
RATT: Leveraging Unlabeled Data to Guarantee Generalization

Saurabh Garg¹ Sivaraman Balakrishnan^{1,2} J. Zico Kolter³ Zachary C. Lipton¹

Abstract

To assess generalization, machine learning scientists typically either (i) bound the generalization gap and then (after training) plug in the empirical risk to obtain a bound on the true risk; or (ii) validate empirically on holdout data. However, (i) typically yields vacuous guarantees for overparameterized models; and (ii) shrinks the training set and its guarantee erodes with each re-use of the holdout set. In this paper, we leverage unlabeled data to produce generalization bounds. After augmenting our (labeled) training set with randomly labeled data, we train in the standard fashion. Whenever classifiers achieve low error on the clean data but high error on the random data, our bound ensures that the true risk is low. We prove that our bound is valid for 0-1 empirical risk minimization and with linear classifiers trained by gradient descent. Our approach is especially useful in conjunction with deep learning due to the early learning phenomenon whereby networks fit true labels before noisy labels but requires one intuitive assumption. Empirically, on canonical computer vision and NLP tasks, our bound provides non-vacuous generalization guarantees that track actual performance closely. This work enables practitioners to certify generalization even when (labeled) holdout data is unavailable and provides insights into the relationship between random label noise and generalization.

1. Introduction

Typically, machine learning scientists establish generalization in one of two ways. One approach, favored by learning theorists, places an *a priori* bound on the gap between the empirical and true risks, usually in terms of the complex-

ity of the hypothesis class. After fitting the model on the available data, one can plug in the empirical risk to obtain a guarantee on the true risk. The second approach, favored by practitioners, involves splitting the available data into training and holdout partitions, fitting the models on the former and estimating the population risk with the latter.

Surely, both approaches are useful, with the former providing theoretical insights and the latter guiding the development of a vast array of practical technology. Nevertheless, both methods have drawbacks. Most *a priori* generalization bounds rely on uniform convergence and thus fail to explain the ability of overparameterized networks to generalize (Zhang et al., 2016; Nagarajan & Kolter, 2019b). On the other hand, provisioning a holdout dataset restricts the amount of labeled data available for training. Moreover, risk estimates based on holdout sets lose their validity with successive re-use of the holdout data due to adaptive overfitting (Murphy, 2012; Dwork et al., 2015; Blum & Hardt, 2015). However, recent empirical studies suggest that on large benchmark datasets, adaptive overfitting is surprisingly absent (Recht et al., 2019).

In this paper, we propose Randomly Assign, Train and Track (RATT), a new method that leverages unlabeled data to provide a *post-training* bound on the true risk (i.e., the population error). Here, we assign random labels to a fresh batch of unlabeled data, augmenting the clean training dataset with these randomly labeled points. Next, we train on this data, following standard risk minimization practices. Finally, we track the error on the randomly labeled portion of training data, estimating the error on the mislabeled portion and using this quantity to upper bound the population error.

Counterintuitively, we guarantee generalization by guaranteeing overfitting. Specifically, we prove that Empirical Risk Minimization (ERM) with 0-1 loss leads to lower error on the *mislabeled training data* than on the *mislabeled population*. Thus, if despite minimizing the loss on the combined training data, we nevertheless have high error on the mislabeled portion, then the (mislabeled) population error will be even higher. Then, by complementarity, the (clean) population error must be low. Finally, we show how to obtain this guarantee using randomly labeled (vs mislabeled data), thus enabling us to incorporate unlabeled data.

To expand the applicability of our idea beyond ERM on 0-1

¹Machine Learning Department, Carnegie Mellon University ²Department of Statistics and Data Science, Carnegie Mellon University ³Computer Science Department, Carnegie Mellon University. Correspondence to: Saurabh Garg <sgarg2@andrew.cmu.edu>.

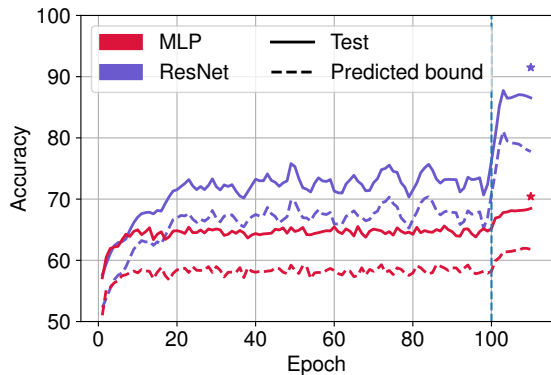


Figure 1. Predicted lower bound on the clean population error with ResNet and MLP on binary CIFAR. Results aggregated over 5 seeds. “*” denotes the best test performance achieved when training with only clean data and the same hyperparameters (except for the stopping point). The bound predicted by RATT (RHS in (2)) closely tracks the population accuracy on clean data.

error, we prove corresponding results for a linear classifier trained by gradient descent to minimize squared loss. Furthermore, leveraging the connection between early stopping and ℓ_2 -regularization in linear models (Ali et al., 2018; 2020; Suggala et al., 2018), our results extend to early-stopped gradient descent. Because we make no assumptions on the data distribution, our results on linear models hold for more complex models such as kernel regression and neural networks in the Neural Tangent Kernel (NTK) regime (Jacot et al., 2018; Du et al., 2018; 2019; Allen-Zhu et al., 2019b; Chizat et al., 2019).

Addressing practical deep learning models, our guarantee requires an additional (reasonable) assumption. Our experiments show that the bound yields non-vacuous guarantees that track test error across several major architectures on a range of benchmark datasets for computer vision and Natural Language Processing (NLP). Because, in practice, overparameterized deep networks exhibit an *early learning phenomenon*, fitting clean data before mislabeled data (Liu et al., 2020; Arora et al., 2019; Li et al., 2019), our procedure yields tight bounds in the early phases of learning. Experimentally, we confirm the early learning phenomenon in standard Stochastic Gradient Descent (SGD) training and illustrate the effectiveness of weight decay combined with large initial learning rates in avoiding interpolation to mislabeled data while maintaining fit on the training data, strengthening the guarantee provided by our method.

To be clear, we do not advocate RATT as a blanket replacement for the holdout approach. Our main contribution is to introduce a new theoretical perspective on generalization and to provide a method that may be applicable even when the holdout approach is unavailable. Of interest, unlike generalization bounds based on uniform-convergence that restrict the complexity of the hypothesis class (Neyshabur

et al., 2018; 2015; 2017b; Bartlett et al., 2017; Nagarajan & Kolter, 2019a), our *post hoc* bounds depend only on the fit to mislabeled data. We emphasize that our theory does not guarantee *a priori* that early learning should take place but only *a posteriori* that when it does, we can provide non-vacuous bounds on the population error. Conceptually, this finding underscores the significance of the early learning phenomenon in the presence of noisy labels and motivates further work to explain why it occurs.

2. Preliminaries

By $\|\cdot\|$, and $\langle \cdot, \cdot \rangle$ we denote the Euclidean norm and inner product, respectively. For a vector $v \in \mathbb{R}^d$, we use v_j to denote its j^{th} entry, and for an event E we let $\mathbb{I}[E]$ denote the binary indicator of the event.

Suppose we have a multiclass classification problem with the input domain $\mathcal{X} \subseteq \mathbb{R}^d$ and label space $\mathcal{Y} = \{1, 2, \dots, k\}^1$. By \mathcal{D} , we denote the distribution over $\mathcal{X} \times \mathcal{Y}$. A dataset $S := \{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}^n$ contains n points sampled i.i.d. from \mathcal{D} . By \mathcal{S} , \mathcal{T} , and $\tilde{\mathcal{S}}$, we denote the (uniform) empirical distribution over points in datasets S , T , and $\tilde{\mathcal{S}}$, respectively. Let \mathcal{F} be a class of hypotheses mapping \mathcal{X} to \mathbb{R}^k . A *training algorithm* \mathcal{A} : takes a dataset S and returns a classifier $f(\mathcal{A}, S) \in \mathcal{F}$. When the context is clear, we drop the parentheses for convenience. Given a classifier f and datum (x, y) , we denote the 0-1 error (i.e., classification error) on that point by $\mathcal{E}(f(x), y) := \mathbb{I}[y \notin \arg \max_{j \in \mathcal{Y}} f_j(x)]$. We express the *population error* on \mathcal{D} as $\mathcal{E}_{\mathcal{D}}(f) := \mathbb{E}_{(x,y) \sim \mathcal{D}} [\mathcal{E}(f(x), y)]$ and the *empirical error* on S as $\mathcal{E}_S(f) := \mathbb{E}_{(x,y) \sim S} [\mathcal{E}(f(x), y)] = \frac{1}{n} \sum_{i=1}^n \mathcal{E}(f(x_i), y_i)$.

Throughout, we consider a *random label assignment* procedure: draw $x \sim \mathcal{D}_{\mathcal{X}}$ (the underlying distribution over \mathcal{X}), and then assign a label sampled uniformly at random. We denote a randomly labeled dataset by $\tilde{\mathcal{S}} := \{(x_i, y_i)\}_{i=1}^m \sim \tilde{\mathcal{D}}^m$, where $\tilde{\mathcal{D}}$ is the distribution of randomly labeled data. By \mathcal{D}' , we denote the mislabeled distribution that corresponds to selecting examples (x, y) according to \mathcal{D} and then re-assigning the label by sampling among the incorrect labels $y' \neq y$ (renormalizing the label marginal).

3. Generalization Bound for RATT with ERM

We now present our generalization bound and proof sketches for ERM on the 0-1 loss (full proofs in App. A). For any dataset T , ERM returns the classifier \hat{f} that minimizes the empirical error:

$$\hat{f} := \arg \min_{f \in \mathcal{F}} \mathcal{E}_T(f). \quad (1)$$

¹For binary classification, we use $\mathcal{Y} = \{-1, 1\}$.

We focus first on binary classification. Assume we have a clean dataset $S \sim \mathcal{D}^n$ of n points and a randomly labeled dataset $\tilde{S} \sim \tilde{\mathcal{D}}^m$ of m ($< n$) points with labels in \tilde{S} are assigned uniformly at random. We show that with 0-1 loss minimization on the union of S and \tilde{S} , we obtain a classifier whose error on \mathcal{D} is upper bounded by a function of the empirical errors on clean data \mathcal{E}_S (lower is better) and on randomly labeled data $\mathcal{E}_{\tilde{S}}$ (higher is better):

Theorem 1. *For any classifier \hat{f} obtained by ERM (1) on dataset $S \cup \tilde{S}$, for any $\delta > 0$, with probability at least $1 - \delta$, we have*

$$\begin{aligned} \mathcal{E}_{\mathcal{D}}(\hat{f}) &\leq \mathcal{E}_S(\hat{f}) + 1 - 2\mathcal{E}_{\tilde{S}}(\hat{f}) \\ &\quad + \left(\sqrt{2}\mathcal{E}_{\tilde{S}}(\hat{f}) + 2 + \frac{m}{2n} \right) \sqrt{\frac{\log(4/\delta)}{m}}. \end{aligned} \quad (2)$$

In short, this theorem tells us that if after training on both clean and randomly labeled data, we achieve low error on the clean data but high error (close to $1/2$) on the randomly labeled data, then low population error is guaranteed. Note that because the labels in \tilde{S} are assigned randomly, the error $\mathcal{E}_{\tilde{S}}(f)$ for any fixed predictor f (not dependent on \tilde{S}) will be approximately $1/2$. Thus, if ERM produces a classifier that has not fit to the randomly labeled data, then $(1 - 2\mathcal{E}_{\tilde{S}}(\hat{f}))$ will be approximately 0, and our error will be determined by the fit to clean data. The final term accounts for finite sample error—notably, it (i) does not depend on the complexity of the hypothesis class; and (ii) approaches 0 at a $\mathcal{O}(1/\sqrt{m})$ rate (for $m < n$).

Our proof strategy unfolds in three steps. First, in Lemma 1 we bound $\mathcal{E}_{\mathcal{D}}(\hat{f})$ in terms of the error on the mislabeled subset of \tilde{S} . Next, in Lemmas 2 and 3, we show that the error on the mislabeled subset can be accurately estimated using only clean and randomly labeled data.

To begin, assume that we actually knew the original labels for the randomly labeled data. By \tilde{S}_C and \tilde{S}_M , we denote the clean and mislabeled portions of the randomly labeled data, respectively (with $\tilde{S} = \tilde{S}_M \cup \tilde{S}_C$). Note that for binary classification, a lower bound on mislabeled population error $\mathcal{E}_{\mathcal{D}'}(\hat{f})$ directly upper bounds the error on the original population $\mathcal{E}_{\mathcal{D}}(\hat{f})$. Thus we only need to prove that the empirical error on the mislabeled portion of our data is lower than the error on unseen mislabeled data, i.e., $\mathcal{E}_{\tilde{S}_M}(\hat{f}) \leq \mathcal{E}_{\mathcal{D}'}(\hat{f}) = 1 - \mathcal{E}_{\tilde{S}_C}(\hat{f})$ (upto $\mathcal{O}(1/\sqrt{m})$).

Lemma 1. *Assume the same setup as in Theorem 1. Then for any $\delta > 0$, with probability at least $1 - \delta$ over the random draws of mislabeled data \tilde{S}_M , we have*

$$\mathcal{E}_{\mathcal{D}}(\hat{f}) \leq 1 - \mathcal{E}_{\tilde{S}_M}(\hat{f}) + \sqrt{\frac{\log(1/\delta)}{m}}. \quad (3)$$

Proof Sketch. The main idea of our proof is to regard the clean portion of the data ($S \cup \tilde{S}_C$) as fixed. Then, there exists

a classifier f^* that is optimal over draws of the mislabeled data \tilde{S}_M . Formally,

$$f^* := \arg \min_{f \in \mathcal{F}} \mathcal{E}_{\tilde{\mathcal{D}}}(f),$$

where $\tilde{\mathcal{D}}$ is a combination of the *empirical distribution* over correctly labeled data $S \cup \tilde{S}_C$ and the (population) distribution over mislabeled data \mathcal{D}' . Recall that $\hat{f} := \arg \min_{f \in \mathcal{F}} \mathcal{E}_{S \cup \tilde{S}}(f)$. Since, \hat{f} minimizes 0-1 error on $S \cup \tilde{S}$, we have $\mathcal{E}_{S \cup \tilde{S}}(\hat{f}) \leq \mathcal{E}_{S \cup \tilde{S}}(f^*)$. Moreover, since f^* is independent of \tilde{S}_M , we have with probability at least $1 - \delta$ that

$$\mathcal{E}_{\tilde{S}_M}(f^*) \leq \mathcal{E}_{\mathcal{D}'}(f^*) + \sqrt{\frac{\log(1/\delta)}{m}}.$$

Finally, since f^* is the optimal classifier on $\tilde{\mathcal{D}}$, we have $\mathcal{E}_{\tilde{\mathcal{D}}}(f^*) \leq \mathcal{E}_{\tilde{\mathcal{D}}}(\hat{f})$. Combining the above steps and using the fact that $\mathcal{E}_{\mathcal{D}} = 1 - \mathcal{E}_{\mathcal{D}'}$, we obtain the desired result. \square

While the LHS in (3) depends on the unknown portion \tilde{S}_M , our goal is to use unlabeled data (with randomly assigned labels) for which the mislabeled portion cannot be readily identified. Fortunately, we do not need to identify the mislabeled points to estimate the error on these points in aggregate $\mathcal{E}_{\tilde{S}_M}(\hat{f})$. Note that because the label marginal is uniform, approximately half of the data will be correctly labeled and the remaining half will be mislabeled. Consequently, we can utilize the value of $\mathcal{E}_{\tilde{S}}(\hat{f})$ and an estimate of $\mathcal{E}_{\tilde{S}_C}(\hat{f})$ to lower bound $\mathcal{E}_{\tilde{S}_M}(\hat{f})$. We formalize this as follows:

Lemma 2. *Assume the same setup as Theorem 1. Then for any $\delta > 0$, with probability at least $1 - \delta$ over the random draws of \tilde{S} , we have $\left| 2\mathcal{E}_{\tilde{S}}(\hat{f}) - \mathcal{E}_{\tilde{S}_C}(\hat{f}) - \mathcal{E}_{\tilde{S}_M}(\hat{f}) \right| \leq 2\mathcal{E}_{\tilde{S}}(\hat{f}) \sqrt{\frac{\log(4/\delta)}{2m}}$.*

To complete the argument, we show that due to the exchangeability of the clean data S and the clean portion of the randomly labeled data \tilde{S}_C , we can estimate the error on the latter $\mathcal{E}_{\tilde{S}_C}(\hat{f})$ by the error on the former $\mathcal{E}_S(\hat{f})$.

Lemma 3. *Assume the same setup as Theorem 1. Then for any $\delta > 0$, with probability at least $1 - \delta$ over the random draws of \tilde{S}_C and S , we have $\left| \mathcal{E}_{\tilde{S}_C}(\hat{f}) - \mathcal{E}_S(\hat{f}) \right| \leq \left(1 + \frac{m}{2n}\right) \sqrt{\frac{\log(2/\delta)}{m}}$.*

Lemma 3 establishes a tight bound on the difference of the error of classifier \hat{f} on \tilde{S}_C and on S . The proof uses Hoeffding's inequality for randomly sampled points from a fixed population (Hoeffding, 1994; Bardenet et al., 2015).

Having established these core components, we can now summarize the proof strategy for Theorem 1. We bound the

population error on clean data (the term on the LHS of (2)) in three steps: (i) use Lemma 1 to upper bound the error on clean distribution $\mathcal{E}_{\mathcal{D}}(\hat{f})$, by the error on mislabeled training data $\mathcal{E}_{\tilde{S}_M}(\hat{f})$; (ii) approximate $\mathcal{E}_{\tilde{S}_M}(\hat{f})$ by $\mathcal{E}_{\tilde{S}_C}(\hat{f})$ and the error on randomly labeled training data (i.e., $\mathcal{E}_{\tilde{S}}(\hat{f})$) using Lemma 2; and (iii) use Lemma 3 to estimate $\mathcal{E}_{\tilde{S}_C}(\hat{f})$ using the error on clean training data ($\mathcal{E}_S(\hat{f})$).

Comparison with Rademacher bound Our bound in Theorem 1 shows that we can upper bound the clean population error of a classifier by estimating its accuracy on the clean and randomly labeled portions of the training data. Next, we show that our bound’s dominating term is upper bounded by the *Rademacher complexity* (Shalev-Shwartz & Ben-David, 2014), a standard distribution-dependent complexity measure.

Proposition 1. *Fix a randomly labeled dataset $\tilde{S} \sim \tilde{\mathcal{D}}^m$. Then for any classifier $f \in \mathcal{F}$ (possibly dependent on \tilde{S})² and for any $\delta > 0$, with probability at least $1 - \delta$ over random draws of \tilde{S} , we have*

$$1 - 2\mathcal{E}_{\tilde{S}}(f) \leq \mathbb{E}_{\epsilon, x} \left[\sup_{f \in \mathcal{F}} \left(\frac{\sum_i \epsilon_i f(x_i)}{m} \right) \right] + \sqrt{\frac{2 \log(\frac{2}{\delta})}{m}},$$

where ϵ is drawn from a uniform distribution over $\{-1, 1\}^m$ and x is drawn from $\mathcal{D}_{\mathcal{X}}^m$.

In other words, the proposition above highlights that the accuracy on the randomly labeled data is never larger than the Rademacher complexity of \mathcal{F} w.r.t. the underlying distribution over \mathcal{X} , implying that our bound is never looser than a bound based on Rademacher complexity. The proof follows by application of the bounded difference condition and McDiarmid’s inequality (McDiarmid, 1989). We now discuss extensions of Theorem 1 to regularized ERM and multiclass classification.

Extension to regularized ERM Consider any function $R : \mathcal{F} \rightarrow \mathbb{R}$, e.g., a regularizer that penalizes some measure of complexity for functions in class \mathcal{F} . Consider the following regularized ERM:

$$\hat{f} := \arg \min_{f \in \mathcal{F}} \mathcal{E}_S(f) + \lambda R(f), \quad (4)$$

where λ is a regularization constant. If the regularization coefficient is independent of the training data $S \cup \tilde{S}$, then our guarantee (Theorem 1) holds. Formally,

Theorem 2. *For any regularization function R , assume we perform regularized ERM as in (4) on $S \cup \tilde{S}$ and obtain a classifier \hat{f} . Then, for any $\delta > 0$, with probability at*

²We restrict \mathcal{F} to functions which output a label in $\mathcal{Y} = \{-1, 1\}$.

least $1 - \delta$, we have $\mathcal{E}_{\mathcal{D}}(\hat{f}) \leq \mathcal{E}_S(\hat{f}) + 1 - 2\mathcal{E}_{\tilde{S}}(\hat{f}) + \left(\sqrt{2}\mathcal{E}_{\tilde{S}}(\hat{f}) + 2 + \frac{m}{2n}\right) \sqrt{\frac{\log(1/\delta)}{m}}$.

A key insight here is that the proof of Theorem 1 treats the clean data S as fixed and considers random draws of the mislabeled portion. Thus a data-independent regularization function does not alter our chain of arguments and hence, has no impact on the resulting inequality. We prove this result formally in App. A.

We note one immediate corollary from Theorem 2: when learning any function f parameterized by w with L_2 -norm penalty on the parameters w , the population error with \hat{f} is determined by the error on the clean training data as long as the error on randomly labeled data is high (close to $1/2$).

Extension to multiclass classification Thus far, we have addressed binary classification. We now extend these results to the multiclass setting. As before, we obtain datasets S and \tilde{S} . Here, random labels are assigned uniformly among all classes.

Theorem 3. *For any regularization function R , assume we perform regularized ERM as in (4) on $S \cup \tilde{S}$ and obtain a classifier f . For a multiclass classification problem with k classes, for any $\delta > 0$, with probability at least $1 - \delta$, we have*

$$\mathcal{E}_{\mathcal{D}}(\hat{f}) \leq \mathcal{E}_S(\hat{f}) + (k-1) \left(1 - \frac{k}{k-1} \mathcal{E}_{\tilde{S}}(\hat{f}) \right) + c \sqrt{\frac{\log(\frac{4}{\delta})}{2m}}, \quad (5)$$

for some constant $c \leq (2k + \sqrt{k} + \frac{m}{n\sqrt{k}})$.

We first discuss the implications of Theorem 3. Besides empirical error on clean data, the dominating term in the above expression is given by $(k-1) \left(1 - \frac{k}{k-1} \mathcal{E}_{\tilde{S}}(\hat{f}) \right)$. For any predictor f (not dependent on \tilde{S}), the term $\mathcal{E}_{\tilde{S}}(\hat{f})$ would be approximately $(k-1)/k$ and for \hat{f} , the difference now evaluates to the accuracy of the randomly labeled data. Note that for binary classification, (5) simplifies to Theorem 1.

The core of our proof involves obtaining an inequality similar to (3). While for binary classification, we could upper bound $\mathcal{E}_{\tilde{S}_M}$ with $1 - \mathcal{E}_{\mathcal{D}}$ (in the proof of Lemma 1), for multiclass classification, error on the mislabeled data and accuracy on the clean data in the population are not so directly related. To establish an inequality analogous to (3), we break the error on the (unknown) mislabeled data into two parts: one term corresponds to predicting the true label on mislabeled data, and the other corresponds to predicting neither the true label nor the assigned (mis-)label. Finally, we relate these errors to their population counterparts to establish an inequality similar to (3).

4. Generalization Bound for RATT with Gradient Descent

In the previous section, we presented results with ERM on 0-1 loss. While minimizing the 0-1 loss is hard in general, these results provide important theoretical insights. In this section, we show parallel results for linear models trained with Gradient Descent (GD).

To begin, we introduce the setup and some additional notation. For simplicity, we begin discussion with binary classification with $\mathcal{X} = \mathbb{R}^d$. Define a linear function $f(x; w) := w^T x$ for some $w \in \mathbb{R}^d$ and $x \in \mathcal{X}$. Given training set S , we suppose that the parameters of the linear function are obtained via gradient descent on the following L_2 regularized problem:

$$\mathcal{L}_S(w; \lambda) := \sum_{i=1}^n (w^T x_i - y_i)^2 + \lambda \|w\|_2^2, \quad (6)$$

where $\lambda \geq 0$ is a regularization parameter. Our choice to analyze squared loss minimization for linear networks is motivated in part by its analytical convenience, and follows recent theoretical work which analyze neural networks trained via squared loss minimization in the Neural Tangent Kernel (NTK) regime when they are well approximated by linear networks (Jacot et al., 2018; Arora et al., 2019; Du et al., 2019; Hu et al., 2019). Moreover, recent research suggests that for classification tasks, squared loss minimization performs comparably to cross-entropy loss minimization (Muthukumar et al., 2020; Hui & Belkin, 2020).

For a given training set S , we use $S_{(i)}$ to denote the training set S with the i^{th} point removed. We now introduce one stability condition:

Condition 1 (Hypothesis Stability). *We have β hypothesis stability if our training algorithm \mathcal{A} satisfies the following for all $i \in \{1, 2, \dots, n\}$:*

$$\mathbb{E}_{\mathcal{S}, (x, y) \in \mathcal{D}} \left[\left| \mathcal{E}(f(x), y) - \mathcal{E}(f_{(i)}(x), y) \right| \right] \leq \frac{\beta}{n},$$

where $f_{(i)} := f(\mathcal{A}, S_{(i)})$ and $f := f(\mathcal{A}, S)$.

This condition is similar to a notion of stability called *hypothesis stability* (Bousquet & Elisseeff, 2002; Kearns & Ron, 1999; Elisseeff et al., 2003). Intuitively, Condition 1 states that empirical leave-one-out error and average population error of leave-one-out classifiers are close. This condition is mild and does not guarantee generalization. We discuss the implications in more detail in App. B.3.

Now we present the main result of this section. As before, we assume access to a clean dataset $\mathcal{S} = \{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}^n$ and randomly labeled dataset $\tilde{\mathcal{S}} = \{(x_i, y_i)\}_{i=n+1}^{n+m} \sim \tilde{\mathcal{D}}^m$. Let $\mathbf{X} = [x_1, x_2, \dots, x_{m+n}]$ and $\mathbf{y} = [y_1, y_2, \dots, y_{m+n}]$. Fix a positive learning rate η such

that $\eta \leq 1 / (\|\mathbf{X}^T \mathbf{X}\|_{\text{op}} + \lambda^2)$ and an initialization $w_0 = 0$. Consider the following gradient descent iterates to minimize objective (6) on $S \cup \tilde{S}$:

$$w_t = w_{t-1} - \eta \nabla_w \mathcal{L}_{S \cup \tilde{S}}(w_{t-1}; \lambda) \quad \forall t = 1, 2, \dots \quad (7)$$

Then we have $\{w_t\}$ converge to the limiting solution $\hat{w} = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}$. Define $\hat{f}(x) := f(x; \hat{w})$.

Theorem 4. *Assume that this gradient descent algorithm satisfies Condition 1 with $\beta = \mathcal{O}(1)$. Then for any $\delta > 0$, with probability at least $1 - \delta$ over the random draws of datasets \tilde{S} and S , we have:*

$$\begin{aligned} \mathcal{E}_{\mathcal{D}}(\hat{f}) \leq & \mathcal{E}_S(\hat{f}) + 1 - 2\mathcal{E}_{\tilde{S}}(\hat{f}) + \sqrt{\frac{4}{\delta} \left(\frac{1}{m} + \frac{3\beta}{m+n} \right)} \\ & + \left(\sqrt{2}\mathcal{E}_{\tilde{S}}(\hat{f}) + 1 + \frac{m}{2n} \right) \sqrt{\frac{\log(4/\delta)}{m}}. \end{aligned} \quad (8)$$

With a mild regularity condition, we establish the same bound on GD training with squared loss, notably the same dominating term on the population error, as in Theorem 1. In App. B.2, we present the extension to multiclass classification, where we again obtain a result parallel to Theorem 3.

Proof Sketch. Because squared loss minimization does not imply 0-1 error minimization, we cannot use arguments from Lemma 1. This is the main technical difficulty. To compare the 0-1 error at a train point with an unseen point, we use the closed-form expression for \hat{w} . We show that the train error on mislabeled points is less than the population error on the distribution of mislabeled data (parallel to Lemma 1).

For a mislabeled training point (x_i, y_i) in $\tilde{\mathcal{S}}$, we show that

$$\mathbb{I}[y_i x_i^T \hat{w} \leq 0] \leq \mathbb{I}[y_i x_i^T \hat{w}_{(i)} \leq 0], \quad (9)$$

where $\hat{w}_{(i)}$ is the classifier obtained by leaving out the i^{th} point from the training set. Intuitively, this condition states that the train error at a training point is less than the leave-one-out error at that point, i.e. the error obtained by removing that point and re-training. Using Condition 1, we then relate the average leave-one-out error (over the index i of the RHS in (9)) to the population error on the mislabeled distribution to obtain an inequality similar to (3). \square

Extensions to kernel regression Since the result in Theorem 4 does not impose any regularity conditions on the underlying distribution over $\mathcal{X} \times \mathcal{Y}$, our guarantees extend straightforwardly to kernel regression by using the transformation $x \rightarrow \phi(x)$ for some feature transform function ϕ . Furthermore, recent literature has pointed out a concrete connection between neural networks and kernel regression with

the so-called *Neural Tangent Kernel* (NTK) which holds in a certain regime where weights do not change much during training (Jacot et al., 2018; Du et al., 2019; 2018; Chizat et al., 2019). Using this concrete correspondence, our bounds on the clean population error (Theorem 4) extend to wide neural networks operating in the NTK regime.

Extensions to early stopped GD Often in practice, gradient descent is stopped early. We now provide theoretical evidence that our guarantees may continue to hold for an early stopped GD iterate. Concretely, we show that in expectation, the outputs of the GD iterates are close to that of a problem with data-independent regularization (as considered in Theorem 2). First, we introduce some notation. By $\mathcal{L}_S(w)$, we denote the objective in (6) with $\lambda = 0$. Consider the GD iterates defined in (7). Let $\tilde{w}_\lambda = \arg \min_w \mathcal{L}_S(w; \lambda)$. Define $f_t(x) := f(x; w_t)$ as the solution at the t^{th} iterate and $\tilde{f}_\lambda(x) := f(x; \tilde{w}_\lambda)$ as the regularized solution. Let κ be the condition number of the population covariance matrix and let s_{\min} be the minimum positive singular value of the empirical covariance matrix.

Proposition 2 (informal). *For $\lambda = \frac{1}{t\eta}$, we have*

$$\mathbb{E}_{x \sim \mathcal{D}_X} \left[(f_t(x) - \tilde{f}_\lambda(x))^2 \right] \leq c(t, \eta) \cdot \mathbb{E}_{x \sim \mathcal{D}_X} [f_t(x)^2],$$

where $c(t, \eta) \approx \kappa \cdot \min(0.25, \frac{1}{s_{\min}^2 t^2 \eta^2})$. An equivalent guarantee holds for a point x sampled from the training data.

The proposition above states that for large enough t , GD iterates stay close to a regularized solution with data-independent regularization constant. Together with our guarantees in Theorem 4 for regularized solution with $\lambda = \frac{1}{t\eta}$, Proposition 2 shows that our guarantees with RATT may hold on early stopped GD. See the formal result in App. B.4.

Remark Proposition 2 only bounds the expected squared difference between the t^{th} gradient descent iterate and a corresponding regularized solution. The expected squared difference and the expected difference of classification errors (what we wish to bound) are not related, in general. However, they can be related under standard low-noise (margin) assumptions. For instance, under the Tsybakov noise condition (Tsybakov et al., 1997; Yao et al., 2007), we can lower-bound the expression on the LHS of Proposition 2 with the difference of expected classification error.

Extensions to deep learning Note that the main lemma underlying our bound on (clean) population error states that when training on a mixture of clean and randomly labeled data, we obtain a classifier whose empirical error on the mislabeled training data is lower than its population error on the distribution of mislabeled data. We prove this for ERM on 0-1 loss (Lemma 1). For linear models (and networks in NTK regime), we obtained this result by assuming hypothesis stability and relating training error at a datum with the

leave-one-out error (Theorem 4). However, to extend our bound to deep models we must assume that training on the mixture of random and clean data leads to overfitting on the random mixture. Formally:

Assumption 1. *Let \hat{f} be a model obtained by training with an algorithm \mathcal{A} on a mixture of clean data S and randomly labeled data \tilde{S} . Then with probability $1 - \delta$ over the random draws of mislabeled data \tilde{S}_M , we assume that the following condition holds:*

$$\mathcal{E}_{\tilde{S}_M}(\hat{f}) \leq \mathcal{E}_{\mathcal{D}'}(\hat{f}) + c \sqrt{\frac{\log(1/\delta)}{2m}},$$

for a fixed constant $c > 0$.

Under Assumption 1, our results in Theorem 1, 2 and 3 extend beyond ERM with the 0-1 loss to general learning algorithms. We include the formal result in App. B.5. Note that given the ability of neural networks to interpolate the data, this assumption seems uncontroversial in the later stages of training. Moreover, concerning the early phases of training, recent research has shown that learning dynamics for complex deep networks resemble those for linear models (Nakkiran et al., 2019; Hu et al., 2020), much like the wide neural networks that we do analyze. Together, these arguments help to justify Assumption 1 and hence, the applicability of our bound in deep learning. Motivated by our analysis on linear models trained with gradient descent, we discuss conditions in App. B.6 which imply Assumption 1 for constant values $\delta > 0$. In the next section, we empirically demonstrate applicability of our bounds for deep models.

5. Empirical Study and Implications

Having established our framework theoretically, we now demonstrate its utility experimentally. First, for linear models and wide networks in the NTK regime where our guarantee holds, we confirm that our bound is not only valid, but closely tracks the generalization error. Next, we show that in practical deep learning settings, optimizing cross-entropy loss by SGD, the expression for our (0-1) ERM bound nevertheless tracks test performance closely and in numerous experiments on diverse models and datasets is never violated empirically.

Datasets To verify our results on linear models, we consider a toy dataset, where the class conditional distribution $p(x|y)$ for each label is Gaussian. For binary tasks, we use binarized CIFAR-10 (first 5 classes vs rest) (Krizhevsky & Hinton, 2009), binary MNIST (0-4 vs 5-9) (LeCun et al., 1998) and IMDB sentiment analysis dataset (Maas et al., 2011). For multiclass setup, we use MNIST and CIFAR-10.

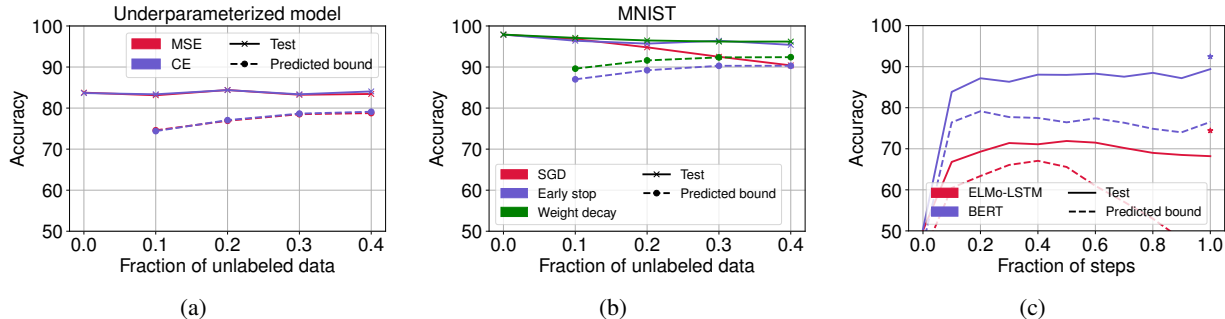


Figure 2. We plot the accuracy and corresponding bound (RHS in (1)) at $\delta = 0.1$. for binary classification tasks. Results aggregated over 3 seeds. (a) Accuracy vs fraction of unlabeled data (w.r.t clean data) in the toy setup with a linear model trained with GD. (b) Accuracy vs fraction of unlabeled data for a 2-layer wide network trained with SGD on binary MNIST. With SGD and no regularization (red curve in (b)), we interpolate the training data and hence the predicted lower bound is 0. However, with early stopping (or weight decay) we obtain tight guarantees. (c) Accuracy vs gradient iteration on IMDB dataset with unlabeled fraction fixed at 0.2. In plot (c), “*” denotes the best test accuracy with the same hyperparameters and training only on clean data. See App. C for exact hyperparameter values.

Architectures To simulate the NTK regime, we experiment with 2-layered wide networks both (i) with the second layer fixed at random initialization; (ii) and updating both layers’ weights. For vision datasets (e.g., MNIST and CIFAR10), we consider (fully connected) multilayer perceptrons (MLPs) with ReLU activations and ResNet18 (He et al., 2016). For the IMDB dataset, we train Long Short-Term Memory Networks (LSTMs; Hochreiter & Schmidhuber (1997)) with ELMo embeddings (Peters et al., 2018) and fine-tune an off-the-shelf uncased BERT model (Devlin et al., 2018; Wolf et al., 2020).

Methodology To bound the population error, we require access to both clean and unlabeled data. For toy datasets, we obtain unlabeled data by sampling from the underlying distribution over \mathcal{X} . For image and text datasets, we hold out a small fraction of the clean training data and discard their labels to simulate unlabeled data. We use the random labeling procedure described in Sec. 2. After augmenting clean training data with randomly labeled data, we train in the standard fashion. See App. C for experimental details.

Underparameterized linear models On toy Gaussian data, we train linear models with GD to minimize cross-entropy loss and mean squared error. Varying the fraction of randomly labeled data we observe that the accuracy on clean unseen data is barely impacted (Fig. 2(a)). This highlights that in low dimensional models adding randomly labeled data with the clean dataset (in toy setup) has minimal effect on the performance on unseen clean data. Moreover, we find that RATT offers a tight lower bound on the unseen clean data accuracy. We observe the same behavior with Stochastic Gradient Descent (SGD) training (ref. App. C). Observe that the predicted bound goes up as the fraction of unlabeled data increases. While the accuracy as dictated by the dominating term in the RHS of (2) decreases with an increase

in the fraction of unlabeled data, we observe a relatively sharper decrease in $\mathcal{O}_p(1/\sqrt{m})$ term of the bound, leading to an overall increase in the predicted accuracy bound. In this toy setup, we also evaluated a kernel regression bound from Bartlett & Mendelson (2002) (Theorem 21), however, the predicted kernel regression bound remains vacuous.

Wide Nets Next, we consider MNIST binary classification with a wide 2-layer fully-connected network. In experiments with SGD training on MSE loss without early stopping or weight decay regularization, we find that adding extra randomly label data hurts the unseen clean performance (Fig. 2(b)). Additionally, due to the perfect fit on the training data, our bound is rendered vacuous. However, with early stopping (or weight decay), we observe close to zero performance difference with additional randomly labeled data. Alongside, we obtain tight bounds on the accuracy on unseen clean data paying only a small price to negligible for incorporating randomly labeled data. Similar results hold for SGD and GD and when cross-entropy loss is substituted for MSE (ref. App. C).

Deep Nets We verify our findings on (i) ResNet-18 and 5-layer MLPs trained with binary CIFAR (Fig. 1); and (ii) ELMo-LSTM and BERT-Base models fine-tuned on the IMDB dataset (Fig. 2(c)). See App. C for additional results with deep models on binary MNIST. We fix the amount of unlabeled data at 20% of the clean dataset size and train all models with standard hyperparameters. Consistently, we find that our predicted bounds are never violated in practice. And as training proceeds, the fit on the mislabeled data increases with perfect overfitting in the interpolation regime rendering our bounds vacuous. However, with early stopping, our bound predicts test performance closely. For example, on IMDB dataset with BERT fine-tuning we predict 79.8 as the accuracy of the classifier, when the true

Dataset	Model	Pred. Acc	Test Acc.	Best Acc.
MNIST	MLP	93.1	97.4	97.9
	ResNet	96.8	98.8	98.9
CIFAR10	MLP	48.4	54.2	60.0
	ResNet	76.4	88.9	92.3

Table 1. Results on multiclass classification tasks. With pred. acc. we refer to the dominating term in RHS of (5). At the given sample size and $\delta = 0.1$, the remaining term evaluates to 30.7, decreasing our predicted accuracy by the same. We note that test acc. denotes the corresponding accuracy on unseen clean data. Best acc. is the best achievable accuracy with just training on just the clean data (and same hyperparameters except the stopping point). Note that across all tasks our predicted bound is tight and the gap between the best accuracy and test accuracy is small. Exact hyperparameters are included in App. C.

performance is 88.04 (and the best achievable performance on unseen data is 92.45). Additionally, we observe that our method tracks the performance from the beginning of the training and not just towards the end.

Finally, we verify our multiclass bound on MNIST and CIFAR10 with deep MLPs and ResNets (see results in Table 1 and per-epoch curves in App. C). As before, we fix the amount of unlabeled data at 20% of the clean dataset to minimize cross-entropy loss via SGD. In all four settings, our bound predicts non-vacuous performance on unseen data. In App. C, we investigate our approach on CIFAR100 showing that even though our bound grows pessimistic with greater numbers of classes, the error on the mislabeled data nevertheless tracks population accuracy.

6. Discussion and Connections to Prior Work

Implicit bias in deep learning Several recent lines of research attempt to explain the generalization of neural networks despite massive overparameterization via the *implicit bias* of gradient descent (Soudry et al., 2018; Gunasekar et al., 2018a;b; Ji & Telgarsky, 2019; Chizat & Bach, 2020). Noting that even for overparameterized linear models, there exist multiple parameters capable of overfitting the training data (with arbitrarily low loss), of which some generalize well and others do not, they seek to characterize the favored solution. Notably, Soudry et al. (2018) find that for linear networks, gradient descent converges (slowly) to the max margin solution. A complementary line of work focuses on the early phases of training, finding both empirically (Rolnick et al., 2017; Arpit et al., 2017) and theoretically (Arora et al., 2019; Li et al., 2020; Liu et al., 2020) that even in the presence of a small amount of mislabeled data, gradient descent is biased to fit the clean data first during initial phases of training. However, to best of our knowledge, no prior work leverages this phenomenon to obtain generalization guarantees on the clean data, which is the primary

focus of our work. Our method exploits this phenomenon to produce non-vacuous generalization bounds. Even when we cannot prove *a priori* that models will fit the clean data well while performing badly on the mislabeled data, we can observe that it indeed happens (often in practice), and thus, *a posteriori*, provide tight bounds on the population error. Moreover, by using regularizers like early stopping or weight decay, we can accentuate this phenomenon, enabling our framework to provide even tighter guarantees.

Non-vacuous generalization bounds In light of the inapplicability of traditional complexity-based bounds to deep neural networks (Zhang et al., 2016; Nagarajan & Kolter, 2019b), researchers have investigated alternative strategies to provide non-vacuous generalization bounds for deep nets (Neyshabur et al., 2015; 2017b;a; 2018; Dziugaite & Roy, 2017; Bartlett et al., 2017; Xu & Raginsky, 2017; Arora et al., 2018; Li & Liang, 2018; Allen-Zhu et al., 2019a; Pensia et al., 2018; Zhou et al., 2018; Nagarajan & Kolter, 2019a; Nakkiran et al., 2020). However, these bounds typically remain numerically loose relative to the true generalization error. However, (Dziugaite & Roy, 2017; Zhou et al., 2018) provide non-vacuous generalization guarantees. Specifically, they transform a base network into consequent networks that do not interpolate the training data either by adding stochasticity to the network weights (Dziugaite & Roy, 2017) or by compressing the original neural network (Zhou et al., 2018). In a similar spirit, our work provides guarantees on overparameterized networks by using early stopping or weight decay regularization, preventing a perfect fit on the training data. Notably, in our framework, the model can perfectly fit the clean portion of the data, so long as they nevertheless fit the mislabeled data poorly.

Leveraging noisy data to provide generalization guarantees In parallel work, Bansal et al. (2020) presented an upper bound on the generalization gap of linear classifiers trained on representations learned via self-supervision. Under certain noise-robustness and rationality assumptions on the training procedure, the authors obtained bounds dependent on the complexity of the linear classifier and independent of the complexity of representations. By contrast, we present generalization bounds for supervised learning that are non-vacuous by virtue of the early learning phenomenon. While both frameworks highlight how robustness to random label corruptions can be leveraged to obtain bounds that do not depend directly on the complexity of the underlying hypothesis class, our framework, methodology, claims, and generalization results are very different from theirs.

Other related work. A long line of work relates early stopped GD to a corresponding regularized solution (Friedman & Popescu, 2003; Yao et al., 2007; Suggala et al., 2018; Ali et al., 2018; Neu & Rosasco, 2018; Ali et al.,

2020). In the most relevant work, Ali et al. (2018) and Suggala et al. (2018) address a regression task, theoretically relating the solutions of early-stopped GD and a regularized problem, obtained with a data-independent regularization coefficient. Towards understanding generalization numerous stability conditions have been discussed (Kearns & Ron, 1999; Bousquet & Elisseeff, 2002; Mukherjee et al., 2006; Shalev-Shwartz et al., 2010). Hardt et al. (2016) studies the uniform stability property to obtain generalization guarantees with early-stopped SGD. While we assume a benign stability condition to relate leave-one-out performance with population error, we do not rely on any stability condition that implies generalization.

7. Conclusion and Future work

Our work introduces a new approach for obtaining generalization bounds that do not directly depend on the underlying complexity of the model class. For linear models, we provably obtain a bound in terms of the fit on randomly labeled data added during training. Our findings raise a number of questions to be explored next. While our empirical findings and theoretical results with 0-1 loss hold absent further assumptions and shed light on why the bound may apply for more general models, we hope to extend our proof that overfitting (in terms classification error) to the finite sample of mislabeled data occurs with SGD training on broader classes of models and loss functions. We hope to build on some early results (Nakkiran et al., 2019; Hu et al., 2020) which provide evidence that deep models behave like linear models in the early phases of training. We also wish to extend our framework to the interpolation regime. Since many important aspects of neural network learning take place within early epochs (Achille et al., 2017; Frankle et al., 2020), including gradient dynamics converging to very small subspace (Gur-Ari et al., 2018), we might imagine operationalizing our bounds in the interpolation regime by discarding the randomly labeled data after initial stages of training.

Acknowledgements

SG thanks Divyansh Kaushik for help with NLP code. This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA8750-19-1-1000. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation therein. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Laboratory, DARPA or the U.S. Government. SB acknowledges funding from the NSF grants DMS-1713003 and CIF-1763734, as well as Amazon AI and

a Google Research Scholar Award. ZL acknowledges Amazon AI, Salesforce Research, Facebook, UPMC, Abridge, the PwC Center, the Block Center, the Center for Machine Learning and Health, and the CMU Software Engineering Institute (SEI), for their generous support of ACMI Lab’s research on machine learning under distribution shift.

References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation*, 2016.
- Abou-Moustafa, K. and Szepesvári, C. An exponential efron-stein inequality for lq stable learning rules. *arXiv preprint arXiv:1903.05457*, 2019.
- Achille, A., Rovere, M., and Soatto, S. Critical learning periods in deep neural networks. *arXiv preprint arXiv:1711.08856*, 2017.
- Ali, A., Kolter, J. Z., and Tibshirani, R. J. A continuous-time view of early stopping for least squares. *arXiv preprint arXiv:1810.10082*, 2018.
- Ali, A., Dobriban, E., and Tibshirani, R. J. The implicit regularization of stochastic gradient flow for least squares. *arXiv preprint arXiv:2003.07802*, 2020.
- Allen-Zhu, Z., Li, Y., and Liang, Y. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pp. 6158–6169, 2019a.
- Allen-Zhu, Z., Li, Y., and Song, Z. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pp. 242–252. PMLR, 2019b.
- Arora, S., Ge, R., Neyshabur, B., and Zhang, Y. Stronger generalization bounds for deep nets via a compression approach. *arXiv preprint arXiv:1802.05296*, 2018.
- Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*, 2019.
- Arpit, D., Jastrzebski, S., Ballas, N., Krueger, D., Bengio, E., Kanwal, M. S., Maharaj, T., Fischer, A., Courville, A., Bengio, Y., et al. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 233–242. JMLR. org, 2017.

- Bansal, Y., Kaplun, G., and Barak, B. For self-supervised learning, rationality implies generalization, provably. *arXiv preprint arXiv:2010.08508*, 2020.
- Bardenet, R., Maillard, O.-A., et al. Concentration inequalities for sampling without replacement. *Bernoulli*, 21(3): 1361–1385, 2015.
- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. Spectrally-normalized margin bounds for neural networks. In *Advances in neural information processing systems*, pp. 6240–6249, 2017.
- Blum, A. and Hardt, M. The ladder: A reliable leaderboard for machine learning competitions. In *International Conference on Machine Learning*, pp. 1006–1014. PMLR, 2015.
- Bousquet, O. and Elisseeff, A. Stability and generalization. *Journal of machine learning research*, 2(Mar):499–526, 2002.
- Chizat, L. and Bach, F. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pp. 1305–1338. PMLR, 2020.
- Chizat, L., Oyallon, E., and Bach, F. On lazy training in differentiable programming. In *Advances in Neural Information Processing Systems*, pp. 2937–2947, 2019.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pp. 1675–1685. PMLR, 2019.
- Du, S. S., Zhai, X., Póczos, B., and Singh, A. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., and Roth, A. L. Preserving statistical validity in adaptive data analysis. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pp. 117–126, 2015.
- Dziugaite, G. K. and Roy, D. M. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. *arXiv preprint arXiv:1703.11008*, 2017.
- Elisseeff, A., Pontil, M., et al. Leave-one-out error and stability of learning algorithms with applications. *NATO science series sub series iii computer and systems sciences*, 190:111–130, 2003.
- Frankle, J., Schwab, D. J., and Morcos, A. S. The early phase of neural network training. *arXiv preprint arXiv:2002.10365*, 2020.
- Friedman, J. and Popescu, B. E. Gradient directed regularization for linear regression and classification. Technical report, Technical Report, Statistics Department, Stanford University, 2003.
- Gunasekar, S., Lee, J., Soudry, D., and Srebro, N. Implicit bias of gradient descent on linear convolutional networks. *arXiv preprint arXiv:1806.00468*, 2018a.
- Gunasekar, S., Woodworth, B., Bhojanapalli, S., Neyshabur, B., and Srebro, N. Implicit regularization in matrix factorization. In *2018 Information Theory and Applications Workshop (ITA)*, pp. 1–10. IEEE, 2018b.
- Gur-Ari, G., Roberts, D. A., and Dyer, E. Gradient descent happens in a tiny subspace. *arXiv preprint arXiv:1812.04754*, 2018.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pp. 1225–1234. PMLR, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep Residual Learning for Image Recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Hoeffding, W. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pp. 409–426. Springer, 1994.
- Hu, W., Li, Z., and Yu, D. Simple and effective regularization methods for training on noisily labeled data with generalization guarantee. *arXiv preprint arXiv:1905.11368*, 2019.
- Hu, W., Xiao, L., Adlam, B., and Pennington, J. The surprising simplicity of the early-time learning dynamics of neural networks. *arXiv preprint arXiv:2006.14599*, 2020.
- Hui, L. and Belkin, M. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks. *arXiv preprint arXiv:2006.07322*, 2020.
- Jacot, A., Gabriel, F., and Hongler, C. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.

- Ji, Z. and Telgarsky, M. The implicit bias of gradient descent on nonseparable data. In *Conference on Learning Theory*, pp. 1772–1798. PMLR, 2019.
- Kearns, M. and Ron, D. Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *Neural computation*, 11(6):1427–1453, 1999.
- Krizhevsky, A. and Hinton, G. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer, 2009.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86, 1998.
- Li, M., Soltanolkotabi, M., and Oymak, S. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. *arXiv preprint arXiv:1903.11680*, 2019.
- Li, M., Soltanolkotabi, M., and Oymak, S. Gradient descent with early stopping is provably robust to label noise for overparameterized neural networks. In *International Conference on Artificial Intelligence and Statistics*, pp. 4313–4324. PMLR, 2020.
- Li, Y. and Liang, Y. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pp. 8157–8166, 2018.
- Liu, S., Niles-Weed, J., Razavian, N., and Fernandez-Granda, C. Early-learning regularization prevents memorization of noisy labels. *arXiv preprint arXiv:2007.00151*, 2020.
- Maas, A., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pp. 142–150, 2011.
- McDiarmid, C. *On the method of bounded differences*, pp. 148–188. London Mathematical Society Lecture Note Series. Cambridge University Press, 1989.
- Mukherjee, S., Niyogi, P., Poggio, T., and Rifkin, R. Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Advances in Computational Mathematics*, 25(1):161–193, 2006.
- Murphy, K. P. *Machine Learning: A Probabilistic Perspective*. MIT Press, 2012.
- Muthukumar, V., Narang, A., Subramanian, V., Belkin, M., Hsu, D., and Sahai, A. Classification vs regression in overparameterized regimes: Does the loss function matter? *arXiv preprint arXiv:2005.08054*, 2020.
- Nagarajan, V. and Kolter, J. Z. Deterministic pac-bayesian generalization bounds for deep networks via generalizing noise-resilience. *arXiv preprint arXiv:1905.13344*, 2019a.
- Nagarajan, V. and Kolter, J. Z. Uniform convergence may be unable to explain generalization in deep learning. In *Advances in Neural Information Processing Systems*, pp. 11615–11626, 2019b.
- Nakkiran, P., Kaplun, G., Kalimeris, D., Yang, T., Edelman, B. L., Zhang, F., and Barak, B. Sgd on neural networks learns functions of increasing complexity. *arXiv preprint arXiv:1905.11604*, 2019.
- Nakkiran, P., Neyshabur, B., and Sedghi, H. The deep bootstrap: Good online learners are good offline generalizers. *arXiv preprint arXiv:2010.08127*, 2020.
- Neu, G. and Rosasco, L. Iterate averaging as regularization for stochastic gradient descent. In *Conference On Learning Theory*, pp. 3222–3242. PMLR, 2018.
- Neyshabur, B., Tomioka, R., and Srebro, N. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pp. 1376–1401, 2015.
- Neyshabur, B., Bhojanapalli, S., McAllester, D., and Srebro, N. Exploring generalization in deep learning. *arXiv preprint arXiv:1706.08947*, 2017a.
- Neyshabur, B., Bhojanapalli, S., and Srebro, N. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017b.
- Neyshabur, B., Li, Z., Bhojanapalli, S., LeCun, Y., and Srebro, N. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2018.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, 2019.
- Pensia, A., Jog, V., and Loh, P.-L. Generalization error bounds for noisy, iterative algorithms. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 546–550. IEEE, 2018.

- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *Proc. of NAACL*, 2018.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pp. 5389–5400. PMLR, 2019.
- Rolnick, D., Veit, A., Belongie, S., and Shavit, N. Deep learning is robust to massive label noise. *arXiv preprint arXiv:1705.10694*, 2017.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shalev-Shwartz, S., Shamir, O., Srebro, N., and Sridharan, K. Learnability, stability and uniform convergence. *The Journal of Machine Learning Research*, 11:2635–2670, 2010.
- Sherman, J. and Morrison, W. J. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Suggala, A., Prasad, A., and Ravikumar, P. K. Connecting optimization and regularization paths. In *Advances in Neural Information Processing Systems*, pp. 10608–10619, 2018.
- Tsybakov, A. B. et al. On nonparametric estimation of density level sets. *The Annals of Statistics*, 25(3):948–969, 1997.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 38–45. Association for Computational Linguistics, 2020.
- Xu, A. and Raginsky, M. Information-theoretic analysis of generalization capability of learning algorithms. *arXiv preprint arXiv:1705.07809*, 2017.
- Yao, Y., Rosasco, L., and Caponnetto, A. On early stopping in gradient descent learning. *Constructive Approximation*, 26(2):289–315, 2007.
- Zhang, C., Bengio, S., Hardt, M., Recht, B., and Vinyals, O. Understanding deep learning requires rethinking generalization. *arXiv preprint arXiv:1611.03530*, 2016.
- Zhou, W., Veitch, V., Austern, M., Adams, R. P., and Orbanz, P. Non-vacuous generalization bounds at the imagenet scale: a pac-bayesian compression approach. *arXiv preprint arXiv:1804.05862*, 2018.