\tilde{Z}^{(2)}}[W, J, U_J](g) - \log \left[\left(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J] \right) (\exp g) \right]. \quad (41)$$

Define $f(w, j, u_j) \triangleq \frac{\lambda}{m} \sum_{i=1}^m (-1)^{u_{j_i}} (\ell(z_{1,j_i}, w) - \ell(z_{2,j_i}, w))$ where $\lambda \geq 0$. By Eq. (41), we can write

$$I^{\tilde{Z}^{(2)}}(W, J; U_J) \geq \mathbb{P}^{\tilde{Z}^{(2)}}[W, U_J, J](f) - \log \left[\left(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}] \right) (\exp f) \right]. \quad (42)$$

Considering the second term of the RHS of Eq. (42), Hoeffding's lemma implies that

$$\left(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}] \right) (\exp f) \quad (43)$$

$$= \mathbb{E}^{\tilde{Z}^{(2)}} \exp \left(\frac{\lambda}{m} \sum_{i=1}^m (-1)^{\tilde{U}_{J_i}} \left(\ell(Z_{1,\tilde{J}_i}, \tilde{W}) - \ell(Z_{2,\tilde{J}_i}, \tilde{W}) \right) \right) \quad (44)$$

$$= \mathbb{E}^{\tilde{Z}^{(2)}} \mathbb{E}^{\tilde{Z}^{(2)}, \tilde{W}, \tilde{J}} \prod_{i=1}^m \exp \left(\frac{\lambda}{m} (-1)^{\tilde{U}_{J_i}} \left(\ell(Z_{1,\tilde{J}_i}, \tilde{W}) - \ell(Z_{2,\tilde{J}_i}, \tilde{W}) \right) \right) \quad (45)$$

$$\leq \mathbb{E}^{\tilde{Z}^{(2)}} \mathbb{E}^{\tilde{Z}^{(2)}, \tilde{W}, \tilde{J}} \prod_{i=1}^m \exp \left(\frac{\lambda^2 \left(\ell(Z_{1,\tilde{J}_i}, \tilde{W}) - \ell(Z_{2,\tilde{J}_i}, \tilde{W}) \right)^2}{2m^2} \right) \quad (46)$$

$$\leq \exp \left(\frac{\lambda^2}{2m} \right), \quad (47)$$

where we use the fact that

$$\left(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_J] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{J}] \right) (f) = 0.$$

Substituting the bound in Eq. (47) into Eq. (42), rearranging and taking expectations, we obtain

$$\mathbb{E} \frac{1}{m} \sum_{i=1}^m (-1)^{U_{J_i}} (\ell(Z_{1,J_i}, W) - \ell(Z_{2,J_i}, W)) \leq \mathbb{E} \inf_{\lambda \geq 0} \frac{I^{\tilde{Z}^{(2)}}(W, J; U_J) + \frac{\lambda^2}{2m}}{\lambda} \quad (48)$$

$$= \mathbb{E} \sqrt{\frac{2}{m} I^{\tilde{Z}^{(2)}}(W, J; U_J)}. \quad (49)$$

Moreover, we have a.s.

$$I^{\tilde{Z}^{(2)}}(J, W; U_J) - \underbrace{I^{\tilde{Z}^{(2)}}(J; U_J)}_0 = I^{\tilde{Z}^{(2)}}(W; U_J | J). \quad (50)$$

Here, $I^{\tilde{Z}^{(2)}}(J; U_J) = 0$ since J is independent of U_J given $\tilde{Z}^{(2)}$. Plugging Eq. (50) into Eq. (49), we obtain the desired result. \square

Proof of Theorem 3.2. Consider

$$I(W; U_{J^{(m_1)}} | \tilde{Z}^{(2)}, J^{(m_1)}) = H(U_{J^{(m_1)}} | J^{(m_1)}, \tilde{Z}^{(2)}) - H(U_{J^{(m_1)}} | J^{(m_1)}, \tilde{Z}^{(2)}, W) \quad (51)$$

$$= \frac{1}{\binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} H(U_{K_1} | \tilde{Z}^{(2)}) - H(U_{J^{(m_1)}} | J^{(m_1)}, \tilde{Z}^{(2)}, W) \quad (52)$$

$$= \frac{1}{\binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} H(U_{K_1} | \tilde{Z}^{(2)}) - \frac{1}{\binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} H(U_{K_1} | \tilde{Z}^{(2)}, W). \quad (53)$$

Eq. (52) follows because $\tilde{Z}^{(2)} \perp\!\!\!\perp J^{(m_1)}$, while Eq. (53) follows because the event $\{J^{(m_1)} = K_1\}$ is independent of $(W, U_{K_1}, \tilde{Z}^{(2)})$. Then

$$\frac{1}{m_1} I(W; U_{J^{(m_1)}} | \tilde{Z}^{(2)}, J^{(m_1)}) = \frac{1}{m_1 \binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} [H(U_{K_1}) - H(U_{K_1} | W, \tilde{Z}^{(2)})] \quad (54)$$

$$= \frac{1}{n} H(U) - \frac{1}{m_1 \binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} H(U_{K_1} | W, \tilde{Z}^{(2)}) \quad (55)$$

$$= \frac{1}{m_2 \binom{n}{m_2}} \sum_{K_2 \in [n]_{m_2}} H(U_{K_2}) - \frac{1}{m_1 \binom{n}{m_1}} \sum_{K_1 \in [n]_{m_1}} H(U_{K_1} | W, \tilde{Z}^{(2)}) \quad (56)$$

$$\leq \frac{1}{m_2 \binom{n}{m_2}} \sum_{K_2 \in [n]_{m_2}} [H(U_{K_2}) - H(U_{K_2} | W, \tilde{Z}^{(2)})] \quad (57)$$

$$= \frac{1}{m_2} I(W; U_{J^{(m_2)}} | \tilde{Z}^{(2)}, J^{(m_2)}). \quad (58)$$

Eq. (54) follows from Eq. (53) and the fact that $U \perp\!\!\!\perp \tilde{Z}^{(2)}$, while Eq. (55) follows from each element of U being IID. Eq. (57) follows from Lemma F.1, which is a modified version of the Han's inequality [22]. Finally, the last step follows from using the same line of reasoning as in Eq. (51) to Eq. (53).

Having established Eq. (4), the claim follows from

$$\mathbb{E} \sqrt{\frac{2}{m} I^{\tilde{Z}^{(2)}}(W; U_J | J)} \leq \sqrt{\frac{2}{m} I(W; U_J | \tilde{Z}^{(2)}, J)} \quad (59)$$

$$\leq \sqrt{\frac{2}{n} I(W; U | \tilde{Z}^{(2)})}, \quad (60)$$

where Eq. (59) is Jensen's inequality, and Eq. (60) is the direct application of Eq. (4) with $m_1 = m$ and $m_2 = n$. This proves the desired result. \square

Proof of Theorem 3.4. By the Donsker–Varadhan variational formula [5, Prop. 4.15] and the disintegration theorem [11, Thm. 6.4], with probability one, for all measurable functions g such that $(\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_i])(\exp g) < \infty$, we have

$$I^{\tilde{Z}^{(2)}}(U_i, W) = \text{KL}(\mathbb{P}^{\tilde{Z}^{(2)}}[U_i, W] \| \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_i] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}]) \quad (61)$$

$$\geq \mathbb{P}^{\tilde{Z}^{(2)}}[U_i, W][g(W, \tilde{Z}^{(2)}, U_i)] - \log \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_i] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}][\exp(g(\tilde{W}, \tilde{Z}^{(2)}, \tilde{U}_i))] \quad (62)$$

where $(W, U_i, \tilde{Z}^{(2)}) \stackrel{d}{=} (\tilde{W}, \tilde{U}_i, \tilde{Z}^{(2)})$ and $\tilde{W} \perp\!\!\!\perp \tilde{U}_i | \tilde{Z}^{(2)}$. For $i \in [n]$, let

$$g_i(W, \tilde{Z}^{(2)}, U_i) \triangleq \lambda (-1)^{U_i} (\ell(Z_{1,i}, W) - \ell(Z_{2,i}, W)).$$

Hoeffding's lemma implies that

$$\mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{U}_i] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[\tilde{W}][\exp(g_i(\tilde{W}, \tilde{Z}^{(2)}, \tilde{U}_i))] \leq \exp\left(\frac{\lambda^2}{2}\right), \quad (63)$$

where in the last line we have used $g_i \in [-\lambda, \lambda]$ a.s. From (62), we obtain

$$\mathbb{E}^{\tilde{Z}^{(2)}} (-1)^{U_i} (\ell(Z_{1,i}, W) - \ell(Z_{2,i}, W)) \leq \inf_{\lambda \geq 0} \frac{\text{KL}(\mathbb{P}^{\tilde{Z}^{(2)}}[U_i, W] \| \mathbb{P}^{\tilde{Z}^{(2)}}[U_i] \otimes \mathbb{P}^{\tilde{Z}^{(2)}}[W]) + \frac{\lambda^2}{2}}{\lambda} \quad (64)$$

$$= \sqrt{2I^{\tilde{Z}^{(2)}}(W; U_i)}. \quad (65)$$

Then, averaging over i and taking expectations,

$$\mathbb{E}[R_D(W) - \hat{R}_S(W)] = \mathbb{E} \frac{1}{n} \sum_{i=1}^n \mathbb{E}^{\tilde{Z}^{(2)}} (-1)^{U_i} (\ell(Z_{1,i}, W) - \ell(Z_{2,i}, W)) \quad (66)$$

$$\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \sqrt{2I^{\tilde{Z}^{(2)}}(W; U_i)}. \quad (67)$$

\square

Proof of Theorem 3.7. For any two random measures $P(\tilde{Z}^{(2)}, U_{J^c}, J)$ and $Q(\tilde{Z}^{(2)}, U)$ on \mathcal{W} , the Donsker–Varadhan variational formula [5, Prop. 4.15] and the disintegration theorem [11, Thm. 6.4], give that with probability one

$$\text{KL}(Q(\tilde{Z}^{(2)}, U) \| P(\tilde{Z}^{(2)}, U_{J^c}, J)) = \sup_{g \in \mathcal{G}} \left(Q(\tilde{Z}^{(2)}, U)[g] - \log P(\tilde{Z}^{(2)}, U_{J^c}, J)[\exp g] \right) \quad (68)$$

where $\mathcal{G} = \{g : P(\tilde{Z}^{(2)}, U_{J^c}, J)(\exp g) < \infty\}$.

Let $g = \frac{\lambda}{m} \sum_{j \in J} (-1)^{U_j} (\ell(Z_{1,j}, W) - \ell(Z_{2,j}, W))$. First, note that

$$\mathbb{E}^{\tilde{Z}^{(2)}, U_{J^c}, J} \left[\frac{\lambda}{m} \sum_{j \in J} (-1)^{U_j} (\ell(Z_{1,j}, W) - \ell(Z_{2,j}, W)) \right] = 0.$$

This is because $\{U_j\}_{j \in J}$ are independent of $\tilde{Z}^{(2)}, U_{J^c}$, and J . Moreover, g is $[-\lambda, \lambda]$ -bounded. Therefore, we can use the Hoeffding’s lemma to obtain

$$\log P(\tilde{Z}^{(2)}, U_{J^c}, J)(\exp g) \leq \frac{\lambda^2}{2}.$$

Hence, from Eq. (68), we conclude that

$$\begin{aligned} & Q(\tilde{Z}^{(2)}, U) \left[\frac{1}{m} \sum_{j \in J} (-1)^{U_j} (\ell(Z_{1,j}, W) - \ell(Z_{2,j}, W)) \right] \\ & \leq \inf_{\lambda > 0} \frac{\text{KL}(Q(\tilde{Z}^{(2)}, U, J) \| P(\tilde{Z}^{(2)}, U_{J^c}, J))}{\lambda} + \frac{\lambda}{2} = \sqrt{2\text{KL}(Q(\tilde{Z}^{(2)}, U) \| P(\tilde{Z}^{(2)}, U_{J^c}, J))} \end{aligned}$$

almost surely. Finally, since $J \perp\!\!\!\perp (\tilde{Z}^{(2)}, U)$ we get

$$\begin{aligned} & Q(\tilde{Z}^{(2)}, U) \left[\frac{1}{m} \sum_{j \in J} (-1)^{U_j} (\ell(Z_{1,j}, W) - \ell(Z_{2,j}, W)) \right] \\ & = Q(\tilde{Z}^{(2)}, U) \left[\frac{1}{n} \sum_{i=1}^n (-1)^{U_i} (\ell(Z_{1,i}, W) - \ell(Z_{2,i}, W)) \right] \\ & = \mathbb{E} [R_{\mathcal{D}}(W) - \hat{R}_S(W)] \end{aligned}$$

The desired result follows. \square

E Proofs of Section 4

Proof of Theorem 4.2. Considering the generalization bound in Theorem 3.7 and Lemma 3.6, we can write

$$\begin{aligned} \mathbb{E} [R_{\mathcal{D}}(W) - \hat{R}_S(W)] & \leq \mathbb{E} \sqrt{2\text{KL}(Q_T(S) \| P_T(\tilde{Z}^{(2)}, U_{J^c}, J))} \\ & \leq \mathbb{E} \sqrt{\sum_{t=1}^T 2\mathbb{E}^{\tilde{Z}^{(2)}, U, J} \text{KL}(Q_{t|} \| P_{t|})}. \end{aligned} \quad (69)$$

First, note that from Eq. (11) it follows that

$$Q_{t|} = \mathcal{N}(\mu_{Q_{t|}}, \frac{2\eta_t}{\beta_t} \mathbb{I}_d),$$

where the mean is given by

$$\begin{aligned} & \mu_{Q_{t|}} = \\ & W_{t-1} - \eta_{t-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{t-1}) - \frac{\eta_{t-1}}{n} (\mathbb{1}\{U_J = 1\} \nabla \tilde{\ell}(Z_{1,J}, W_{t-1}) + \mathbb{1}\{U_J = 2\} \nabla \tilde{\ell}(Z_{2,J}, W_{t-1})). \end{aligned}$$

Next, we propose the following construction of $P_{t|}$. Note that $P_{t|}$ is \mathcal{F}_t -measurable random probability measure where

$$\mathcal{F}_t = \sigma(S_{J^c}, Z_{1,J}, Z_{2,J}, J, W_{0:t-1}).$$

Hence we can exploit the information in the trajectory up to time t to construct $P_{t|}$. In particular, we use the information in \mathcal{F}_t to perform a binary hypothesis testing in which the two hypotheses are defined as

$$\begin{aligned}\mathcal{H}_1 : U_J &= 1, \\ \mathcal{H}_2 : U_J &= 2.\end{aligned}$$

Equivalently, \mathcal{H}_1 and \mathcal{H}_2 can also be described as the hypotheses that $Z_{1,J}$ is a member of the training set and $Z_{2,J}$ is a member of the training set, respectively. Denote $\pi_t = (\pi_{t,1}, \pi_{t,2})$ as a probability vector whose i -th element shows the belief of the prior at time t that the true hypothesis is \mathcal{H}_i for $i \in \{1, 2\}$. Then, we consider the prior as

$$P_{t|} = \mathcal{N}(\mu_{P_{t|}}, \frac{2\eta_{t-1}}{\beta_{t-1}} \mathbb{I}_d), \quad (70)$$

where

$$\mu_{P_{t|}} = W_{t-1} - \eta_{t-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{t-1}) - \frac{\eta_{t-1}}{n} (\pi_{t,1} \nabla \tilde{\ell}(Z_{1,J}, W_{t-1}) + \pi_{t,2} \nabla \tilde{\ell}(Z_{2,J}, W_{t-1})). \quad (71)$$

Here $\pi_1 = (\frac{1}{2}, \frac{1}{2})$. Then, we construct the belief vector π_t for $t \geq 2$ using the log-likelihood ratio as

$$\pi_t = \left(\theta \left(\log \frac{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_1]}{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_2]} \right), 1 - \theta \left(\log \frac{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_1]}{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_2]} \right) \right), \quad (72)$$

where $\theta : \mathbb{R} \rightarrow [0, 1]$. Also, we might expect that the optimal θ satisfies $\theta(0) = \frac{1}{2}$, $\lim_{x \rightarrow \infty} \theta(x) = 1$, and $\lim_{x \rightarrow -\infty} \theta(x) = 0$.

Denote probability density function $\mathbb{P}^{\tilde{Z}^{(2)}, U_{J^c}, \mathcal{H}_k, W_0}[W_{1:t-1}]$ as $f_k(W_{1:t-1})$ for $k \in \{1, 2\}$. Due to Markov structure of the update rule in Eq. (11), we have

$$\begin{aligned}f_k(W_{1:t-1}) &= \\ \prod_{i=1}^{t-1} \left(\frac{\beta_{i-1}}{4\pi\eta_{i-1}} \right)^{\frac{d}{2}} \exp \left(- \frac{\beta_{i-1} \|W_i - W_{i-1} + \eta_{i-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{i-1}) + \frac{\eta_{i-1}}{n} \nabla \tilde{\ell}(Z_{k,J}, W_{i-1})\|^2}{4\eta_{i-1}} \right).\end{aligned} \quad (73)$$

Here, Eq. (73) is obtained by the Markov property of the update rule in Eq. (11). Then, since the prior distribution on \mathcal{H}_1 and \mathcal{H}_2 is uniform, we have

$$\log \frac{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_1]}{\mathbb{P}^{\mathcal{F}_t}[\mathcal{H}_2]} = \log \frac{f_1(W_{1:t-1})}{f_2(W_{1:t-1})} \quad (74)$$

$$= Y_{t,2} - Y_{t,1}, \quad (75)$$

where $Y_{t,1}$ and $Y_{t,2}$ are given by

$$\begin{aligned}Y_{t,1} &\triangleq \sum_{i=1}^{t-1} \frac{\beta_{i-1}}{4\eta_{i-1}} \|W_i - W_{i-1} + \eta_{i-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{i-1}) + \frac{\eta_{i-1}}{n} \nabla \tilde{\ell}(Z_{1,J}, W_{i-1})\|^2, \\ Y_{t,2} &\triangleq \sum_{i=1}^{t-1} \frac{\beta_{i-1}}{4\eta_{i-1}} \|W_i - W_{i-1} + \eta_{i-1} \frac{n-1}{n} \nabla \tilde{R}_{S_{J^c}}(W_{i-1}) + \frac{\eta_{i-1}}{n} \nabla \tilde{\ell}(Z_{2,J}, W_{i-1})\|^2.\end{aligned} \quad (76)$$

Therefore, the belief vector is given by

$$\pi_t = \left(\theta((Y_{t,2} - Y_{t,1})), 1 - \theta((Y_{t,2} - Y_{t,1})) \right), \quad (77)$$

where $Y_{0,1} = Y_{0,2} = 0$ and for $t \geq 2$, $Y_{t,1}$ and $Y_{t,2}$ are given by Eq. (76). To conclude the proof, we obtain

$$\text{KL}(Q_T(S) \| P_T(\tilde{Z}^{(2)}, U_{J^c}, J)) \leq \sum_{t=1}^T \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \text{KL}(Q_t \| P_t) \quad (78)$$

$$= \sum_{t=1}^T \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \frac{\beta_{t-1} \eta_{t-1} \|(\mathbb{1}\{U_J = 1\} - \pi_{t,1}) \nabla \tilde{\ell}(Z_{1,J}, W_{t-1}) + (\mathbb{1}\{U_J = 2\} - \pi_{t,2}) \nabla \tilde{\ell}(Z_{2,J}, W_{t-1})\|^2}{4n^2} \quad (79)$$

$$= \sum_{t=1}^T \mathbb{E}^{\tilde{Z}^{(2)}, U, J} \frac{\beta_{t-1} \eta_{t-1} (\mathbb{1}\{U_J = 1\} - \pi_{t,1})^2 \|\nabla \tilde{\ell}(Z_{1,J}, W_{t-1}) - \nabla \tilde{\ell}(Z_{2,J}, W_{t-1})\|^2}{4n^2} \quad (80)$$

Finally, plugging Eq. (80) into Eq. (69), we get the desired result in Eq. (14). \square

F Conditional Han's Inequality

Lemma F.1. *Let (X_1, \dots, X_n, Y) be $n+1$ -dimensional random variable where X_1, \dots, X_N are discrete random variables. Then,*

$$\frac{1}{k \binom{n}{k}} \sum_{T \in [n]_k} H(X_T | Y)$$

is decreasing in k .

Proof. For notational convenience, let $\bar{H}_k(X_{[n]} | Y) = \frac{1}{\binom{n}{k}} \sum_{T \in [n]_k} H(X_T | Y)$. Note that if we manage to show that

$$\bar{H}_k(X_{[n]} | Y) - \bar{H}_{k-1}(X_{[n]} | Y) \leq \bar{H}_{k+1}(X_{[n]} | Y) - \bar{H}_k(X_{[n]} | Y), \quad (81)$$

then the result in Lemma F.1 follows. To show Eq. (81), we can write

$$\begin{aligned} & H(X_1, \dots, X_{k+1} | Y) + H(X_1, \dots, X_{k-1} | Y) \\ &= H(X_1, \dots, X_k | Y) + H(X_{k+1} | X_1, \dots, X_k, Y) + H(X_1, \dots, X_{k-1} | Y) \end{aligned} \quad (82)$$

$$\leq H(X_1, \dots, X_k | Y) + H(X_{k+1} | X_1, \dots, X_{k-1}, Y) + H(X_1, \dots, X_{k-1} | Y) \quad (83)$$

$$= H(X_1, \dots, X_k | Y) + H(X_1, \dots, X_{k-1}, X_{k+1} | Y). \quad (84)$$

Here in Eq. (83), we drop X_k from the condition in the second term. Therefore, we have

$$H(X_1, \dots, X_{k+1} | Y) + H(X_1, \dots, X_{k-1} | Y) \leq H(X_1, \dots, X_k | Y) + H(X_1, \dots, X_{k-1}, X_{k+1} | Y). \quad (85)$$

Then, by averaging Eq. (85) over all $n!$ permutation of $\{1, \dots, n\}$, we get the desired result in Eq. (81). \square

G Details of Experiments

In this section, we discuss the details behind the experiments as well as the details of minimizing the generalization bound in Theorem 4.2.

G.1 Network architectures and learning curve

Tables 2 to 5 summarize the hyper-parameters we used for the experiments. Also, in Fig. 3 we plot the learning curves for the experiments reported in Section 4.2.1.

G.2 Optimizing the bound over the choice of θ function

Our generalization bound in Theorem 4.2 consists of an infimum over the functions in Θ . To study the impact of infimum, we consider the family of functions Θ given by

$$\Theta = \{\theta_a(x) | \exists a > 0 \text{ such that } \theta_a(x) = \frac{1}{2}(1 + \text{erf}(\frac{x}{a})) \text{ or } \theta_a(x) = \frac{1}{2}(1 + \tanh(\frac{x}{a}))\}.$$

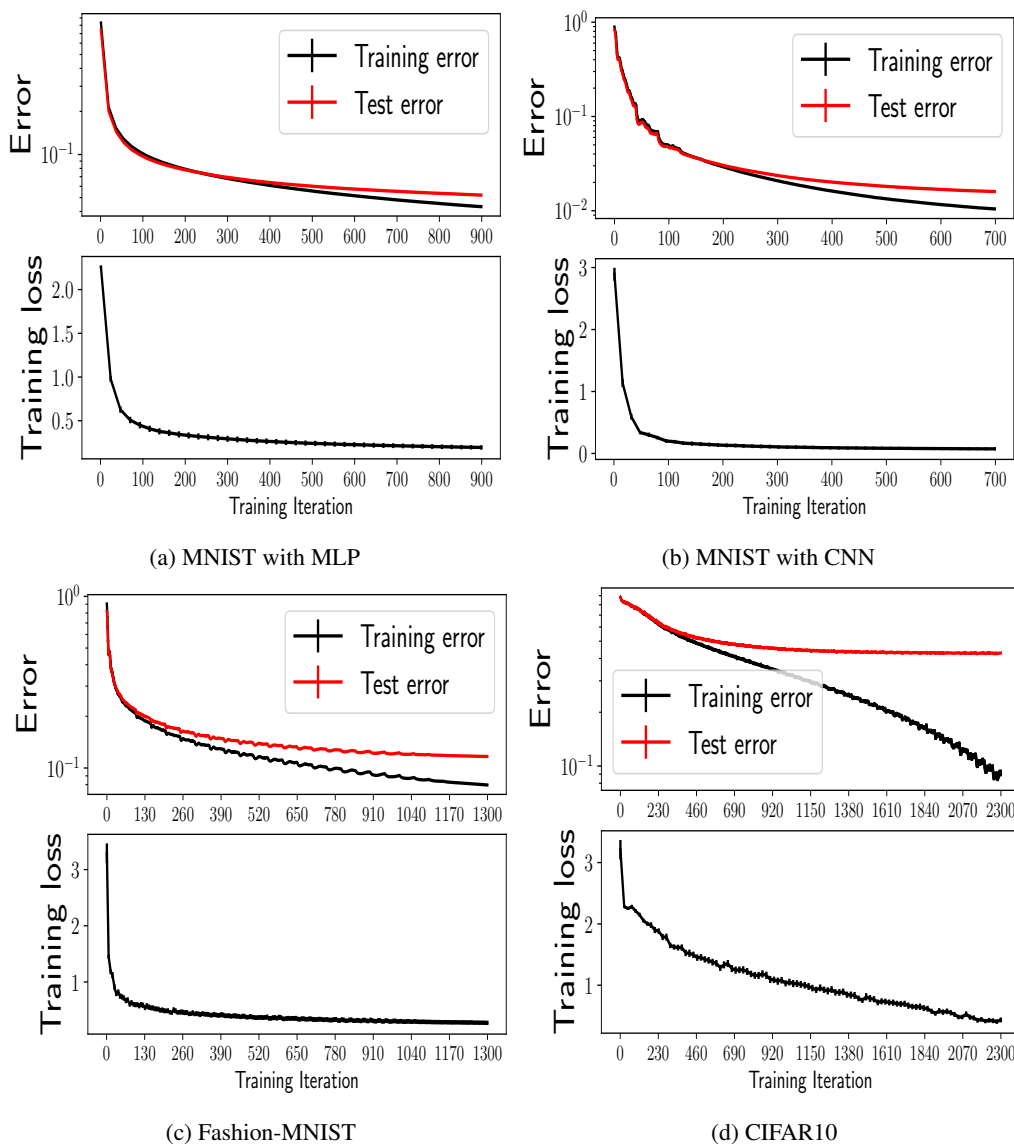


Figure 3: Learning curves. These plots show the training error, error on the test set, and the training loss. The loss functions is cross-entropy. Note y-axes for the error plots are log-scale.

Dataset	MNIST
Architecture	MLP(784-500-500-10)
η_t	$0.06 \times (0.95)^{\lceil \frac{t}{50} \rceil}$
$\frac{2\eta_t}{\beta_t}$	$10^{-8} + (3 \times 10^{-6} - 10^{-8}) \times \exp(-0.5 \lceil \frac{t}{50} \rceil)$
Number of iterations	900
Final training error	$4.33 \pm 0.01\%$
Generalization error	$0.88 \pm 0.01\%$
Number of training examples	20000
Number of runs	100

Table 2: Details of Experiments reported for MNIST with MLP

Dataset	MNIST
Architecture	CL($5 \times 5(32)$)-MaxPool(2×2)-CL($5 \times 5(64)$) MaxPool(2×2)-FC(128)-FC(10)
η_t	$0.05 \times (0.90)^{\lceil \frac{t}{40} \rceil}$
$\frac{2\eta_t}{\beta_t}$	$10^{-8} + (10^{-5} - 10^{-8}) \times \exp(-0.5 \lceil \frac{t}{40} \rceil)$
Number of iterations	700
Final training error	$2.59 \pm 0.01\%$
Generalization error	$0.55 \pm 0.01\%$
Number of training examples	20000
Number of runs	100

Table 3: Details of Experiments reported for MNIST with CNN

Dataset	Fashion-FMNIST
Architecture	CL($5 \times 5(32)$)-MaxPool(2×2)-CL($5 \times 5(64)$) MaxPool(2×2)-FC(200)-FC(10)
η_t	$0.07 \times (0.95)^{\lceil \frac{t}{50} \rceil}$
$\frac{2\eta_t}{\beta_t}$	$5 \times 10^{-8} + (7 \times 10^{-6} - 5 \times 10^{-8}) \times \exp(-0.3 \lceil \frac{t}{50} \rceil)$
Number of iterations	1300
Final training error	$7.96 \pm 0.03\%$
Generalization error	$3.71 \pm 0.03\%$
Number of training examples	20000
Number of runs	100

Table 4: Details of Experiments reported for Fashion-MNIST with CNN

Dataset	CIFAR10
Architecture	CL($3 \times 3(32)$)-MaxPool(2×2)-CL($3 \times 3(64)$) CL($3 \times 3(32)$)-MaxPool(2×2)-FC(128)-FC(10)
η_t	$0.15 \times (0.98)^{\lceil \frac{t}{50} \rceil}$
$\frac{2\eta_t}{\beta_t}$	$10^{-9} + (3 \times 10^{-5} - 10^{-9}) \times \exp(-0.3 \lceil \frac{t}{50} \rceil)$
Number of iterations	2300
Final training error	$9.39 \pm 0.46\%$
Generalization error	$32.89 \pm 0.44\%$
Number of training examples	15000
Number of runs	100

Table 5: Details of Experiments reported for CIFAR10 with CNN

Then, we divide the samples of the optimization trajectory into two sets of equal size: training set and the test set. Then, we optimize over a to find the $\theta_{a^*}(x)$ that achieves the minimum expected generalization over the training set. The numbers reported in Table 1 are based on the evaluation of $\theta_{a^*}(x)$ over the test set. Thus, the number reported in Table 1 are unbiased estimate of the generalization bound in Theorem 4.2.