
Retraining with Predicted Hard Labels Provably Increases Model Accuracy

Rudrajit Das^{*1} Inderjit S Dhillon²³ Alessandro Epasto¹ Adel Javanmard¹⁴ Jieming Mao¹ Vahab Mirrokni¹
Sujay Sanghavi³ Peilin Zhong¹

Abstract

The performance of a model trained with noisy labels is often improved by simply *retraining* the model with its *own predicted hard labels* (i.e., 1/0 labels). Yet, a detailed theoretical characterization of this phenomenon is lacking. In this paper, we theoretically analyze retraining in a linearly separable binary classification setting with randomly corrupted labels given to us and prove that retraining can improve the population accuracy obtained by initially training with the given (noisy) labels. To the best of our knowledge, this is the first such theoretical result. Retraining finds application in improving training with local label differential privacy (DP), which involves training with noisy labels. We empirically show that retraining selectively on the samples for which the predicted label matches the given label significantly improves label DP training at no extra privacy cost; we call this consensus-based retraining. For example, when training ResNet-18 on CIFAR-100 with $\epsilon = 3$ label DP, we obtain more than 6% improvement in accuracy with consensus-based retraining.

1 Introduction

We study the simple idea of **retraining** an *already trained* model with its **own predicted hard labels** (i.e., 1/0 labels and *not* the raw probabilities) when the given labels with which the model is initially trained are **noisy**. This is a simple yet effective way to boost a model’s performance in the presence of noisy labels. More formally, suppose we train a

discriminative model \mathcal{M} (for a classification problem) on a dataset of n samples and *noisy* label pairs $\{(\mathbf{x}_j, \tilde{y}_j)\}_{j=1}^n$. Let $\hat{\theta}_0$ be the final learned weight/checkpoint of \mathcal{M} and let $\tilde{y}_j = \mathcal{M}(\hat{\theta}_0, \mathbf{x}_j)$ be the current checkpoint’s predicted *hard* label for sample \mathbf{x}_j . Now, we propose to *retrain* \mathcal{M} with the \tilde{y}_j ’s in one of the following two ways:

- (i) **Full retraining:** Retrain \mathcal{M} with $\{(\mathbf{x}_j, \tilde{y}_j)\}_{j=1}^n$, i.e., retrain \mathcal{M} with the *predicted* labels of *all* the samples.
- (ii) **Consensus-based retraining:** Define $\mathcal{S}_{\text{cons}} := \{j \in \{1, \dots, n\} \mid \tilde{y}_j = \hat{y}_j\}$ to be the set of samples for which the predicted label matches the given noisy label; we call this the *consensus set*. Retrain \mathcal{M} with $\{(\mathbf{x}_j, \tilde{y}_j)\}_{j \in \mathcal{S}_{\text{cons}}}$, i.e., retrain \mathcal{M} with the *predicted* labels of *only the consensus set*.

Intuitively, retraining with predicted hard labels can be beneficial when the underlying classes are “**well-separated**”. In such a case, the model can potentially correctly predict the labels of many samples in the training set far away from the decision boundary which were originally incorrectly labeled and presented to it. As a result, the model’s accuracy (w.r.t. the actual labels) on the training data can be *significantly higher* than the accuracy of the noisy labels presented to it. Hence, retraining with predicted labels can potentially improve the model’s performance. This intuition is illustrated in Figure 1 where we consider a *separable* binary classification problem with noisy labels. The exact details are in Appendix A but importantly, Figures 1a and 1b correspond to versions of this problem with “large” and “small” separation, respectively. Please see the figure caption for detailed discussion but in summary, Figure 1 shows us that *the success of retraining depends on the degree of separation between the classes*.

The motivation for *consensus-based retraining* is that matching the predicted and given labels can potentially yield a smaller but *much more accurate* subset compared to the entire set; such a filtering effect can further improve the model’s performance. As we show in Section 5 (see Tables 3 and 11), this intuition bears out in practice.

There are plenty of ideas revolving around training a model with its own predictions, the two most common ones being

^{*}Work done while Rudrajit was a PhD student at the University of Texas at Austin and student researcher at Google Research. ¹Google Research ²Google ³University of Texas at Austin ⁴University of Southern California. Correspondence to: Rudrajit Das <dasrudrajit@google.com>, Alessandro Epasto <aepasto@google.com>, Adel Javanmard <ajavanma@usc.edu>.

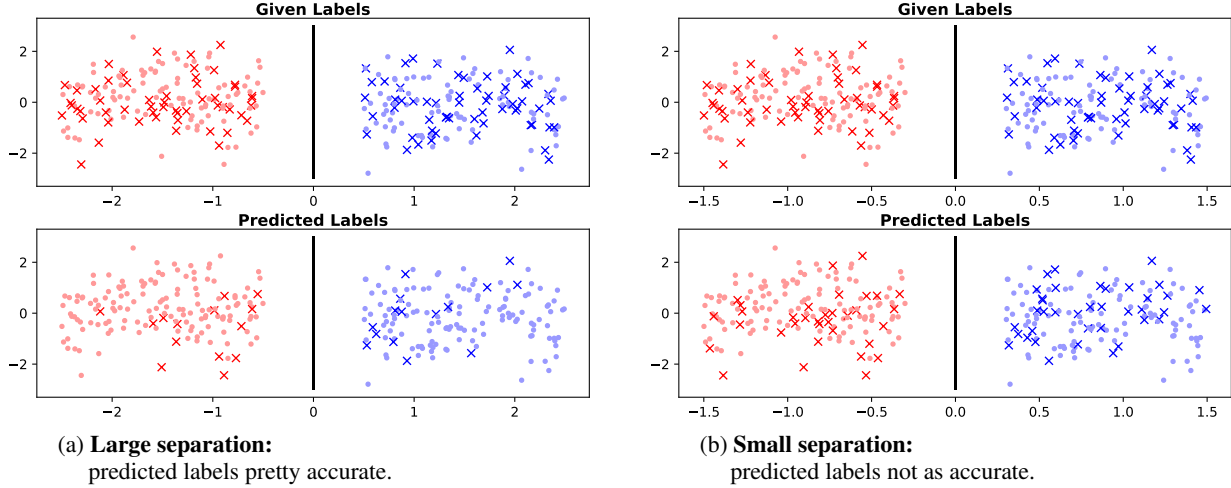


Figure 1. Retraining Intuition. Samples to the right (respectively, left) of the separator (black vertical line in the middle) and colored blue (respectively, red) have *actual* label +1 (respectively, -1). For both classes, the incorrectly labeled samples are marked by crosses (\times), whereas the correctly labeled samples are marked by dots (\circ) of the appropriate color. The amount of label noise and the number of training samples are the *same* in 1a and 1b. The top and bottom plots show the joint scatter plot of the training samples with the (noisy) labels given to us and the labels predicted by the model after training with the given labels, respectively. Notice that in 1a, *the model correctly predicts the labels of several samples that were given to it with the wrong label – especially, those that are far away from the separator*. This is not quite the case in 1b. This difference gets reflected in the *performance on the test set* after retraining. Specifically, in 1a, *retraining increases the test accuracy to 97.67% from 89%*. However, retraining yields no improvement in 1b. **So the success of retraining depends on the inter-class separation; in particular, retraining is beneficial when the classes are well-separated.**

self-training (Scudder, 1965; Yarowsky, 1995; Lee et al., 2013) and *self-distillation* (Furlanello et al., 2018; Mobahi et al., 2020); we discuss these and important differences from retraining in Section 2. However, from a *theoretical perspective*, we are not aware of any work proving that retraining a model with its predicted *hard* labels can be beneficial in the presence of label noise in any setting. In Section 4, **we derive the first theoretical result** (to our knowledge) showing that full retraining with hard labels *improves model accuracy*.

The primary reason for our interest in retraining is that it turned out to be a simple yet effective way to improve training with local¹ *label differential privacy* (DP) whose goal is to safeguard the privacy of labels in a supervised ML problem by injecting label noise (see Section 3 for a formal definition). Label DP is used in scenarios where only the labels are considered sensitive, e.g., advertising, recommendation systems, etc. (Ghazi et al., 2021). Importantly, *retraining can be applied on top of any label DP training algorithm at no extra privacy cost*. Our main *algorithmic contribution* is empirically demonstrating the *efficacy of consensus-based retraining in improving label DP training* (Section 5). Three things are worth clarifying here. First, as a meta-idea, retraining is not particularly new; however,

its application – especially with consensus-based filtering – as a *light-weight way to improve label DP training at no extra privacy cost* is new to our knowledge. Second, we are *not* advocating consensus-based retraining as a SOTA general-purpose algorithm for learning with noisy labels. Third, we do not view full retraining to be an algorithmic contribution; we consider it for theoretical analysis and as a baseline for consensus-based retraining.

Our **main contributions** can be summarized as follows:

(a) In Section 4, we consider a linearly separable binary classification problem wherein the data (feature) dimension is d , and we are given randomly flipped labels with the label flipping probability being $p < \frac{1}{2}$ independently for each sample. *Our main result is proving that full retraining with the predicted hard labels improves the population accuracy obtained by initially training with the given labels*, provided that p is close enough to $\frac{1}{2}$ and the dataset size n satisfies $\frac{d}{(1-2p)^2} \log \frac{d}{(1-2p)^2} \lesssim n \lesssim \frac{d^2}{(1-2p)^2}$ (“ \lesssim ” means bounded asymptotically, ignoring constant factor multiples); see Remark 4.10 for details. In addition, our results show that retraining becomes more beneficial as the amount of label noise (i.e., p) increases or as the degree of separation between the classes increases. To the best of our knowledge, **these are the first theoretical results** quantifying the benefits of retraining with predicted hard labels in the presence

¹In this paper, we focus only on local label DP. So throughout the paper, we will mostly omit the word “local” for conciseness.

of label noise. *The analysis of retraining is particularly challenging* because of the dependence of the predicted labels on the entire training set and the non-uniform/sample-dependent nature of label noise in the predicted labels; see the discussion after the statement of Theorem 4.8.

(b) In Section 5, we show the promise of **consensus-based retraining** (i.e., retraining on only those samples for which the predicted label matches the given noisy label) as a simple way to improve the performance of any label DP algorithm, at no extra privacy cost. For e.g., when training ResNet-18 on CIFAR-100 with $\epsilon = 3$ label DP, we obtain 6.4% improvement in accuracy with consensus-based retraining (see Table 2). The corresponding improvement for a small BERT model trained on AG News Subset (a news classification dataset) with $\epsilon = 0.5$ label DP is 11.7% (see Table 10).

2 Related Work

Self-training (ST). Retraining is similar in spirit to ST (Scudder, 1965; Yarowsky, 1995; Lee et al., 2013; Sohn et al., 2020) which is the process of progressively training a model with its own predicted hard labels in the *semi-supervised* setting. Our focus in this work is on the *fully supervised* setting. This is different from ST (in the semi-supervised setting) which typically selects samples based on the model’s confidence and hence, we call our algorithmic idea of interest *retraining* to distinguish it from ST. In fact, we show that our consensus-based sample selection strategy leads to better performance than confidence-based sample selection in Appendix J. There is a vast body of work on ST and related ideas; see Amini et al. (2022) for a survey. On the theoretical side also, there are several papers showing and quantifying different kinds of benefits of ST and related ideas (Carmon et al., 2019; Raghunathan et al., 2020; Kumar et al., 2020; Chen et al., 2020; Oymak & Gulcu, 2020; Wei et al., 2020; Zhang et al., 2022). But *none of these works characterize the pros/cons of ST or any related algorithm in the presence of noisy labels*. In contrast, we show that retraining can provably improve accuracy in the presence of label noise in Section 4. Empirically, ST-based ideas have been proposed to improve learning with noisy labels (Reed et al., 2014; Tanaka et al., 2018; Han et al., 2019; Nguyen et al., 2019; Li et al., 2020; Goel et al., 2022); but these works do not have rigorous theory. In the context of theory on label noise and model’s confidence, Zheng et al. (2020) show that if the model’s predicted score for the observed label is small, then the observed label is likely not equal to the true label. Note that the results of Zheng et al. (2020) pertain to the correctness of the observed labels, whereas our results pertain to the correctness of the predicted labels.

Self-distillation (SD). Retraining is also similar in principle to SD (Furlanello et al., 2018; Mobahi et al., 2020), where a teacher model is first trained with provided hard labels and

then its predicted *soft* labels are used to train a student model with the same architecture as the teacher. In contrast, we use predicted *hard* labels in retraining. Also, SD is usually employed with a temperature parameter (Hinton et al., 2015) to force the teacher and student models to be different; we do not have any such parameter in retraining. SD is known to ameliorate learning in the presence of noisy labels (Li et al., 2017) and this has been theoretically analyzed by Dong et al. (2019); Das & Sanghavi (2023). Dong et al. (2019) propose their own SD algorithm that uses *dynamically updated soft* labels and provide some complicated conditions of when their algorithm can learn the correct labels in the presence of noisy labels. In contrast, we analyze retraining with *fixed hard* labels. Das & Sanghavi (2023) analyze the standard SD algorithm in the presence of noisy labels with *fixed soft* labels but their analysis in the classification setting requires some strong assumptions such as access to the population, feature maps of all points in the same class having the same inner product, etc. We do not require such strong assumptions in this paper (in fact, we present sample complexity bounds).

Label differential privacy (DP). Label DP (described in detail in Section 3) is a relaxation of full-data DP wherein the privacy of only the labels (and not the features) is safeguarded (Chaudhuri & Hsu, 2011; Beimel et al., 2013; Wang & Xu, 2019; Ghazi et al., 2021; Malek Esmaeili et al., 2021; Ghazi et al., 2022; Badanidiyuru Varadaraja et al., 2023). In this work, we are not trying to propose a SOTA label DP algorithm (with an ingenious noise-injection scheme); instead, we advocate retraining as a simple post-processing step that can be applied on top of any label DP algorithm (regardless of the noise-injection scheme) to improve its performance, at no extra privacy cost. Similar to our goal, Tang et al. (2022) apply techniques from unsupervised and semi-supervised learning to improve label DP training. In particular, one of their steps involves keeping the given noisy label of a sample only if it matches a pseudo-label generated by unsupervised learning. This is similar in spirit to our consensus-based retraining scheme but a crucial difference is that we do not perform any unsupervised learning; we show that matching the given noisy label to the model’s own predicted label is itself pretty effective. Further, unlike our work, Tang et al. (2022) do not have any rigorous theory.

3 Preliminaries

Notation. For two functions $f(n)$ and $g(n)$, we write $f(n) \lesssim g(n)$ if there exists n_0 and a constant $C > 0$ such that for all $n \geq n_0$, $f(n) \leq Cg(n)$. For any positive integer $m \geq 1$, we denote the set $\{1, \dots, m\}$ by $[m]$. Let e_i denote the i^{th} canonical vector, namely, the vector of all zeros except a one in the i^{th} coordinate. We denote the ℓ_2 norm of a vector v by $\|v\|_{\ell_2}$, and the operator norm of a matrix

M by $\|M\|$. The unit d -dimensional sphere (i.e., the set of d -dimensional vectors with unit norm) is denoted by \mathbb{S}^{d-1} . For a random variable X , its sub-gaussian norm, denoted by $\|X\|_{\psi_2}$, is defined as $\|X\|_{\psi_2} := \sup_{q \geq 1} q^{-1/2} (\mathbb{E}|X|^q)^{1/q}$.² In addition, for random vector $\mathbf{X} \in \mathbb{R}^d$, its sub-gaussian norm is defined as $\|\mathbf{X}\|_{\psi_2} = \sup_{\mathbf{z} \in \mathbb{S}^{d-1}} \|\langle \mathbf{X}, \mathbf{z} \rangle\|_{\psi_2}$. We denote the CDF and complementary CDF (CCDF) of a standard normal variable (i.e., distributed as $\mathcal{N}(0, 1)$) by $\Phi(\cdot)$ and $\Phi^c(\cdot)$, respectively.

Definition 3.1 (Label differential privacy (DP)). A randomized algorithm \mathcal{A} taking as input a dataset and with range \mathcal{R} is said to be ϵ -labelDP if for any two datasets D and D' differing in the *label* of a single example and for any $S \subseteq \mathcal{R}$, it holds that $\mathbb{P}(\mathcal{A}(D) \in S) \leq e^\epsilon \mathbb{P}(\mathcal{A}(D') \in S)$.

Local label DP training involves injecting noise into the labels and then training with these noisy labels. The simplest way of injecting label noise to ensure label DP is randomized response (RR) introduced by Warner (1965). Specifically, suppose we require ϵ -labelDP for a problem with C classes, then the distribution of the output \hat{y} of RR when y is the true label is given by: $\mathbb{P}(\hat{y} = y) = \frac{e^\epsilon}{e^\epsilon + C - 1}$, and $\mathbb{P}(\hat{y} = z) = \frac{1}{e^\epsilon + C - 1}$ for any $z \neq y$. Our label noise model in Section 4 (Equation (4.2)) is actually RR for 2 classes. Based on vanilla RR, more sophisticated ways to inject label noise for better performance under label DP have been proposed (Ghazi et al., 2021; Malek Esmaeili et al., 2021). Our empirical results in Section 5 are with RR and the method of Ghazi et al. (2021).

4 Full Retraining in the Presence of Label Noise: Theoretical Analysis

Here we will analyze full retraining (as introduced in Section 1) for a linear setting with noisy labels. Since full retraining is the only kind of retraining we consider here, we will omit the word “full” subsequently in this section.

Problem setting: We consider binary classification in a linearly separable³ setting similar to the classical Gaussian mixture model. We will first describe the classical Gaussian mixture model and then describe our setting of interest, which is endowed with a positive margin.

In the classical Gaussian mixture model, each data point belongs to one of two classes $\{+1, -1\}$ with corresponding probabilities π_+ , π_- , such that $\pi_+ + \pi_- = 1$. The feature vector $\mathbf{x} \in \mathbb{R}^d$ for a data point with label $y \in \{+1, -1\}$ is drawn independently from $\mathcal{N}(y\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu} \in \mathbb{R}^d$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$. In other words the mean of the feature vector is $\pm\boldsymbol{\mu}$ depending on the data point’s label, and $\boldsymbol{\Sigma}$ is the

covariance matrix of the feature vector. We let $\gamma := \|\boldsymbol{\mu}\|_{\ell_2}$.

Gaussian mixture model with positive margin: Here each data point (\mathbf{x}, y) is generated independently by first sampling the label $y \in \{+1, -1\}$ with probabilities π_+ , π_- , and then generating the feature vector $\mathbf{x} \in \mathbb{R}^d$ as:

$$\mathbf{x} = y(1 + u)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}, \text{ where} \quad (4.1)$$

- $u > 0$ is drawn from a sub-gaussian distribution with unit sub-gaussian norm (i.e., $\|u\|_{\psi_2} = 1$).
- $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ independent of u .
- $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$ is the covariance matrix in the space orthogonal to $\boldsymbol{\mu}$. We have $\boldsymbol{\Sigma}^{1/2}\boldsymbol{\mu} = \mathbf{0}$ and $\boldsymbol{\Sigma}$ is of rank $d - 1$. Also, $\lambda_{\min} > 0$ and λ_{\max} are the minimum non-zero eigenvalue and the maximum eigenvalue of $\boldsymbol{\Sigma}$.

Note that under this data model, the projection of datapoints on the space orthogonal to $\boldsymbol{\mu}$ is Gaussian. Along $\boldsymbol{\mu}$, we have $\langle \mathbf{x}, \boldsymbol{\mu} \rangle = y(1 + u)\|\boldsymbol{\mu}\|_{\ell_2}^2$, and so the randomness of data along $\boldsymbol{\mu}$ is due to u . Further, since $u > 0$, $\text{sign}(\langle \mathbf{x}, \boldsymbol{\mu} \rangle) = y$. Thus, $\boldsymbol{\mu}$ is a separator for the data w.r.t. the labels. In addition, we have a margin of at least $\gamma = \|\boldsymbol{\mu}\|_{\ell_2}$; this is because $\frac{y\langle \mathbf{x}, \boldsymbol{\mu} \rangle}{\|\boldsymbol{\mu}\|_{\ell_2}} = (1 + u)\|\boldsymbol{\mu}\|_{\ell_2} \geq \gamma$. We will mostly refer to the margin γ as the “degree of separation”.

We are given access to a training set $\mathcal{T} = \{(\mathbf{x}_i, \hat{y}_i)\}_{i \in [n]}$ where for each $i \in [n]$, \hat{y}_i is a noisy version of the true label y_i (which we do not observe). Specifically:

$$\hat{y}_i = \begin{cases} y_i & \text{with probability } 1 - p, \\ -y_i & \text{with probability } p, \end{cases} \quad (4.2)$$

for some $p < 1/2$, and independently for each $i \in [n]$.⁴

4.1 Initial Training

Given the training set $\mathcal{T} = \{(\mathbf{x}_i, \hat{y}_i)\}_{i \in [n]}$, we consider the following linear classifier model⁵ (also used in Carmon et al. (2019)):

$$\hat{\boldsymbol{\theta}}_0 = \frac{1}{n} \sum_{i=1}^n \hat{y}_i \mathbf{x}_i. \quad (4.3)$$

This classifier’s predicted label for a sample $\mathbf{x} \in \mathbb{R}^d$ is $\text{sign}(\langle \mathbf{x}, \hat{\boldsymbol{\theta}}_0 \rangle)$. Note that $\langle \mathbf{x}, \hat{\boldsymbol{\theta}}_0 \rangle = \frac{1}{n} \sum_i \hat{y}_i \langle \mathbf{x}_i, \mathbf{x} \rangle$ is a

⁴In the context of ϵ -labelDP for a binary classification problem with randomized response, $p = \frac{1}{e^\epsilon + 1}$.

⁵This is a simplification to the least squares’ solution (LSS) obtained by excluding the empirical covariance matrix’s inverse. The LSS can be analyzed by bounding the deviation of the empirical covariance matrix from the population covariance matrix (which shrinks as $n \rightarrow \infty$), then adapting our current analysis to features pre-multiplied by the population covariance matrix’s inverse, and accounting for this deviation. This would only make the math more tedious without adding any meaningful insights.

²Refer to Vershynin (2010) for other equivalent definitions.

³More specifically, there exists a linear separator that perfectly classifies the data points according to their labels.

weighted average of the noisy labels in the training set, with the weights being $\frac{1}{n}\langle \mathbf{x}_i, \mathbf{x} \rangle$. So it is similar to kernel methods with the inner product kernel.

Our next result bounds the probability of $\hat{\theta}_0$ misclassifying a fixed test point \mathbf{x} , with respect to the randomness in the training set \mathcal{T} .

Theorem 4.1 (Initial training). *Consider $\mathbf{x} \notin \mathcal{T}$ and let y be its true label. We have*

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\theta}_0 \rangle) \neq y) \geq \alpha_0(\mathbf{x}) := \frac{1}{2\sqrt{2\pi}} \exp\left(-\frac{5(1 + \sqrt{n}(1-2p))^2 \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2}\right) \text{ and}$$

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\theta}_0 \rangle) \neq y) \leq \tilde{\alpha}_0(\mathbf{x}) := \frac{1}{2} \exp\left(-\frac{n(1-2p)^2 \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{8\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2}\right) + 2 \exp\left(-\frac{n}{32}(1-2p)^2\right).$$

The proof of Theorem 4.1 is in Appendix B. Notice that the learned classifier $\hat{\theta}_0$ is more likely to be wrong on the samples that are less aligned with (or closer to orthogonal to) the ground truth separator, i.e., $\boldsymbol{\mu}$. We can view $\hat{\theta}_0$ as a *noisy* label provider (a.k.a. pseudo-labeler) where the degree of label noise is **non-uniform** or **sample-dependent** unlike the original noisy source used to learn $\hat{\theta}_0$. Specifically, for a sample \mathbf{x} with true label y and predicted label $\tilde{y} = \text{sign}(\langle \mathbf{x}, \hat{\theta}_0 \rangle)$, we have:

$$\tilde{y} = \begin{cases} y & \text{with probability at least } \geq 1 - \tilde{\alpha}_0(\mathbf{x}) \text{ and} \\ -y & \text{with probability at most } \leq \tilde{\alpha}_0(\mathbf{x}), \end{cases}$$

where $\tilde{\alpha}_0(\mathbf{x})$ is as defined in Theorem 4.1. We will next define the population error of a classifier $\hat{\theta}$ as

$$\text{err}(\hat{\theta}) := \mathbb{E}_{(\mathbf{x}, y)} [\mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\theta} \rangle) \neq y)], \quad (4.4)$$

where the probability (inside the expectation) is w.r.t. the randomness in the training set used to learn $\hat{\theta}$. Similarly, the population accuracy of $\hat{\theta}$ is defined as

$$\text{acc}(\hat{\theta}) := \mathbb{E}_{(\mathbf{x}, y)} [\mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\theta} \rangle) = y)] = 1 - \text{err}(\hat{\theta}). \quad (4.5)$$

Depending on the context, our results and discussions will sometimes be presented in terms of error (lower is better) and at other times in terms of accuracy (higher is better).

Our next result bounds the error of the classifier $\hat{\theta}_0$.

Theorem 4.2 (Initial training’s population error). *We have*

$$\text{err}(\hat{\theta}_0) \geq \frac{(1 - e^{-\frac{d}{16}})}{4\sqrt{2\pi}} \exp\left(-\frac{160(1 + \sqrt{n}(1-2p))^2 \gamma^4}{\lambda_{\min}^2 d}\right), \quad (4.6)$$

$$\text{err}(\hat{\theta}_0) \leq \frac{1}{2} \exp\left(-\frac{n(1-2p)^2 \gamma^4}{16\lambda_{\max}^2 d}\right) + e^{-\frac{d}{8}} + 2 \exp\left(-\frac{n}{32}(1-2p)^2\right). \quad (4.7)$$

The proof of Theorem 4.2 is in Appendix C.

Remark 4.3 (Tightness of error bounds). When $n \lesssim \frac{\lambda_{\max}^2 d^2}{(1-2p)^2 \gamma^4}$, $\gamma \lesssim (\lambda_{\max}^2 d)^{1/4}$ and $\frac{\lambda_{\max}}{\lambda_{\min}} \lesssim 1$, then the lower and upper bounds for $\text{err}(\hat{\theta}_0)$ in Theorem 4.2 match (up to constant factors).

Based on Theorem 4.2, we have the following corollary.

Corollary 4.4 (Initial training’s sample complexity). *For any $\delta > 2e^{-d/8} + 4 \exp(-\frac{n}{32}(1-2p)^2)$, having number of samples $n \geq 8\lambda_{\max}^2 \frac{\log 1/\delta}{(1-2p)^2} \frac{d}{\gamma^4}$ ensures $\text{acc}(\hat{\theta}_0) > 1 - \delta$. In particular, when the label flipping probability satisfies $p > 2e^{-d/8} + 4 \exp(-\frac{n}{32}(1-2p)^2)$, then $n \geq 8\lambda_{\max}^2 \frac{\log 1/p}{(1-2p)^2} \frac{d}{\gamma^4}$ ensures $\text{acc}(\hat{\theta}_0) > 1 - p$, i.e., our learned classifier $\hat{\theta}_0$ has better accuracy than the source providing noisy labels (used to learn $\hat{\theta}_0$).*

Remark 4.5 (Effect of degree of separation). As the parameter quantifying the degree of separation γ decreases, the error bound in Theorem 4.2 increases and the sample complexity required to outperform the noisy label source in Corollary 4.4 increases. This is consistent with our intuition that a classification task should become harder as the degree of separation reduces; we saw this in Figure 1.

We conclude this subsection by deriving an information-theoretic lower bound on the sample complexity of *any* classifier to argue that $\hat{\theta}_0$ attains the optimal sample complexity with respect to d and p .

Theorem 4.6 (Information-theoretic lower bound on sample complexity). *With a slight generalization of notation, let $\text{acc}(\hat{\theta}; \boldsymbol{\mu})$ denote the accuracy of the classifier $\hat{\theta}$ as per Equation (4.5), when the ground truth separator is $\boldsymbol{\mu}$. We also consider the case of $\boldsymbol{\Sigma} = \mathcal{P}_{\boldsymbol{\mu}}^\perp$, viz., the projection matrix onto the space orthogonal to $\boldsymbol{\mu}$. For any classifier $\hat{\theta}$ learned from $\mathcal{T} := \{(\mathbf{x}_j, \tilde{y}_j)\}_{j \in [n]}$, in order to achieve $\inf_{\boldsymbol{\mu} \in \mathbb{S}^{d-1}} \text{acc}(\hat{\theta}; \boldsymbol{\mu}) \geq 1 - \delta$, the condition $n = \Omega\left(\frac{(1-\delta)}{(1-2p)^2} d\right)$ is necessary in our problem setting.*

It is worth mentioning that there is a similar lower bound in Gentile & Helmbold (1998) for a different classification setting. In contrast, Theorem 4.6 is tailored to our setting and moreover, the proof technique is also different and interesting in its own right. Specifically, for the proof of Theorem 4.6, we follow a standard technique in proving minimax lower bounds which is to reduce the problem of interest to an appropriate multi-way hypothesis testing problem; this is accompanied by the application of the conditional version of Fano’s inequality and some ideas from high-dimensional

geometry. This proof is in Appendix D. There are also a few other noteworthy lower bounds (Cai & Wei, 2021; Im & Grigas, 2023; Zhu et al., 2024) in more general settings than ours; naturally, their bounds are weaker than Theorem 4.6. We discuss these bounds in Appendix E.

Remark 4.7 (Minimax optimality of sample complexity). Note that the dependence of the sample complexity on d and p in Corollary 4.4 matches that of the lower bound in Theorem 4.6. Thus, our sample complexity bound in Corollary 4.4 is optimal with respect to d and p .

4.2 Retraining

We first label the training set using $\hat{\theta}_0$. Denote by $\tilde{y}_i = \text{sign}(\langle x_i, \hat{\theta}_0 \rangle)$ the predicted label for sample x_i . We then retrain using these predicted labels. Our learned classifier here is similar to the one in Section 4.1, except that the observed labels are replaced by the predicted labels. Specifically, our retraining classifier model is the following:

$$\hat{\theta}_1 = \frac{1}{n} \sum_i \tilde{y}_i x_i. \quad (4.8)$$

$\hat{\theta}_1$'s predicted label for a sample $x \in \mathbb{R}^d$ is $\text{sign}(\langle x, \hat{\theta}_1 \rangle)$. Our next result bounds the probability of $\hat{\theta}_1$ misclassifying a fixed test point x , with respect to the randomness in the training set \mathcal{T} (similar to Theorem 4.1).

Theorem 4.8 (Retraining). Suppose that $\frac{n}{d} > \frac{4\lambda_{\max}}{\gamma^2(1-2p)}$, $nd > \frac{\gamma^4}{\lambda_{\max}^2}$, and $d \geq 7$. Consider $x \notin \mathcal{T}$ and let y be its true label. We have

$$\mathbb{P}(\text{sign}(\langle x, \hat{\theta}_1 \rangle) \neq y) \leq \alpha_1(x), \text{ where}$$

$$\begin{aligned} \alpha_1(x) := & 2 \exp \left(- \frac{n(1-2q')^2 \langle x, \mu \rangle^2}{64(\langle x, \mu \rangle^2 + \|\Sigma^{1/2}x\|_{\ell_2}^2)} \right) \\ & + 4 \exp \left(- \frac{n}{32} (1-2p)^2 \right) \\ & + \frac{n}{2} \left(\exp \left(- \frac{\gamma^4}{8\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}, \end{aligned}$$

with $q' = \exp \left(- \frac{n(1-2p)\gamma^2}{40\lambda_{\max}} \right)$ and $p' = \left(1 + \frac{3\gamma^4}{8\lambda_{\max}^2 nd} \right) p$.

The proof of Theorem 4.8 is in Appendix F.

Technical challenges: The proof of Theorem 4.8 is especially challenging due to the following reasons:

(i) The source of the predicted labels $\{\tilde{y}_i\}_{i=1}^n$ used by $\hat{\theta}_1$, viz., $\hat{\theta}_0$, is a *non-uniform noisy label provider*, as we discussed after Theorem 4.1.

(ii) More importantly, $\hat{\theta}_0$ depends on all the samples in the training set and hence each predicted label \tilde{y}_i is dependent

on the entire training set. *This dependence of the predicted labels on all the data points makes the analysis highly challenging because standard techniques under independence are not applicable here.*⁶

The high-level proof idea for Theorem 4.8 is that for each sample $x_\ell = y_\ell(1 + u_\ell)\mu + \Sigma^{1/2}z_\ell$ in the training set, we carefully curate a dummy label \tilde{y}'_ℓ that does not depend on $\{z_k\}_{k \neq \ell}$ but depends on all $\{u_k\}_{k \in [n]}$; see (F.3) in Appendix F. We show that with high probability, the labels predicted by $\hat{\theta}_0$ on the training set $\{\tilde{y}_\ell\}_{\ell \in [n]}$ match the corresponding dummy labels $\{\tilde{y}'_\ell\}_{\ell \in [n]}$. We analyze the misclassification error of $\hat{\theta}_1$ under this high probability “good” event with the dummy labels $\{\tilde{y}'_\ell\}_{\ell \in [n]}$, which are more conducive to analysis (because they do not depend on $\{z_k\}_{k \neq \ell}$), and also bound the probability of the “bad” event (i.e., \exists at least one ℓ such that $\tilde{y}_\ell \neq \tilde{y}'_\ell$). We provide an outline of these steps in detail at the beginning of Appendix F.

We will now provide an upper bound on the error (defined in (4.4)) of $\hat{\theta}_1$.

Theorem 4.9 (Retraining’s population error). Suppose the conditions of Theorem 4.8 hold. Then, we have

$$\begin{aligned} \text{err}(\hat{\theta}_1) \leq & 2 \exp \left(- \frac{n(1-2q')^2 \gamma^4}{64(\gamma^4 + 2\lambda_{\max}^2 d)} \right) \\ & + 2e^{-d/8} + 4 \exp \left(- \frac{n}{32} (1-2p)^2 \right) \\ & + \frac{n}{2} \left(\exp \left(- \frac{\gamma^4}{8\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}, \quad (4.9) \end{aligned}$$

with $q' = \exp \left(- \frac{n(1-2p)\gamma^2}{40\lambda_{\max}} \right)$ and $p' = \left(1 + \frac{3\gamma^4}{8\lambda_{\max}^2 nd} \right) p$.

The proof of Theorem 4.9 is in Appendix G. As expected, (4.9) is decreasing in γ , because the classification task becomes easier as the degree of separation increases.

Retraining vs. Initial training. We will now compare the error bounds of the initial training model $\hat{\theta}_0$ (4.6) from Theorem 4.2 with that of the retraining model $\hat{\theta}_1$ (4.9) from Theorem 4.9. Note that the dominant term in the upper bound on $\text{err}(\hat{\theta}_1)$ is the last term in (4.9) which has $(1-2p')$ inside the exponent, and p' approaches p as n and d grow. In contrast, the lower bound on $\text{err}(\hat{\theta}_0)$ in (4.6) has $(1-2p)^2$ inside the exponent. *This is the main observation that indicates retraining can improve model accuracy, and this improvement becomes increasingly significant as $p \rightarrow \frac{1}{2}$.* We will now formalize this observation by providing sufficient conditions under which $\text{err}(\hat{\theta}_1) < \text{err}(\hat{\theta}_0)$ or equivalently, $\text{acc}(\hat{\theta}_1) > \text{acc}(\hat{\theta}_0)$.

Remark 4.10 (When does retraining improve accuracy?). We consider an asymptotic regime where $n, d \rightarrow \infty$. Also,

⁶The analysis of $\hat{\theta}_0$ (Theorem 4.1) is relatively simpler because the given noisy labels $\{\tilde{y}_i\}_{i=1}^n$ are independent across samples.

suppose that $\frac{\lambda_{\max}}{\lambda_{\min}} \lesssim 1$. By comparing (4.6) with (4.9), we obtain the following sufficient conditions. There exists an absolute constant c_0 such that if $p \in (\frac{1}{2} - c_0, \frac{1}{2})$ and

$$\frac{\lambda_{\min}^2 d}{\gamma^4 (1-2p)^2} \log \left(\frac{\lambda_{\min}^2 d}{\gamma^4 (1-2p)^2} \right) \lesssim n \lesssim \frac{\lambda_{\min}^2 d^2}{\gamma^4 (1-2p)^2},$$

then the *accuracy of retraining is **greater** than the accuracy of initial training*.

Some comments regarding Remark 4.10 are in order.

Range of n in Remark 4.10. We believe that the lower bound on n is tight w.r.t. d and p , ignoring log factors. We claim this because $\Omega\left(\frac{d}{(1-2p)^2}\right)$ samples are necessary for $\text{acc}(\hat{\theta}_0) > 1 - p$ as per Theorem 4.6, and we can only expect the retraining classifier $\hat{\theta}_1$ to be better than the initial training classifier $\hat{\theta}_0$ if the accuracy of $\hat{\theta}_0$'s predicted labels (which is used to train $\hat{\theta}_1$) is more than the accuracy of the given labels = $1 - p$. In contrast, the upper bound on n may be an artifact of our analysis. Having said that, the range of $\frac{d^2}{(1-2p)^2} \lesssim n$ is not very interesting because it is far more than the number of learnable parameters = d .

Effect of degree of separation γ . As γ increases, the minimum value of n in Remark 4.10 for which we obtain an improvement with retraining reduces. This is consistent with the high-level insight of Fig. 1, viz., retraining is more beneficial when the separation between the classes is large.

5 Improving Label DP Training with Retraining (RT)

Motivated by our theoretical results in Section 4 which show that retraining (abbreviated as RT henceforth) can improve accuracy in the presence of label noise, we propose to apply our proposals in Section 1, viz., full RT and more importantly, *consensus-based RT* to improve local label DP training because it involves training with noisy labels. Note that this can be done **on top of any label DP mechanism** and that too **at no additional privacy cost** (both the predicted labels and originally provided noisy labels are private). Here we empirically evaluate full and consensus-based RT on four classification datasets (available on TensorFlow) trained with label DP. These include three vision datasets, namely CIFAR-10, CIFAR-100, and DomainNet (Peng et al., 2019), and one language dataset, namely AG News Subset (Zhang et al., 2015). Unless otherwise mentioned, all our empirical results here are with the standard cross-entropy loss. All the results are averaged over three runs. We only provide important experimental details here; the other details can be found in Appendix H.

CIFAR-10/100. We train a ResNet-18 model (from scratch) on CIFAR-10 and CIFAR-100 with label DP. Label DP training is done with the prior-based method of Ghazi et al.

(2021) – specifically, Algorithm 3 with two stages. Our training set consists of 45k examples and we assume access to a validation set with clean labels consisting of 5k examples which we use for deciding when to stop training, setting hyper-parameters, etc.⁷ For CIFAR-10 and CIFAR-100 with three different values of ϵ (DP parameter; see Def. 3.1), we list the test accuracies of the baseline (i.e., the method of Ghazi et al. (2021)), full RT and consensus-based RT in Tables 1 and 2, respectively. Notice that **consensus-based RT is the clear winner**. Also, for the three values of ϵ in Table 1 (CIFAR-10), the size of the consensus set (used in consensus-based RT) is $\sim 31\%$, 55% and 76% , respectively, of the entire training set. The corresponding numbers for Table 2 (CIFAR-100) are $\sim 11\%$, 34% and 56% , respectively. So for small ϵ (high label noise), *consensus-based RT comprehensively outperforms full RT and the baseline with a small fraction of the training set*. Further, in Table 3, we list the accuracies of the predicted labels and the given labels over the entire (training) dataset and the accuracy of the predicted labels (which are the same as the given labels) over the consensus set for CIFAR-10 and CIFAR-100. To summarize, the accuracy of predicted labels over the consensus set is significantly more than the accuracy of predicted and given labels over the entire dataset. This gives us an idea of why consensus-based RT is much better than full RT and baseline, even though the consensus set is smaller than full dataset.

Table 1. **CIFAR-10.** Test set accuracies (mean \pm standard deviation). *Consensus-based RT is better than full RT which is better than the baseline* (initial training).

ϵ	Baseline	Full RT	Consensus-based RT
1	57.78 \pm 1.13	60.07 \pm 0.63	63.84 \pm 0.56
2	79.06 \pm 0.59	81.34 \pm 0.40	83.31 \pm 0.28
3	85.18 \pm 0.50	86.67 \pm 0.28	87.67 \pm 0.28

Table 2. **CIFAR-100.** Test set accuracies (mean \pm standard deviation). Overall, *consensus-based RT is **significantly** better than full RT which is somewhat better than the baseline* (initial training).

ϵ	Baseline	Full RT	Consensus-based RT
3	23.53 \pm 1.01	24.42 \pm 1.22	29.98 \pm 1.11
4	44.53 \pm 0.81	46.99 \pm 0.66	51.30 \pm 0.98
5	55.75 \pm 0.36	56.98 \pm 0.43	59.47 \pm 0.26

One may wonder how an increase in the number of model parameters affects results because of potential overfitting to label noise. So next, we train ResNet-34 (which has

⁷In practice, we do not need access to the validation set. Instead, it can be stored by a secure agent which returns us a private version of the validation accuracy and this will not be too far off from the true validation accuracy when the validation set is large enough. We also show some results without a validation set in Appendix K.

Table 3. **CIFAR-10 (top) and CIFAR-100 (bottom).** Accuracies of predicted labels and given labels over the entire (training) dataset and accuracies of predicted labels over the consensus set. Note that the *accuracy over the consensus set* \gg *accuracy of over the entire dataset* (with both predicted and given labels). This gives us an idea of why consensus-based RT is much better than full RT and baseline, even though the consensus set is smaller than the full dataset ($\sim 31\%$, 55% and 76% of the full dataset for $\epsilon = 1, 2$ and 3 in the case of CIFAR-10, and $\sim 11\%$, 34% and 56% of the full dataset for $\epsilon = 3, 4$ and 5 in the case of CIFAR-100).

CIFAR-10	ϵ	Acc. of <i>predicted</i> labels on <i>full</i> dataset	Acc. of <i>given</i> labels on <i>full</i> dataset	Acc. of <i>predicted</i> labels on <i>consensus</i> set
	1	59.30 ± 0.74	32.61 ± 0.74	76.17 ± 0.15
	2	81.62 ± 0.18	57.11 ± 0.05	92.65 ± 0.22
	3	89.28 ± 0.35	76.73 ± 0.12	95.94 ± 0.23
CIFAR-100	ϵ	Acc. of <i>predicted</i> labels on <i>full</i> dataset	Acc. of <i>given</i> labels on <i>full</i> dataset	Acc. of <i>predicted</i> labels on <i>consensus</i> set
	3	24.90 ± 0.92	22.35 ± 0.41	76.09 ± 0.85
	4	50.85 ± 0.82	46.32 ± 0.34	91.59 ± 1.24
	5	66.51 ± 0.02	68.09 ± 0.33	94.83 ± 0.15

Table 4. **CIFAR-100 w/ ResNet-34.** Test set accuracies (mean \pm standard deviation). Just like Table 2 (ResNet-18), consensus-based RT is the clear winner here. But the performance in case of ResNet-34 is worse than ResNet-18 due to more overfitting to noise because of more parameters; this is expected.

ϵ	Baseline	Full RT	Consensus-based RT
4	37.53 ± 1.58	39.33 ± 1.37	43.87 ± 1.62
5	51.13 ± 0.69	52.33 ± 0.38	55.43 ± 0.42

nearly six times the number of parameters of ResNet-18) on CIFAR-100; the setup is exactly the same as our previous ResNet-18 experiments. Additionally, one may wonder if retraining is useful when a label noise-robust technique is used during initial training, i.e., on top of the method of Ghazi et al. (2021). To that end, we train ResNet-34 on CIFAR-100 with (i) the popular noise-correcting technique of *forward correction* (Patrini et al., 2017) applied to the first stage of Ghazi et al. (2021) (it is not clear how to apply it to the second stage), and (ii) the noise-robust *symmetric cross-entropy (CE) loss function* of Wang et al. (2019) with $\alpha = 0.8, \beta = 0.2$ used instead of the standard CE loss function. We show the results for ResNet-34 without any noise-robust method, forward correction used initially, and with the symmetric CE loss function in Tables 4, 5, and 6, respectively, for $\epsilon = \{4, 5\}$; please see the detailed discussion in the captions. In summary, consensus-based RT is pretty effective even with overfitting by ResNet-34. Also, more importantly, **consensus-based RT is beneficial even after noise-robust initial training.**

DomainNet (<https://www.tensorflow.org/datasets/catalog/domainnet>). This is an image classification dataset much larger than CIFAR where the images belong to one of 345 classes. Here we do linear probing (i.e., fitting a softmax layer) on top of a ResNet-50 pretrained on ImageNet. We reserve 10% of the entire training set for validation and use the rest for training with label DP. Just like the CIFAR experiments, we assume

Table 5. **CIFAR-100 w/ ResNet-34: noise-robust forward correction (Patrini et al., 2017) used initially** (in the baseline). Test set accuracies (mean \pm standard deviation). We see that *retraining (especially, consensus-based RT) yields improvement even after noise-robust initial training*, although the amount of improvement is less compared to Table 4 (no noise-correction).

ϵ	Baseline	Full RT	Consensus-based RT
4	41.80 ± 0.90	42.63 ± 0.58	46.53 ± 0.76
5	53.83 ± 0.83	54.43 ± 0.48	56.60 ± 0.43

Table 6. **CIFAR-100 w/ ResNet-34: noise-robust symmetric CE loss (Wang et al., 2019) used instead of standard CE loss.** Test set accuracies (mean \pm standard deviation). We see that *retraining (especially, consensus-based RT) yields improvement even with a noise-robust loss function.*

ϵ	Baseline	Full RT	Consensus-based RT
4	37.07 ± 2.03	38.17 ± 2.03	43.20 ± 1.77
5	53.10 ± 0.54	53.40 ± 0.33	56.13 ± 0.25

that the validation set comes with clean labels. In Table 7, we show results for $\epsilon = \{3, 4\}$ when label DP training is done with the method of Ghazi et al. (2021) with two stages (same as the CIFAR experiments). Observe that **consensus-based RT is the clear winner**. Further, in Tables 8 and 9, we show the corresponding results when the noise-robust techniques of *forward* and *backward* correction (Patrini et al., 2017) are applied to the first stage of Ghazi et al. (2021) in initial training (it is not clear how to apply forward and backward correction to the second stage). As expected, forward & backward correction lead to better initial model performance (compared to no correction). The main thing to note however is that **consensus-based RT yields significant gains even after noise-robust initial training**, consistent with our earlier results.⁸

⁸It is worth noting that for $\epsilon = 3$, consensus-based RT leads to similar accuracy with and without noise correction.

Table 7. **DomainNet.** Test set accuracies (mean \pm standard deviation). *Consensus-based RT leads to very significant improvement.*

ϵ	Baseline	Full RT	Consensus-based RT
3	23.60 \pm 0.92	29.23 \pm 1.03	36.30 \pm 0.75
4	48.25 \pm 0.05	52.10 \pm 0.10	57.40 \pm 0.20

Table 8. **DomainNet: noise-robust forward correction used initially.** Test set accuracies (mean \pm standard deviation). *Consensus-based RT yields significant improvement even after noise-robust initial training.*

ϵ	Baseline	Full RT	Consensus-based RT
3	31.23 \pm 0.56	33.30 \pm 0.65	36.07 \pm 0.78
4	58.50 \pm 0.08	58.63 \pm 0.12	61.80 \pm 0.08

Table 9. **DomainNet: noise-robust backward correction used initially.** Test set accuracies (mean \pm standard deviation). Once again, *consensus-based RT yields significant improvement even after noise-robust initial training.*

ϵ	Baseline	Full RT	Consensus-based RT
3	30.17 \pm 0.61	31.47 \pm 0.74	35.03 \pm 0.78
4	56.63 \pm 0.37	56.80 \pm 0.37	60.47 \pm 0.46

Finally, we show results on a language dataset.

AG News Subset (https://www.tensorflow.org/datasets/catalog/ag_news_subset). This is a news article classification dataset consisting of 4 categories – world, sports, business or sci/tech. Just like the previous experiments, we keep 10% of the full training set for validation which we assume comes with clean labels, and use the rest for training with label DP. We use the small BERT model available in TensorFlow and the BERT English uncased preprocessor; links to both of these are in Appendix H. We pool the output of the BERT encoder, add a dropout layer with probability = 0.2, followed by a softmax layer. We fine-tune the full model. Here label DP training is done with randomized response. We list the test accuracies of the baseline (i.e., randomized response), full RT and consensus-based RT in Table 10 for three different values of ϵ . Once again, **consensus-based RT is the clear winner**. For the three values of ϵ in Table 10, the size of the consensus set is $\sim 28\%$, 32% and 38% , respectively, of the entire training set. So here, *consensus-based RT appreciably outperforms full RT and baseline with less than two-fifths of the entire training set*. In Table 11 (Appendix I), we list the accuracies of the predicted & given labels over the full dataset and the consensus set; the observations are similar to Table 3 giving us an idea of why consensus-based RT performs the best.

So in summary, *consensus-based retraining is a **straightforward post-processing step** that can be applied on top of a*

Table 10. **AG News Subset.** Test set accuracies (mean \pm standard deviation). *Consensus-based RT is better than full RT which is better than the baseline.*

ϵ	Baseline	Full RT	Consensus-based RT
0.3	54.54 \pm 0.97	60.03 \pm 2.90	65.91 \pm 1.93
0.5	69.21 \pm 0.31	75.63 \pm 1.08	80.95 \pm 1.47
0.8	79.10 \pm 1.43	82.19 \pm 1.54	84.26 \pm 1.03

*base method to **significantly improve** its performance in the presence of noisy labels.*

Additional empirical results in the Appendix. In Appendix J, we show that consensus-based RT outperforms retraining on samples for which the model is the most confident; this is similar to self-training’s method of sample selection in the semi-supervised setting. In Appendix K, we show that RT is beneficial even without a validation set. Finally, going beyond label DP, we show that consensus-based RT is beneficial in the presence of human annotation errors which can be thought of as “real” label noise in Appendix L.

6 Conclusion

In this work, we provided the first theoretical result showing retraining with predicted hard labels can provably increase model accuracy in the presence of label noise. We also showed the efficacy of consensus-based retraining (i.e., retraining on only those samples for which the predicted label matches the given label) in improving local label DP training. We will conclude by discussing some limitations of our work which pave the way for future directions of work. Our theoretical results in this work focused on full retraining. Because consensus-based retraining worked very well empirically, we would like to analyze it theoretically in the future. Also, our theoretical results are under the uniform label noise model. In the future, we would like to analyze retraining under non-uniform label noise models. We also hope to test our ideas on larger-scale models and datasets.

Acknowledgments

Adel Javanmard is supported in part by the NSF Award DMS-2311024, the Sloan fellowship in Mathematics, an Adobe Faculty Research Award, an Amazon Faculty Research Award, and an iORB grant from USC Marshall School of Business. Sujay Sanghavi is partially supported by NSF EnCORE Tripods (2217069) and NSF AI Institute for the Foundations of Machine Learning (2019844). The authors are grateful to anonymous reviewers for their feedback on improving this paper.

Impact Statement

This paper presents work whose goal is to advance the field of machine learning. There are potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Abramowitz, M. and Stegun, I. A. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1968.
- Amini, M.-R., Feofanov, V., Pauletto, L., Devijver, E., and Maximov, Y. Self-training: A survey. *arXiv preprint arXiv:2202.12040*, 2022.
- Badanidiyuru Varadaraja, A., Ghazi, B., Kamath, P., Kumar, R., Leeman, E., Manurangsi, P., Varadarajan, A. V., and Zhang, C. Optimal unbiased randomizers for regression with label differential privacy. *Advances in Neural Information Processing Systems*, 36:60226–60246, 2023.
- Ball, K. et al. An elementary introduction to modern convex geometry. *Flavors of geometry*, 31(1-58):26, 1997.
- Beimel, A., Nissim, K., and Stemmer, U. Private learning and sanitization: Pure vs. approximate differential privacy. In *International Workshop on Approximation Algorithms for Combinatorial Optimization*, pp. 363–378. Springer, 2013.
- Cai, T. T. and Wei, H. Transfer learning for nonparametric classification. *The Annals of Statistics*, 49(1):100–128, 2021.
- Carmon, Y., Raghunathan, A., Schmidt, L., Duchi, J. C., and Liang, P. S. Unlabeled data improves adversarial robustness. *Advances in neural information processing systems*, 32, 2019.
- Chaudhuri, K. and Hsu, D. Sample complexity bounds for differentially private learning. In *Proceedings of the 24th Annual Conference on Learning Theory*, pp. 155–186. JMLR Workshop and Conference Proceedings, 2011.
- Chen, Y., Wei, C., Kumar, A., and Ma, T. Self-training avoids using spurious features under domain shift. *Advances in Neural Information Processing Systems*, 33: 21061–21071, 2020.
- Das, R. and Sanghavi, S. Understanding self-distillation in the presence of label noise. In *International Conference on Machine Learning*, pp. 7102–7140. PMLR, 2023.
- Dong, B., Hou, J., Lu, Y., and Zhang, Z. Distillation \approx early stopping? harvesting dark knowledge utilizing anisotropic information retrieval for overparameterized neural network. *arXiv preprint arXiv:1910.01255*, 2019.
- Furlanello, T., Lipton, Z., Tschannen, M., Itti, L., and Anandkumar, A. Born again neural networks. In *International Conference on Machine Learning*, pp. 1607–1616. PMLR, 2018.
- Gentile, C. and Helmbold, D. P. Improved lower bounds for learning from noisy examples: An information-theoretic approach. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 104–115, 1998.
- Ghazi, B., Golowich, N., Kumar, R., Manurangsi, P., and Zhang, C. Deep learning with label differential privacy. *Advances in neural information processing systems*, 34: 27131–27145, 2021.
- Ghazi, B., Kamath, P., Kumar, R., Leeman, E., Manurangsi, P., Varadarajan, A. V., and Zhang, C. Regression with label differential privacy. *arXiv preprint arXiv:2212.06074*, 2022.
- Ghosh, M. Exponential tail bounds for chisquared random variables. *Journal of Statistical Theory and Practice*, 15 (2):35, 2021.
- Goel, A., Jiao, Y., and Massiah, J. Pars: Pseudo-label aware robust sample selection for learning with noisy labels. *arXiv preprint arXiv:2201.10836*, 2022.
- Han, J., Luo, P., and Wang, X. Deep self-learning from noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5138–5147, 2019.
- Hinton, G., Vinyals, O., Dean, J., et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.
- Im, H. and Grigas, P. Binary classification with instance and label dependent label noise. *arXiv preprint arXiv:2306.03402*, 2023.
- Kumar, A., Ma, T., and Liang, P. Understanding self-training for gradual domain adaptation. In *International conference on machine learning*, pp. 5468–5479. PMLR, 2020.
- Lee, D.-H. et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896. Atlanta, 2013.
- Li, J., Socher, R., and Hoi, S. C. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*, 2020.
- Li, Y., Yang, J., Song, Y., Cao, L., Luo, J., and Li, L.-J. Learning from noisy labels with distillation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1910–1918, 2017.

- Malek Esmaeili, M., Mironov, I., Prasad, K., Shilov, I., and Tramer, F. Antipodes of label differential privacy: Pate and alibi. *Advances in Neural Information Processing Systems*, 34:6934–6945, 2021.
- Mobahi, H., Farajtabar, M., and Bartlett, P. Self-distillation amplifies regularization in hilbert space. *Advances in Neural Information Processing Systems*, 33:3351–3361, 2020.
- Nguyen, D. T., Mummadi, C. K., Ngo, T. P. N., Nguyen, T. H. P., Beggel, L., and Brox, T. Self: Learning to filter noisy labels with self-ensembling. *arXiv preprint arXiv:1910.01842*, 2019.
- Oymak, S. and Gulcu, T. C. Statistical and algorithmic insights for semi-supervised learning with self-training. *arXiv preprint arXiv:2006.11006*, 2020.
- Patrini, G., Rozza, A., Krishna Menon, A., Nock, R., and Qu, L. Making deep neural networks robust to label noise: A loss correction approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1944–1952, 2017.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1406–1415, 2019.
- Raghunathan, A., Xie, S. M., Yang, F., Duchi, J., and Liang, P. Understanding and mitigating the tradeoff between robustness and accuracy. In *International Conference on Machine Learning*, pp. 7909–7919. PMLR, 2020.
- Reed, S., Lee, H., Anguelov, D., Szegedy, C., Erhan, D., and Rabinovich, A. Training deep neural networks on noisy labels with bootstrapping. *arXiv preprint arXiv:1412.6596*, 2014.
- Scarlett, J. and Cevher, V. An introductory guide to fano’s inequality with applications in statistical estimation. *arXiv preprint arXiv:1901.00555*, 2019.
- Scudder, H. Probability of error of some adaptive pattern-recognition machines. *IEEE Transactions on Information Theory*, 11(3):363–371, 1965.
- Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C. A., Cubuk, E. D., Kurakin, A., and Li, C.-L. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *Advances in neural information processing systems*, 33:596–608, 2020.
- Tanaka, D., Ikami, D., Yamasaki, T., and Aizawa, K. Joint optimization framework for learning with noisy labels. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5552–5560, 2018.
- Tang, X., Nasr, M., Mahloujifar, S., Shejwalkar, V., Song, L., Houmansadr, A., and Mittal, P. Machine learning with differentially private labels: Mechanisms and frameworks. *Proceedings on Privacy Enhancing Technologies*, 2022.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Wang, D. and Xu, J. On sparse linear regression in the local differential privacy model. In *International Conference on Machine Learning*, pp. 6628–6637. PMLR, 2019.
- Wang, Y., Ma, X., Chen, Z., Luo, Y., Yi, J., and Bailey, J. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 322–330, 2019.
- Warner, S. L. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- Wei, C., Shen, K., Chen, Y., and Ma, T. Theoretical analysis of self-training with deep networks on unlabeled data. *arXiv preprint arXiv:2010.03622*, 2020.
- Wei, J., Zhu, Z., Cheng, H., Liu, T., Niu, G., and Liu, Y. Learning with noisy labels revisited: A study using real-world human annotations. In *International Conference on Learning Representations*, 2022.
- Yarowsky, D. Unsupervised word sense disambiguation rivaling supervised methods. In *33rd annual meeting of the association for computational linguistics*, pp. 189–196, 1995.
- Zhang, S., Wang, M., Liu, S. L., Chen, P.-Y. C., and Xiong, J. How does unlabeled data improve generalization in self-training? a one-hidden-layer theoretical analysis. In *International Conference on Learning Representations*, 2022.
- Zhang, X., Zhao, J., and LeCun, Y. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28, 2015.
- Zheng, S., Wu, P., Goswami, A., Goswami, M., Metaxas, D., and Chen, C. Error-bounded correction of noisy labels. In *International Conference on Machine Learning*, pp. 11447–11457. PMLR, 2020.
- Zhu, Y., Zhang, J., Gangrade, A., and Scott, C. Label noise: Ignorance is bliss. *Advances in Neural Information Processing Systems*, 37:116575–116616, 2024.

Appendix

Table of Contents

- Appendix [A](#): Problem Setting of Figure 1
- Appendix [B](#): Proof of Theorem 4.1
- Appendix [C](#): Proof of Theorem 4.2
- Appendix [D](#): Proof of Theorem 4.6
- Appendix [E](#): Other Relevant Lower Bounds for Learning with Noisy Labels
- Appendix [F](#): Proof of Theorem 4.8
- Appendix [G](#): Proof of Theorem 4.9
- Appendix [H](#): Remaining Experimental Details
- Appendix [I](#): Accuracy over the Entire Dataset and over the Consensus Set in the Case of AG News Subset
- Appendix [J](#): Consensus-Based Retraining Does Better than Confidence-Based Retraining
- Appendix [K](#): Retraining is Beneficial Even Without a Validation Set
- Appendix [L](#): Beyond Label DP: Evaluating Retraining in the Presence of Human Annotation Errors

A Problem Setting of Figure 1

The setting is exactly the same as the problem setting in Section 4 with $d = 50$, $\boldsymbol{\mu} = \gamma \mathbf{e}_1$, $\boldsymbol{\Sigma} = \mathbf{I} - \mathbf{e}_1 \mathbf{e}_1^\top$, and $p = 0.4$. In the direction of $\boldsymbol{\mu}$, we have $\langle \mathbf{x}, \boldsymbol{\mu} \rangle = y(1 + u) \|\boldsymbol{\mu}\|_{\ell_2}^2 = y(1 + u)\gamma^2$. We consider balanced classes (i.e., $\pi_+ = \pi_- = \frac{1}{2}$) and choose $u \sim \text{Unif}[0, 4]$, the uniform distribution over the interval $[0, 4]$. In Figure 1a, $\gamma^2 = 0.5$ (large separation) and in Figure 1b, $\gamma^2 = 0.3$ (small separation). The number of training samples in each case is 300 and the retraining is done on the same training set on which the model is initially trained. The learned classifiers from initial training and retraining are the same as in Section 4 (i.e., Equations (4.3) and (4.8), respectively). Finally, the test accuracy of both initial training and retraining in Figure 1b is 68%.

B Proof of Theorem 4.1

We first prove the lower bound $\alpha_0(\mathbf{x})$ and then the upper bound $\tilde{\alpha}_0(\mathbf{x})$.

B.1 Lower Bound $\alpha_0(\mathbf{x})$

Let $\xi_i = y_i \hat{y}_i$. So $\xi_i = 1$ with probability $1 - p$ and $\xi_i = -1$ with probability p . Under our data model, for each sample we can write $\mathbf{x}_i = y_i(1 + u_i)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \mathbf{z}_i$ with $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$. Hence,

$$n\hat{\boldsymbol{\theta}}_0 = \sum_{i \in [n]} \hat{y}_i \mathbf{x}_i = \left(\sum_{i \in [n]} \xi_i (1 + u_i) \right) \boldsymbol{\mu} + \sum_{i \in [n]} \hat{y}_i \boldsymbol{\Sigma}^{1/2} \mathbf{z}_i.$$

We also note that

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\boldsymbol{\theta}}_0 \rangle) \neq y) = \mathbb{P}(y \langle \mathbf{x}, \hat{\boldsymbol{\theta}}_0 \rangle \leq 0) = \mathbb{P}\left(\left(\sum_{i \in [n]} \xi_i (1 + u_i) \right) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle + \langle \mathbf{x}, \sum_{i \in [n]} y \hat{y}_i \boldsymbol{\Sigma}^{1/2} \mathbf{z}_i \rangle \leq 0 \right)$$

Define $\tilde{\mathbf{z}} := \sum_{i \in [n]} y \hat{y}_i \boldsymbol{\Sigma}^{1/2} \mathbf{z}_i$. Conditioning on $\{\hat{y}_i\}_{i \in [n]}$ we have $\tilde{\mathbf{z}} \sim \mathcal{N}(\mathbf{0}, n\boldsymbol{\Sigma})$, and so $\langle \mathbf{x}, \tilde{\mathbf{z}} \rangle \sim \mathcal{N}(\mathbf{0}, n \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)$.

We write

$$\mathbb{P}\left(\left(\sum_{i \in [n]} \xi_i (1 + u_i) \right) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle + \langle \mathbf{x}, \tilde{\mathbf{z}} \rangle \leq 0 \mid \{\xi_i, y_i, u_i\}_{i \in [n]} \right) = \Phi\left(-\frac{(\sum_{i \in [n]} \xi_i (1 + u_i)) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right),$$

where $\Phi(t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^t e^{-s^2/2} ds$ denotes the cdf of standard normal distribution.

Combining the previous two equations, we arrive at

$$\begin{aligned} \mathbb{P}(\text{sign}(\langle \mathbf{x}, \hat{\boldsymbol{\theta}}_0 \rangle) \neq y) &= \mathbb{E} \left[\Phi \left(-\frac{(\sum_{i \in [n]} \xi_i (1 + u_i)) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right) \right] \\ &\geq \mathbb{E} \left[\Phi \left(-\frac{|\sum_{i \in [n]} \xi_i (1 + u_i)| |\langle \mathbf{x}, \boldsymbol{\mu} \rangle|}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right) \right] \\ &\geq \Phi \left(-\frac{\mathbb{E}[|\sum_{i \in [n]} \xi_i (1 + u_i)|] |\langle \mathbf{x}, \boldsymbol{\mu} \rangle|}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right) \end{aligned} \tag{B.1}$$

where the expectation is with respect to $\{\xi_i, u_i\}_{i \in [n]}$ and the last step follows by Jensen's inequality and the fact that $\Phi(t)$ is a convex function for $t \leq 0$.

We next note that by Cauchy–Schwarz inequality,

$$\begin{aligned}
 \mathbb{E} \left[\left| \sum_i \xi_i (1 + u_i) \right| \right] &\leq \sqrt{\mathbb{E}[(\sum_i \xi_i (1 + u_i))^2]} \\
 &= \sqrt{\sum_{i,j} \mathbb{E}[\xi_i \xi_j (1 + u_i)(1 + u_j)]} \\
 &= \sqrt{\sum_i \mathbb{E}[(1 + u_i)^2] + \sum_{i \neq j} \mathbb{E}[\xi_i] \mathbb{E}[\xi_j] \mathbb{E}[1 + u_i] \mathbb{E}[1 + u_j]} \\
 &= \sqrt{5n + 4n(n-1)(1-2p)^2} \leq \sqrt{5n} + \sqrt{5n}(1-2p),
 \end{aligned}$$

where we used that ξ_i, ξ_j are independent for $i \neq j$, and independent from u_i 's. In addition, $\mathbb{E}[\xi_i] = \mathbb{E}[y_i \widehat{y}_i] = \mathbb{P}(y_i = \widehat{y}_i) - \mathbb{P}(y_i \neq \widehat{y}_i) = 1 - 2p$. Also $\mathbb{E}[u_i] \leq 1$ and $\mathbb{E}[u_i^2] \leq 2$, since u_i has unit sub-gaussian norm. Hence, we get

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_0 \rangle) \neq y) \geq \Phi \left(-\frac{\sqrt{5}(1 + \sqrt{n}(1-2p)) |\langle \mathbf{x}, \boldsymbol{\mu} \rangle|}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right). \quad (\text{B.2})$$

We next use a classical lower bound on Φ . Let $\Phi^c(t) = 1 - \Phi(t) = \Phi(-t)$. Using Equation (7.1.13) of [Abramowitz & Stegun \(1968\)](#), we have for any $t > 0$:

$$\Phi^c(t) > \sqrt{\frac{2}{\pi}} \left(\frac{e^{-\frac{t^2}{2}}}{t + \sqrt{t^2 + 4}} \right).$$

Since $t + \sqrt{t^2 + 4} < 2(t+1)$ for $t > 0$, we get:

$$\Phi^c(t) > \frac{1}{\sqrt{2\pi}} \left(\frac{e^{-\frac{t^2}{2}}}{t+1} \right) > \frac{1}{2\sqrt{2\pi}} \exp(-t^2), \quad (\text{B.3})$$

where we used that $t+1 \leq t^2 + 2 \leq 2e^{t^2/2}$. Combining (B.3) and (B.2) we get

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_0 \rangle) \neq y) \geq \frac{1}{2\sqrt{2\pi}} \exp \left(-\frac{5(1 + \sqrt{n}(1-2p))^2 \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2} \right),$$

which completes the derivation of $\alpha_0(\mathbf{x})$.

B.2 Upper Bound $\widetilde{\alpha}_0(\mathbf{x})$

We continue from (B.1), which reads

$$\mathbb{P}(\text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_0 \rangle) \neq y) = \mathbb{E} \left[\Phi \left(-\frac{(\sum_{i \in [n]} \xi_i (1 + u_i)) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \right) \right], \quad (\text{B.4})$$

where the expectation is with respect to the randomness in training data, namely u_i and ξ_i for $i \in [n]$, while the test point (\mathbf{x}, y) is fixed.

Note that under our data model, $\langle y \mathbf{x}, \boldsymbol{\mu} \rangle = (1 + u) \|\boldsymbol{\mu}\|_{\ell_2}^2 = (1 + u) \gamma^2 > 0$ since u is a non-negative random variable. We next define the following event:

$$\mathcal{E}_0 := \left\{ \sum_{i \in [n]} \xi_i (1 + u_i) \geq n(1 - 2p)/2 \right\}. \quad (\text{B.5})$$

Recall that $\mathbb{E}[\xi_i] = 1 - 2p$ and by the Hoeffding-type inequality for sub-gaussian random variables,

$$\mathbb{P} \left(\left| \sum_i \xi_i (1 + u_i) - n(1 - 2p)(1 + \mathbb{E}[u_i]) \right| \geq t \right) \leq 2 \exp \left(-\frac{t^2}{8n} \right),$$

(Note that $\|\xi_i(1+u_i)\|_{\psi_2} \leq 1 + \|u_i\|_{\psi_2} = 2$.) By choosing $t = n(1-2p)/2$, and using the fact that $u_i > 0$ for all i we get

$$\mathbb{P}\left(\sum_i \xi_i(1+u_i) < n(1-2p)/2\right) \leq 2 \exp\left(-\frac{n}{32}(1-2p)^2\right).$$

Hence, $\mathbb{P}(\mathcal{E}_0^c) \leq 2 \exp\left(-\frac{n}{32}(1-2p)^2\right)$.

We then have

$$\begin{aligned} \mathbb{E}\left[\Phi\left(-\frac{(\sum_{i \in [n]} \xi_i(1+u_i))\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right)\right] &= \mathbb{E}\left[\Phi\left(-\frac{(\sum_{i \in [n]} \xi_i(1+u_i))\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right)(\mathbb{1}_{\mathcal{E}_0} + \mathbb{1}_{\mathcal{E}_0^c})\right] \\ &\leq \mathbb{E}\left[\Phi\left(-\frac{n(1-2p)\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{2\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right)\mathbb{1}_{\mathcal{E}_0}\right] + \mathbb{P}(\mathcal{E}_0^c) \\ &= \Phi\left(-\frac{n(1-2p)\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{2\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right) \mathbb{P}(\mathcal{E}_0) + \mathbb{P}(\mathcal{E}_0^c) \\ &\leq \Phi\left(-\frac{n(1-2p)\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{2\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right) + \mathbb{P}(\mathcal{E}_0^c), \end{aligned}$$

where in the first inequality we used the observation that $\langle y\mathbf{x}, \boldsymbol{\mu} \rangle \geq 0$ as explained above, the definition of \mathcal{E}_0 and that $\Phi(-z)$ is a decreasing function in z .

We next recall the tail bound of normal distribution in Equation (7.1.13) of [Abramowitz & Stegun \(1968\)](#):

$$\Phi^c(t) \leq \sqrt{\frac{2}{\pi}} \frac{e^{-t^2/2}}{t + \sqrt{t^2 + \frac{8}{\pi}}} \leq \frac{1}{2} e^{-t^2/2},$$

for all $t \geq 0$, where $\Phi^c(t) = 1 - \Phi(t) = \Phi(-t)$ is the complementary CDF of normal variables. Using this bound and the bound we derived on $\mathbb{P}(\mathcal{E}_0^c)$ we get

$$\mathbb{E}\left[\Phi\left(-\frac{(\sum_{i \in [n]} \xi_i(1+u_i))\langle y\mathbf{x}, \boldsymbol{\mu} \rangle}{\sqrt{n} \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}}\right)\right] \leq \frac{1}{2} \exp\left(-\frac{n(1-2p)^2\langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{8 \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2}\right) + 2 \exp\left(-\frac{n}{32}(1-2p)^2\right).$$

By invoking (B.4), we obtain the desired bound $\tilde{\alpha}_0(\mathbf{x})$.

C Proof of Theorem 4.2

We will first derive the lower bound on $\text{err}(\hat{\boldsymbol{\theta}}_0)$. Using Theorem 4.1 (and because the training set \mathcal{T} has measure zero), we have

$$\begin{aligned} \text{err}(\hat{\boldsymbol{\theta}}_0) &\geq \mathbb{E}[\alpha_0(\mathbf{x})] \\ &= \frac{1}{2\sqrt{2\pi}} \mathbb{E}\left[\exp\left(-\frac{5(1+\sqrt{n}(1-2p))^2\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2}\right)\right] \\ &\geq \frac{1}{2\sqrt{2\pi}} \mathbb{E}\left[\exp\left(-\frac{5(1+\sqrt{n}(1-2p))^2\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2}\right) \middle| \frac{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2} \leq v^2\right] \mathbb{P}\left(\frac{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2} \leq v^2\right) \\ &\geq \frac{1}{2\sqrt{2\pi}} \exp(-5(1+\sqrt{n}(1-2p))^2 v^2) \mathbb{P}\left(\frac{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2} \leq v^2\right), \end{aligned} \tag{C.1}$$

for any value of $v > 0$. We next lower bound the probability on the right-hand side. We have

$$\begin{aligned}
 \mathbb{P}\left(\frac{|\langle \mathbf{x}, \boldsymbol{\mu} \rangle|}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}} \leq v\right) &= \mathbb{P}\left(\frac{|y(1+u)\gamma^2|}{\|\boldsymbol{\Sigma} \mathbf{z}\|_{\ell_2}} \leq v\right) \\
 &\geq \mathbb{P}\left(\frac{(1+u)\gamma^2}{\lambda_{\min} \|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2}} \leq v\right) \\
 &= \mathbb{P}\left(\frac{(1+u)\gamma^2}{\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2}} \leq v\lambda_{\min}\right), \tag{C.2}
 \end{aligned}$$

where the first step holds since under our data model we have $\langle \mathbf{x}, \boldsymbol{\mu} \rangle = y(1+u) \|\boldsymbol{\mu}\|_{\ell_2}^2$ and $\boldsymbol{\Sigma}^{1/2} \mathbf{x} = y(1+u) \boldsymbol{\Sigma}^{1/2} \boldsymbol{\mu} + \boldsymbol{\Sigma} \mathbf{z} = \boldsymbol{\Sigma} \mathbf{z}$, given that $\boldsymbol{\Sigma}^{1/2} \boldsymbol{\mu} = 0$. Here $\mathcal{P}_{\boldsymbol{\mu}}^{\perp}$ is the projection onto the orthogonal space of $\boldsymbol{\mu}$.

Continuing from (C.2) we write

$$\begin{aligned}
 \mathbb{P}\left(\frac{(1+u)\gamma^2}{\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2}} \leq v\lambda_{\min}\right) &\geq \mathbb{P}\left((1+u)\gamma^2 \leq v\lambda_{\min} \sqrt{d/2}, \|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2} \geq \sqrt{d/2}\right) \\
 &= \mathbb{P}\left((1+u)\gamma^2 \leq v\lambda_{\min} \sqrt{d/2}\right) \mathbb{P}\left(\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2} \geq \sqrt{d/2}\right), \tag{C.3}
 \end{aligned}$$

given that \mathbf{z} and u are independent. Next note that $\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}$ is distributed as a Gaussian vector in a $(d-1)$ -dimensional space, and therefore $\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2}^2$ follows the chi-squared distribution with $(d-1)$ degrees of freedom. Using the tail bound of the chi-squared distribution (see Theorem 4 of Ghosh (2021)), we get:

$$\mathbb{P}\left(\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2}^2 < d/2\right) \leq e^{-d/16} \implies \mathbb{P}\left(\|\mathcal{P}_{\boldsymbol{\mu}}^{\perp} \mathbf{z}\|_{\ell_2} \geq \sqrt{d/2}\right) \geq 1 - e^{-d/16}. \tag{C.4}$$

We next lower bound the other term on the right-hand side of (C.3). By Markov's inequality for any non-negative random variable X , we have $\mathbb{P}(X \leq 2\mathbb{E}[X]) \geq 1/2$. Also $\mathbb{E}[(1+u)\gamma^2] \leq 2\gamma^2$, since $\mathbb{E}[u] \leq 1$ and therefore, by choosing $v = \frac{4\sqrt{2}\gamma^2}{\lambda_{\min}\sqrt{d}}$, we get

$$\mathbb{P}\left((1+u)\gamma^2 \leq v\lambda_{\min} \sqrt{d/2}\right) \geq 1/2. \tag{C.5}$$

Combining (C.4), (C.5), (C.3) and (C.2) with (C.1) and plugging in our choice of $v = \frac{4\sqrt{2}\gamma^2}{\lambda_{\min}\sqrt{d}}$, we arrive at

$$\text{err}(\widehat{\boldsymbol{\theta}}_0) \geq \frac{1}{4\sqrt{2\pi}} \exp\left(-160(1+\sqrt{n}(1-2p))^2 \frac{\gamma^4}{\lambda_{\min}^2 d}\right) (1 - e^{-d/16}).$$

Let us now derive the upper bound on $\text{err}(\widehat{\boldsymbol{\theta}}_0)$. Using Theorem 4.1 (and because the training set \mathcal{T} has measure zero), we have

$$\begin{aligned}
 \text{err}(\widehat{\boldsymbol{\theta}}_0) &\leq \mathbb{E}[\widetilde{\alpha}_0(\mathbf{x})] \\
 &= \frac{1}{2} \mathbb{E}\left[\exp\left(-\frac{n(1-2p)^2 \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{8 \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2}\right)\right] + 2 \exp\left(-\frac{n}{32}(1-2p)^2\right). \tag{C.6}
 \end{aligned}$$

Under our data model, $\mathbf{x} = y(1+u)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \mathbf{z}$. Consider the probabilistic event $\mathcal{E}_2 := \{\mathbf{z} : \|\mathbf{z}\|_{\ell_2}^2 \leq 2d\}$. For every $\eta > 0$, we have

$$\begin{aligned}
 \mathbb{E}\left[\exp\left(-\frac{\eta \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2}\right)\right] &= \mathbb{E}\left[\exp\left(-\frac{\eta \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2}\right) (\mathbb{1}_{\mathcal{E}_2} + \mathbb{1}_{\mathcal{E}_2^c})\right] \\
 &\leq \exp\left(-\frac{\eta \gamma^4}{2\lambda_{\max} d}\right) + \mathbb{P}(\mathcal{E}_2^c),
 \end{aligned}$$

where the inequality holds because $\langle \mathbf{x}, \boldsymbol{\mu} \rangle = y(1+u)\gamma^2$ due to which $|\langle \mathbf{x}, \boldsymbol{\mu} \rangle| \geq \gamma^2$ (given that u is a non-negative random variable) and because

$$\|\Sigma^{1/2} \mathbf{x}\|_{\ell_2}^2 = \|\Sigma \mathbf{z}\|_{\ell_2}^2 \leq \lambda_{\max}^2 \|\mathbf{z}\|_{\ell_2}^2 \leq 2\lambda_{\max}^2 d,$$

when event \mathcal{E}_2 happens. Further, $\|\mathbf{z}\|_{\ell_2}^2 \sim \chi_d^2$ and so using tail bounds for χ_d^2 (see Theorem 3 of Ghosh (2021)), we have $\mathbb{P}(\mathcal{E}_2^c) \leq e^{-d/8}$. Putting things together, we obtain

$$\mathbb{E} \left[\exp \left(-\frac{\eta \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\|\Sigma^{1/2} \mathbf{x}\|_{\ell_2}^2} \right) \right] \leq \exp \left(-\frac{\eta \gamma^4}{2\lambda_{\max}^2 d} \right) + e^{-d/8}.$$

Setting $\eta = n(1-2p)^2/8$ and using this in (C.6), we get

$$\text{err}(\widehat{\boldsymbol{\theta}}_0) \leq \frac{1}{2} \exp \left(-\frac{n(1-2p)^2 \gamma^4}{16\lambda_{\max}^2 d} \right) + e^{-d/8} + 2 \exp \left(-\frac{n}{32} (1-2p)^2 \right).$$

D Proof of Theorem 4.6

Proof. Since we are interested with dependence of sample size with respect to d and p , we assume $\gamma := \|\boldsymbol{\mu}\|_{\ell_2} = 1$ and $\Sigma = \mathbf{I} - \boldsymbol{\mu}\boldsymbol{\mu}^\top$ the projection onto the orthogonal space of $\boldsymbol{\mu}$. Also note that without loss of generality, we can assume that our estimator $\widehat{\boldsymbol{\theta}}$ has unit norm, because its norm does not affect the sign of $\langle \mathbf{x}, \widehat{\boldsymbol{\theta}} \rangle$. This way, $\boldsymbol{\mu}$ and $\widehat{\boldsymbol{\theta}}$ both belong to \mathbb{S}^{d-1} .

We first follow a standard argument to “reduce” the classification problem to a multi-way hypothesis testing problem. Let $\rho \in (0, 1)$ be an arbitrary but fixed value which we can choose. We define a ρ -packing of \mathbb{S}^{d-1} as a set $\mathcal{M} = \{\boldsymbol{\mu}_1, \dots, \boldsymbol{\mu}_M\} \subset \mathbb{S}^{d-1}$ such that $\langle \boldsymbol{\mu}_l, \boldsymbol{\mu}_k \rangle \leq \rho$ for $l \neq k$. Also define the ρ -packing number of \mathbb{S}^{d-1} as

$$M(\rho, \mathbb{S}^{d-1}) := \sup \{M \in \mathbb{N} : \text{there exists a } \rho\text{-packing } \mathcal{M} \text{ of } \mathbb{S}^{d-1} \text{ with size } M\}. \quad (\text{D.1})$$

For convenience, we define the misclassification error of a classifier $\widehat{\boldsymbol{\theta}} \in \mathbb{S}^{d-1}$ when the ground truth separator is $\boldsymbol{\mu}$ as $\text{err}(\widehat{\boldsymbol{\theta}}; \boldsymbol{\mu}) = 1 - \text{acc}(\widehat{\boldsymbol{\theta}}; \boldsymbol{\mu})$. As per our condition,

$$\delta \geq \sup_{\boldsymbol{\mu} \in \mathbb{S}^{d-1}} \text{err}(\widehat{\boldsymbol{\theta}}; \boldsymbol{\mu}) \geq \sup_{\boldsymbol{\mu} \in \mathcal{M}} \text{err}(\widehat{\boldsymbol{\theta}}; \boldsymbol{\mu}). \quad (\text{D.2})$$

In order to further lower bound the right hand side, we let I be a random variable uniformly distributed on the hypothesis set $\{1, 2, \dots, M\}$ and consider the case of $\boldsymbol{\mu} = \boldsymbol{\mu}_I$. We also define \widehat{I} as the index of the element in \mathcal{M} with maximum inner product with $\widehat{\boldsymbol{\theta}}$ (it does not matter how we break ties). Under our data model we have $\langle \mathbf{x}, \boldsymbol{\mu}_I \rangle = y(1+u)\|\boldsymbol{\mu}_I\|_{\ell_2}^2 = y(1+u)$ and since u is a non-negative random variable, $y = \text{sign}(\langle \mathbf{x}, \boldsymbol{\mu}_I \rangle)$. We then have

$$\begin{aligned} \sup_{\boldsymbol{\mu} \in \mathcal{M}} \text{err}(\widehat{\boldsymbol{\theta}}; \boldsymbol{\mu}) &\geq \max_{i \in [M]} \mathbb{P} \left(\text{sign}(\langle \mathbf{x}, \boldsymbol{\mu}_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}} \rangle) \middle| I = i \right) \\ &\geq \frac{1}{M} \sum_{i=1}^M \mathbb{P} \left(\text{sign}(\langle \mathbf{x}, \boldsymbol{\mu}_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}} \rangle) \middle| I = i \right) \\ &\geq \Phi \left(-2\sqrt{\frac{1+\rho}{1-\rho}} \right) \frac{1}{M} \sum_{i=1}^M \mathbb{P}(\widehat{I} \neq i | I = i) \end{aligned} \quad (\text{D.3})$$

$$= \Phi \left(-2\sqrt{\frac{1+\rho}{1-\rho}} \right) \mathbb{P}(\widehat{I} \neq I), \quad (\text{D.4})$$

with Φ denoting the CDF of a standard normal variable. Equation (D.3) above follows from the lemma below.

Lemma D.1. *For any $i \in [M]$, we have*

$$\mathbb{P} \left(\text{sign}(\langle \mathbf{x}, \boldsymbol{\mu}_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}} \rangle) \middle| I = i \right) \geq \Phi \left(-2\sqrt{\frac{1+\rho}{1-\rho}} \right) \mathbb{P}(\widehat{I} \neq i | I = i).$$

Combining (D.2) and (D.4), we obtain that

$$\delta_0 := \frac{\delta}{\Phi\left(-2\sqrt{\frac{1+\rho}{1-\rho}}\right)} \geq \mathbb{P}(\widehat{I} \neq I). \quad (\text{D.5})$$

Next recall the set $\mathcal{T} := \{(\mathbf{x}_j, \widehat{\mathbf{y}}_j)\}_{j \in [n]}$, and let $\mathbf{X} = [\mathbf{x}_1^T, \dots, \mathbf{x}_n^T]^T \in \mathbb{R}^{n \times d}$ and $\widehat{\mathbf{y}} = [\widehat{y}_1, \dots, \widehat{y}_n]^T$. By an application of Fano's inequality, with conditioning on \mathbf{X} (see e.g., Section 2.3 of [Scarlett & Cevher \(2019\)](#)) we have

$$\mathcal{I}(I; \widehat{I} | \mathbf{X}) \geq (1 - \delta_0) \log(M(\rho, \mathbb{S}^{d-1})) - \log 2, \quad (\text{D.6})$$

where $\mathcal{I}(I; \widehat{I} | \mathbf{X})$ represents the conditional mutual information between I and \widehat{I} . Using the fact that $I \rightarrow \widehat{\mathbf{y}} \rightarrow \widehat{I}$ forms a Markov chain conditioned on \mathbf{X} , and by an application of the data processing inequality we have:

$$\mathcal{I}(I; \widehat{I} | \mathbf{X}) \leq \mathcal{I}(I; \widehat{\mathbf{y}} | \mathbf{X}) \quad (\text{D.7})$$

We will now upper bound $\mathcal{I}(I; \widehat{\mathbf{y}} | \mathbf{X})$. Let $w_j := \frac{1}{2}(\text{sign}(\langle \mathbf{x}_j, \boldsymbol{\mu}_I \rangle) + 1)$ be the 0-1 version of the actual label of \mathbf{x}_j , viz., $\text{sign}(\langle \mathbf{x}_j, \boldsymbol{\mu}_I \rangle)$. As per our setting, we have $\widehat{y}_j = 2(w_j \oplus z_j) - 1$ where $z_j \sim \text{Bernoulli}(p)$ and \oplus denotes modulo-2 addition. Since the noise variables z_j are independent and \widehat{y}_j depends on (I, \mathbf{X}) only through $w_j = \frac{1}{2}(\text{sign}(\langle \mathbf{x}_j, \boldsymbol{\mu}_I \rangle) + 1)$, by using the tensorization property of the mutual information (see e.g., Lemma 2, part (iii) of [Scarlett & Cevher \(2019\)](#)), we have

$$\mathcal{I}(I; \mathbf{y} | \mathbf{X}) \leq \sum_{j=1}^n \mathcal{I}(w_j; \widehat{y}_j) \leq n(\log 2 - H_2(p)), \quad (\text{D.8})$$

where the second inequality follows since \widehat{y}_j is generated by passing w_j through a binary symmetric channel, which has capacity $\log 2 - H_2(p)$ with $H_2(p) := -p \log p - (1-p) \log(1-p)$ denoting the binary entropy function.

We next use the lemma below to further upper bound the right-hand side of (D.8).

Lemma D.2. *For a discrete probability distribution, consider the entropy function given by*

$$H(p_1, \dots, p_k) = \sum_{i=1}^k p_i \log(1/p_i).$$

We have the following bound:

$$H(p_1, \dots, p_k) \geq \log k - k \sum_{i=1}^k (p_i - 1/k)^2.$$

Using Lemma D.2 with $k = 2$ we obtain $\log 2 - H_2(p) \leq 4(p - 1/2)^2 = (1 - 2p)^2$, which along with (D.8), (D.7) and (D.6) gives

$$\frac{n(1 - 2p)^2 + \log 2}{1 - \delta_0} \geq \log(M(\rho, \mathbb{S}^{d-1})). \quad (\text{D.9})$$

In our next lemma, we lower bound $M(\rho, \mathbb{S}^{d-1})$.

Lemma D.3. *Recall the definition of ρ -packing number of \mathbb{S}^{d-1} given by (D.1). We have the following bound:*

$$M(\rho, \mathbb{S}^{d-1}) \geq \exp\left(\frac{d\rho^2}{2}\right).$$

Using Lemma D.3 along with (D.9), we obtain the following lower bound on the sample complexity:

$$n \geq \frac{\frac{\rho^2}{2}(1 - \delta_0)d - \log 2}{(1 - 2p)^2}, \text{ with } \delta_0 = \frac{\delta}{\Phi\left(-2\sqrt{\frac{1+\rho}{1-\rho}}\right)}.$$

Note that $\rho \in (0, 1)$ can be set arbitrarily. Since our claim is on the order of n , the specific value of ρ can be chosen based on δ to ensure $\frac{\rho^2}{2}(1 - \delta_0)$ is of the order of $(1 - \delta)$. This finishes the proof of Theorem 4.6. \square

We will now prove Lemmas D.1, D.2 and D.3.

Proof of Lemma D.1. Let us consider the event $\widehat{I} \neq i$, given that $I = i$. We note that if $\widehat{I} \neq I$, then $\langle \widehat{\theta}, \mu_I \rangle \leq \sqrt{\frac{1+\rho}{2}}$ or equivalently, the angle between $\widehat{\theta}$ and μ_I is $\geq b := \cos^{-1}\left(\sqrt{\frac{1+\rho}{2}}\right)$. Otherwise, by definition of \widehat{I} we have $\langle \widehat{\theta}, \mu_{\widehat{I}} \rangle \geq \langle \widehat{\theta}, \mu_I \rangle > \sqrt{\frac{1+\rho}{2}}$, and therefore the angle between $\mu_{\widehat{I}}$ and μ_I is $< 2b$. Noting that $\cos(2b) = 2\cos^2(b) - 1 = \rho$, we would then have $\langle \mu_{\widehat{I}}, \mu_I \rangle > \rho$, which is a contradiction since \mathcal{M} forms a ρ -packing. We proceed by writing

$$\mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\theta} \rangle) \middle| I = i\right) \geq \mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\theta} \rangle) \middle| I = i, \widehat{I} \neq i\right) \mathbb{P}(\widehat{I} \neq i | I = i).$$

As discussed above, on the event that $I = i, \widehat{I} \neq i$, we have $\theta := \langle \widehat{\theta}, \mu_i \rangle \leq \sqrt{\frac{1+\rho}{2}}$. Consider the decomposition $\mathbf{x} = \langle \mathbf{x}, \mu_i \rangle \mu_i + \mathcal{P}_{\mu_i}^\perp \mathbf{x}$. We then have

$$\mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\theta} \rangle) \middle| I = i, \widehat{I} \neq i\right) = \mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_i \rangle) (\langle \mathbf{x}, \mu_i \rangle \theta + \mathcal{P}_{\mu_i}^\perp \mathbf{x}, \mathcal{P}_{\mu_i}^\perp \widehat{\theta}) \leq 0 \middle| I = i, \widehat{I} \neq i\right).$$

Note that \mathbf{x} is a test data point, independent of the training data \mathcal{T} and so it is independent of $\widehat{\theta}$. In addition, under our data model, $\langle \mathbf{x}, \mu_i \rangle$ is independent of $\mathcal{P}_{\mu_i}^\perp \mathbf{x}$. Hence, $\text{sign}(\langle \mathbf{x}, \mu_i \rangle) (\langle \mathbf{x}, \mu_i \rangle \theta + \mathcal{P}_{\mu_i}^\perp \mathbf{x}, \mathcal{P}_{\mu_i}^\perp \widehat{\theta}) \sim \mathcal{N}(0, \|\mathcal{P}_{\mu_i}^\perp \widehat{\theta}\|_{\ell_2}^2)$. Given that $\|\widehat{\theta}\|_{\ell_2} = 1$ we also have $\|\mathcal{P}_{\mu_i}^\perp \widehat{\theta}\|_{\ell_2}^2 = 1 - \langle \widehat{\theta}, \mu_i \rangle^2 = 1 - \theta^2$. In short, we can write

$$\text{sign}(\langle \mathbf{x}, \mu_i \rangle) (\langle \mathbf{x}, \mu_i \rangle \theta + \mathcal{P}_{\mu_i}^\perp \mathbf{x}, \mathcal{P}_{\mu_i}^\perp \widehat{\theta}) = \sqrt{1 - \theta^2} Z, \quad Z \sim \mathcal{N}(0, 1).$$

Using this characterization, we proceed by writing

$$\begin{aligned} \mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_I \rangle) \neq \text{sign}(\langle \mathbf{x}, \widehat{\theta} \rangle) \middle| I = i, \widehat{I} \neq i\right) &= \mathbb{P}\left(\text{sign}(\langle \mathbf{x}, \mu_i \rangle) (\langle \mathbf{x}, \mu_i \rangle \theta + \sqrt{1 - \theta^2} Z \leq 0\right) \\ &= \mathbb{P}\left(Z \leq -\frac{|\langle \mathbf{x}, \mu_i \rangle| \theta}{\sqrt{1 - \theta^2}}\right) \\ &= 1 - \mathbb{E}\left[\Phi\left(\frac{\theta}{\sqrt{1 - \theta^2}} |\langle \mathbf{x}, \mu_i \rangle|\right)\right] \\ &\stackrel{(a)}{\geq} 1 - \mathbb{E}\left[\Phi\left(\sqrt{\frac{1+\rho}{1-\rho}} |\langle \mathbf{x}, \mu_i \rangle|\right)\right] \\ &\stackrel{(b)}{\geq} 1 - \Phi\left(\sqrt{\frac{1+\rho}{1-\rho}} \mathbb{E}[|\langle \mathbf{x}, \mu_i \rangle|]\right) \\ &\stackrel{(c)}{\geq} 1 - \Phi\left(2\sqrt{\frac{1+\rho}{1-\rho}}\right) = \Phi\left(-2\sqrt{\frac{1+\rho}{1-\rho}}\right), \end{aligned}$$

where (a) follows from the fact that $\theta \leq \sqrt{\frac{1+\rho}{2}}$; (b) holds due to Jensen's inequality and concavity of $\Phi(\cdot)$ on the positive values, and (c) holds because under our data model $\mathbb{E}[|\langle \mathbf{x}, \mu_i \rangle|] = \mathbb{E}[y(1+u) \|\mu_i\|_{\ell_2}^2] = \mathbb{E}[1+u] \leq 2$ since $\|u\|_{\psi_2} = 1$. This completes the proof of claim.

Proof of Lemma D.2. Define $q_i = p_i - 1/k$. Note that q_i can be negative, and we have $\sum_{i=1}^k q_i = 0$. We write

$$\begin{aligned}
 H(p_1, \dots, p_k) &= - \sum_{i=1}^k p_i \log p_i \\
 &= - \sum_{i=1}^k (1/k + q_i) \log(1/k + q_i) \\
 &= - \sum_{i=1}^k (1/k + q_i) [\log(1/k) + \log(1 + kq_i)] \\
 &\geq \log k - \sum_{i=1}^k (1/k + q_i) k q_i \\
 &= \log k - k \sum_{i=1}^k q_i^2 \\
 &= \log k - k \sum_{i=1}^k (p_i - 1/k)^2.
 \end{aligned} \tag{D.10}$$

Note that in Equation (D.10) we used the fact that $1 + kq_i \geq 0$ and $\log x \leq x - 1$ for all $x \geq 0$.

This completes the proof of the lemma.

Proof of Lemma D.3. Define a ρ -cover of \mathbb{S}^{d-1} as a set of $\mathcal{V} := \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ such that for any $\boldsymbol{\theta} \in \mathbb{S}^{d-1}$, there exists some \mathbf{v}_i such that $\langle \boldsymbol{\theta}, \mathbf{v}_i \rangle \geq \rho$. The ρ -covering number of \mathbb{S}^{d-1} is

$$N(\rho, \mathbb{S}^{d-1}) := \inf\{N \in \mathbb{N} : \text{there exists a } \rho\text{-cover } \mathcal{V} \text{ of } \mathbb{S}^{d-1} \text{ with size } N\}.$$

By a simple argument we have $M(\rho, \mathbb{S}^{d-1}) \geq N(\rho, \mathbb{S}^{d-1})$. Concretely, we construct a ρ -packing greedily by adding an element at each step which has inner product at most ρ with all the previously selected elements, until it is no longer possible. This means that any point on \mathbb{S}^{d-1} has inner product larger than ρ by some of the elements in the constructed set (otherwise it contradicts its maximality). Hence, we have a set that is both a ρ -cover and a ρ -packing of \mathbb{S}^{d-1} , and by definition it results in $M(\rho, \mathbb{S}^{d-1}) \geq N(\rho, \mathbb{S}^{d-1})$.

We next lower bound $N(\rho, \mathbb{S}^{d-1})$ via a volumetric argument. Let $\mathcal{V} := \{\mathbf{v}_1, \dots, \mathbf{v}_N\}$ be a ρ -cover of \mathbb{S}^{d-1} . For each element $\mathbf{v}_i \in \mathcal{V}$ we consider the cone around it with apex angle $\cos^{-1}(\rho)$. Its intersection with \mathbb{S}^{d-1} defines a spherical cap which we denote by $\mathcal{C}(\mathbf{v}_i, \rho)$. Since \mathcal{V} forms a ρ -cover of \mathbb{S}^{d-1} , we have

$$\text{Vol}(\mathbb{S}^{d-1}) \leq \text{Vol}(\cup_{i=1}^N \mathcal{C}(\mathbf{v}_i, \rho)) \leq \sum_{i=1}^N \text{Vol}(\mathcal{C}(\mathbf{v}_i, \rho)).$$

We next use Lemma 2.2 from Ball et al. (1997) by which we have $\frac{\text{Vol}(\mathcal{C}(\mathbf{v}_i, \rho))}{\text{Vol}(\mathbb{S}^{d-1})} \leq e^{-d\rho^2/2}$. Using this above, we get

$$1 \leq N e^{-d\rho^2/2}$$

for any ρ -cover \mathcal{V} . Thus, we have $N(\rho, \mathbb{S}^{d-1}) \geq \exp(\frac{d\rho^2}{2})$, which completes the proof of the lemma.

E Comparison with Other Lower Bounds for Learning with Noisy Labels

Cai & Wei (2021); Im & Grigas (2023); Zhu et al. (2024) derive some noteworthy lower bounds in much more general settings than ours; naturally, their bounds are weaker than our bound in Theorem 4.6. In the work of Cai & Wei (2021), our setting corresponds to $n_p = 0, n_q = n$. As per Theorem 3.2 of Cai & Wei (2021), the lower bound on the error is effectively⁹ $n^{-O(\frac{1+\alpha}{d})}$. So when $\alpha \ll d$, this lower bound yields a much worse sample complexity than our result in Theorem 4.6. In Im & Grigas (2023), the lower bound on the error (Theorem 2) does not reduce with n , so even if there are infinite samples, we cannot get zero error in the worst case. As for Zhu et al. (2024), their lower bound on the error (Theorem 1) also has a non-diminishing term depending on ϵ . In the special case of $\epsilon = 0$ (or ϵ being small enough), there is an $n^{-1/2}$ dependence but no dependence on the dimension or a related quantity. However, the upper bound in Theorem 2 therein does have a dependence on a VC dimension-like quantity as expected, so their lower bound is probably loose w.r.t. dimension.

⁹This is because $\beta \leq 1$ as per Definition 2 and $\alpha\beta \leq d$ as per the paragraph after Remark 3 (of Cai & Wei (2021)).

F Proof of Theorem 4.8

Similar to the proof of Theorem 4.1, we want to upper bound $\mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0)$.¹⁰ Plugging in $\widehat{\boldsymbol{\theta}}_1$ from (4.8) we have

$$y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle = \frac{y}{n} \langle \mathbf{x}, \sum_i \widetilde{y}_i \mathbf{x}_i \rangle = \frac{y}{n} \langle \mathbf{x}, \sum_i \text{sign}(\langle \mathbf{x}_i, \widehat{\boldsymbol{\theta}}_0 \rangle) \mathbf{x}_i \rangle$$

A major complication is that $\widehat{\boldsymbol{\theta}}_0$ depends on all the data points in the training set. Expanding $\widehat{\boldsymbol{\theta}}_0$, we get

$$\widehat{\boldsymbol{\theta}}_0 = \frac{1}{n} \sum_i \widehat{y}_i \mathbf{x}_i = \frac{1}{n} \sum_i \xi_i y_i \mathbf{x}_i = \frac{1}{n} \left(\left(\sum_i \xi_i (1 + u_i) \right) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \sum_i \xi_i y_i \mathbf{z}_i \right). \quad (\text{F.1})$$

Thus,

$$\widetilde{y}_\ell = \text{sign}(\langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle) = \text{sign} \left(\langle \mathbf{x}_\ell, \left(\sum_i \xi_i (1 + u_i) \right) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \sum_i \xi_i y_i \mathbf{z}_i \rangle \right). \quad (\text{F.2})$$

Here is the outline of our proof:

- For every $\ell \in [n]$ we define the dummy label

$$\widetilde{y}'_\ell = \text{sign} \left(\langle \mathbf{x}_\ell, \left(\sum_i \xi_i (1 + u_i) \right) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_\ell y_\ell \mathbf{z}_\ell \rangle \right). \quad (\text{F.3})$$

Note that \widetilde{y}'_ℓ depends only on \mathbf{z}_ℓ , whereas \widetilde{y}_ℓ depends on all $\{\mathbf{z}_i\}_{i \in [n]}$.

- We define the event $\mathcal{E} := \{\forall \ell : \widetilde{y}'_\ell = \text{sign}(\langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle)\}$.
- We have

$$\begin{aligned} \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0) &= \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0; \mathcal{E}) + \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0; \mathcal{E}^c) \\ &\leq \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0; \mathcal{E}) + \mathbb{P}(\mathcal{E}^c). \end{aligned} \quad (\text{F.4})$$

- We bound each of the term on the right-hand side separately. Note that under the event \mathcal{E} , we have

$$y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle = \frac{y}{n} \langle \mathbf{x}, \sum_\ell \widetilde{y}'_\ell \mathbf{x}_\ell \rangle.$$

We will control the probability of the RHS above being negative using concentration results on the sum of i.i.d. subgaussian random variables (with appropriate conditioning first).

Next we get into details of this proof outline. We start by bounding $\mathbb{P}(\mathcal{E}^c)$.

Lemma F.1. Fix $\ell \in [n]$. Suppose that $nd \geq \frac{\gamma^4}{\lambda_{\max}^2}$. Define $p' := (1 + \frac{3\gamma^4}{8\lambda_{\max}^2 nd})p$. We then have

$$\mathbb{P}(\widetilde{y}'_\ell \neq \text{sign}(\langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle)) = \mathbb{P}(\widetilde{y}'_\ell \langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle < 0) \leq \frac{1}{2} \left(\exp \left(-\frac{\gamma^4}{8\lambda_{\max}^2} (1 - 2p') \frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}.$$

Using Lemma F.1 along with union bounding over $\ell \in [n]$ (note that the events are dependent), we get

$$\mathbb{P}(\mathcal{E}^c) \leq \frac{n}{2} \left(\exp \left(-\frac{\gamma^4}{8\lambda_{\max}^2} (1 - 2p') \frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}, \quad (\text{F.5})$$

where $p' = (1 + \frac{3\gamma^4}{8\lambda_{\max}^2 nd})p$.

¹⁰Note that $\mathbb{P}(\text{sign}(\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle) \neq y) = \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0)$.

Next we note that under the event \mathcal{E} , we have:

$$y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle = \frac{y}{n} \langle \mathbf{x}, \sum_{\ell} \widetilde{\mathbf{y}}_{\ell} \mathbf{x}_{\ell} \rangle = \frac{1}{n} \sum_{\ell \in [n]} y \text{sign} \left(\langle \mathbf{x}_{\ell}, (\sum_i \xi_i (1 + u_i)) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_{\ell} y_{\ell} \mathbf{z}_{\ell} \rangle \right) \langle \mathbf{x}, \mathbf{x}_{\ell} \rangle. \quad (\text{F.6})$$

We define the shorthand $\beta := \sum_i \xi_i (1 + u_i)$ and condition on β . Then (F.6) will be sum of i.i.d. subgaussian variables

$$T_{\ell} := y \text{sign} \left(\langle \mathbf{x}_{\ell}, \beta \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_{\ell} y_{\ell} \mathbf{z}_{\ell} \rangle \right) \langle \mathbf{x}, \mathbf{x}_{\ell} \rangle, \quad \ell \in [n]. \quad (\text{F.7})$$

We continue by first characterizing its expectation.

Lemma F.2. *Suppose that $\frac{\beta \gamma^2}{2\lambda_{\max} d} > 1$ and $d \geq 7$. Then for $\ell \in [n]$, we have*

$$\mathbb{E}[T_{\ell} | \beta] \geq (1 - 2q) \langle y \mathbf{x}, \boldsymbol{\mu} \rangle > 0,$$

where $q := \exp \left(-\frac{\beta \gamma^2}{20\lambda_{\max}} \right)$.

Recall that $\mathbf{x}_{\ell} = y_{\ell} (1 + u_{\ell}) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \mathbf{z}_{\ell}$ and therefore

$$\begin{aligned} \|T_{\ell} - \mathbb{E}[T_{\ell} | \beta]\|_{\psi_2} &\leq 2\|T_{\ell}\|_{\psi_2} = 2\|\langle \mathbf{x}, \mathbf{x}_{\ell} \rangle\|_{\psi_2} \\ &= 2\|\langle \mathbf{x}, y_{\ell} (1 + u_{\ell}) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \mathbf{z}_{\ell} \rangle\|_{\psi_2} \\ &\leq 4|\langle \mathbf{x}, \boldsymbol{\mu} \rangle| + 2\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}, \end{aligned}$$

where we used the assumption $\|u_{\ell}\|_{\psi_2} = 1$ and the fact that $\|\mathbf{z}\|_{\psi_2} = 1$.

Next by using Hoeffding-type inequality for sum of sub-gaussian random variables (see e.g., Proposition 5.1 of [Vershynin \(2010\)](#)) we have for every $t \geq 0$,

$$\mathbb{P} \left(\left| \sum_{\ell} T_{\ell} - \sum_{\ell} \mathbb{E}[T_{\ell} | \beta] \right| \geq t \mid \beta \right) \leq 2 \exp \left(-\frac{t^2}{64n(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right).$$

Therefore,

$$\begin{aligned} \mathbb{P}(\sum_{\ell} T_{\ell} < 0 \mid \beta) &\leq \mathbb{P} \left(\left| \sum_{\ell} T_{\ell} - \sum_{\ell} \mathbb{E}[T_{\ell} | \beta] \right| \geq \sum_{\ell} \mathbb{E}[T_{\ell} | \beta] \mid \beta \right) \\ &\leq 2 \exp \left(-\frac{n(\mathbb{E}[T_{\ell} | \beta])^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \\ &\leq 2 \exp \left(-\frac{n(1 - 2q)^2 \langle y \mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right), \end{aligned} \quad (\text{F.8})$$

where the last step follows from Lemma F.2. Recall that $q = \exp \left(-\frac{\beta \gamma^2}{20\lambda_{\max}} \right)$.

Our next step is to take expectation of the above with respect to $\beta = \sum_i \xi_i (1 + u_i)$. Before proceeding we define the following event:

$$\mathcal{E}_0 := \{\beta \geq n(1 - 2p)/2\}. \quad (\text{F.9})$$

Note that by Hoeffding-type inequality for sub-gaussian random variables and given that $\mathbb{E}[\xi_i] = 1 - 2p$, we have

$$\mathbb{P} \left(\left| \sum_i \xi_i (1 + u_i) - n(1 - 2p)(1 + \mathbb{E}[u_i]) \right| \geq t \right) \leq 2 \exp \left(-\frac{t^2}{8n} \right).$$

By choosing $t = n(1 - 2p)/2$, and using the fact that $u_i > 0$ for all i , we get

$$\mathbb{P} \left(\sum_i \xi_i (1 + u_i) < n(1 - 2p)/2 \right) \leq 2 \exp \left(-\frac{n}{32} (1 - 2p)^2 \right).$$

Hence,

$$\mathbb{P}(\mathcal{E}_0^c) \leq 2 \exp\left(-\frac{n}{32}(1-2p)^2\right).$$

We now continue by taking expectation of (F.8) with respect to β .

$$\begin{aligned} \mathbb{P}(\sum_{\ell} T_{\ell} < 0) &\leq 2 \mathbb{E} \left[\exp \left(-\frac{n(1-2q)^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \right] \\ &= 2 \mathbb{E} \left[\exp \left(-\frac{n(1-2q)^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) (\mathbb{1}_{\mathcal{E}_0} + \mathbb{1}_{\mathcal{E}_0^c}) \right] \\ &\leq 2 \mathbb{E} \left[\exp \left(-\frac{n(1-2q)^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \mathbb{1}_{\mathcal{E}_0} \right] + 2\mathbb{P}(\mathcal{E}_0^c). \end{aligned} \quad (\text{F.10})$$

On the event \mathcal{E}_0 we have

$$\frac{\beta\gamma^2}{2\lambda_{\max}d} \geq \frac{n(1-2p)\gamma^2}{4\lambda_{\max}d} > 1,$$

by our requirement in the theorem statement. Therefore,

$$q = \exp\left(-\frac{\beta\gamma^2}{20\lambda_{\max}}\right) \leq \exp\left(-\frac{n(1-2p)\gamma^2}{40\lambda_{\max}}\right) := q'.$$

Since $1-2q \geq 1-2q' > 0$ this implies that

$$\begin{aligned} \mathbb{E} \left[\exp \left(-\frac{n(1-2q)^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \mathbb{1}_{\mathcal{E}_0} \right] &\leq \mathbb{E} \left[\exp \left(-\frac{n(1-2q')^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \mathbb{1}_{\mathcal{E}_0} \right] \\ &\leq \exp \left(-\frac{n(1-2q')^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right). \end{aligned}$$

Using the above bound along with the bound on $\mathbb{P}(\mathcal{E}_0^c)$ into (F.10), we arrive at

$$\mathbb{P}(\sum_{\ell} T_{\ell} < 0) \leq 2 \exp \left(-\frac{n(1-2q')^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) + 4 \exp \left(-\frac{n}{32}(1-2p)^2 \right). \quad (\text{F.11})$$

Invoking (F.6) and the definition of T_{ℓ} given by (F.7) we obtain

$$\mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0; \mathcal{E}) \leq 2 \exp \left(-\frac{n(1-2q')^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) + 4 \exp \left(-\frac{n}{32}(1-2p)^2 \right). \quad (\text{F.12})$$

Finally we combine (F.12), (F.5) and (F.4) to get

$$\begin{aligned} \mathbb{P}(y\langle \mathbf{x}, \widehat{\boldsymbol{\theta}}_1 \rangle < 0) &\leq 2 \exp \left(-\frac{n(1-2q')^2 \langle y\mathbf{x}, \boldsymbol{\mu} \rangle^2}{64(\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2)} \right) \\ &\quad + 4 \exp \left(-\frac{n}{32}(1-2p)^2 \right) + \frac{n}{2} \left(\exp \left(-\frac{\gamma^4}{8\lambda_{\max}^2}(1-2p')\frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}. \end{aligned} \quad (\text{F.13})$$

This finishes the proof of Theorem 4.8.

We will now prove the intermediate lemmas F.1 and F.2.

F.1 Proof of Lemma F.1

First note that $\mathbb{P}(\tilde{y}'_\ell \neq \text{sign}(\langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle)) = \mathbb{P}(\tilde{y}'_\ell \langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle < 0)$.

Define $\tilde{\xi}_i := \xi_i(1 + u_i)$. Then using (F.1), we have

$$\tilde{y}'_\ell \langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle = \frac{\tilde{y}'_\ell}{n} \langle \mathbf{x}_\ell, (\sum_i \tilde{\xi}_i) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_\ell y_\ell \mathbf{z}_\ell + \boldsymbol{\Sigma}^{1/2} \sum_{i \neq \ell} \xi_i y_i \mathbf{z}_i \rangle.$$

We have $\sum_{i \neq \ell} \xi_i y_i \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, (n-1) \mathbf{I}_d)$ and so $\langle \tilde{y}'_\ell \mathbf{x}_\ell, \boldsymbol{\Sigma}^{1/2} \sum_{i \neq \ell} \xi_i y_i \mathbf{z}_i \rangle$ is a zero mean Gaussian with variance $(n-1) \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}^2$. Let \mathcal{F}_ℓ be the σ -algebra generated by $(\mathbf{z}_\ell, y_\ell, \{\xi_i, u_i\}_{i \in [n]})$. Note that \mathbf{z}_i , for $i \neq \ell$, is independent of \mathcal{F}_ℓ and \tilde{y}'_ℓ is \mathcal{F}_ℓ -measurable. Hence,

$$\begin{aligned} \mathbb{P}(\tilde{y}'_\ell \langle \mathbf{x}_\ell, \widehat{\boldsymbol{\theta}}_0 \rangle < 0 | \mathcal{F}_\ell) &= \Phi \left(-\frac{\tilde{y}'_\ell \langle \mathbf{x}_\ell, (\sum_i \tilde{\xi}_i) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_\ell y_\ell \mathbf{z}_\ell \rangle}{\sqrt{n-1} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}} \right) \\ &\stackrel{(a)}{=} \Phi \left(-\frac{|\langle \mathbf{x}_\ell, (\sum_i \tilde{\xi}_i) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_\ell y_\ell \mathbf{z}_\ell \rangle|}{\sqrt{n-1} \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}} \right) \\ &\stackrel{(b)}{\leq} \frac{1}{2} \exp \left(-\frac{|\langle \mathbf{x}_\ell, (\sum_i \tilde{\xi}_i) \boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2} \xi_\ell y_\ell \mathbf{z}_\ell \rangle|^2}{2n \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}^2} \right) \\ &\stackrel{(c)}{\leq} \frac{1}{2} \exp \left(-\frac{\langle \mathbf{x}_\ell, \boldsymbol{\mu} \rangle^2 (\sum_i \tilde{\xi}_i)^2}{4n \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}^2} \right) \exp \left(\frac{\langle \mathbf{x}_\ell, \boldsymbol{\Sigma}^{1/2} \mathbf{z}_\ell \rangle^2}{2n \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}^2} \right) \\ &\leq \frac{1}{2} \exp \left(-\frac{\langle \mathbf{x}_\ell, \boldsymbol{\mu} \rangle^2 (\sum_i \tilde{\xi}_i)^2}{4n \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}^2} \right) \exp \left(\frac{\|\mathbf{z}_\ell\|_{\ell_2}^2}{2n} \right) \\ &\stackrel{(d)}{\leq} \frac{1}{2} \exp \left(-\frac{\gamma^4 (\sum_i \tilde{\xi}_i)^2}{4n \lambda_{\max}^2 \|\mathbf{z}_\ell\|_{\ell_2}^2} \right) \exp \left(\frac{\|\mathbf{z}_\ell\|_{\ell_2}^2}{2n} \right). \end{aligned} \tag{F.14}$$

Here, (a) follows by the choice of \tilde{y}'_ℓ ; (b) holds by using Equation (7.1.13) of Abramowitz & Stegun (1968):

$$\Phi^c(t) \leq \sqrt{\frac{2}{\pi}} \frac{e^{-t^2/2}}{t + \sqrt{t^2 + \frac{8}{\pi}}} \leq \frac{1}{2} e^{-t^2/2},$$

for all $t \geq 0$. In addition (c) holds since $(a+b)^2 \geq \frac{a^2}{2} - b^2$. Also, (d) holds because under our data model we have $\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell = \boldsymbol{\Sigma} \mathbf{z}_\ell$ and so $\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2} \leq \lambda_{\max} \|\mathbf{z}_\ell\|_{\ell_2}$. Further, $|\langle \mathbf{x}_\ell, \boldsymbol{\mu} \rangle| = |y(1 + u_\ell)| \|\boldsymbol{\mu}\|_{\ell_2}^2 \geq \gamma^2$ because $u_\ell \geq 0$. Putting these bounds together, we have

$$\frac{\langle \mathbf{x}_\ell, \boldsymbol{\mu} \rangle}{\|\boldsymbol{\Sigma}^{1/2} \mathbf{x}_\ell\|_{\ell_2}} \geq \frac{\gamma^2}{\lambda_{\max} \|\mathbf{z}_\ell\|_{\ell_2}}.$$

We proceed by taking expectation of the right-hand side of (F.14) with respect to $\tilde{\xi}_1, \dots, \tilde{\xi}_n$. Define the shorthand $\lambda := \frac{\gamma^4}{4n \lambda_{\max}^2 \|\mathbf{z}_\ell\|_{\ell_2}^2}$. Fix an arbitrary $\lambda_0 > 0$ for now (we will determine its value later) and define truncated parameter

$\bar{\lambda} = \min(\lambda, \lambda_0)$. We have

$$\mathbb{E} \left[e^{-\lambda(\sum_i \tilde{\xi}_i)^2} \right] \leq \mathbb{E} \left[e^{-\bar{\lambda}(\sum_i \tilde{\xi}_i)^2} \right] \quad (\text{F.15})$$

$$\begin{aligned} &= \mathbb{E} \left[e^{-\bar{\lambda} \sum_i (1+u_i)^2} e^{-\bar{\lambda} \sum_{i \neq j} \xi_i \xi_j (1+u_i)(1+u_j)} \right] \\ &\leq e^{-\bar{\lambda}n} \prod_{i \neq j} \mathbb{E} \left[e^{-\bar{\lambda} \xi_i (1+u_i)} \right] \mathbb{E} \left[e^{-\bar{\lambda} \xi_j (1+u_j)} \right] \\ &= e^{-\bar{\lambda}n} (\mathbb{E} [e^{-\bar{\lambda} \xi (1+u)}])^{n(n-1)} \\ &= e^{-\bar{\lambda}n} (\mathbb{E} [(1-p)e^{-\bar{\lambda}(1+u)} + pe^{\bar{\lambda}(1+u)}])^{n(n-1)} \\ &\leq e^{-\bar{\lambda}n} \left((1-p)e^{-\bar{\lambda}} + pe^{\bar{\lambda}} \mathbb{E} [e^{\bar{\lambda}u}] \right)^{n(n-1)} \\ &\leq e^{-\bar{\lambda}n} \left((1-p)e^{-\bar{\lambda}} + pe^{\bar{\lambda}} e^{\bar{\lambda}^2/2} \right)^{n(n-1)} \\ &= e^{-\bar{\lambda}n^2} \left(1-p + pe^{2\bar{\lambda}+\bar{\lambda}^2/2} \right)^{n(n-1)}. \end{aligned} \quad (\text{F.16})$$

Here, the first and the second inequality follows from $u_i \geq 0$. The third inequality holds since u has unit sub-gaussian norm and so its moment-generating function is bounded as $\mathbb{E}[e^{\lambda u}] \leq e^{\lambda^2/2}$.

We next further upper bound the right-hand side of F.16 to get a simpler expression. In doing that we use the following lemma (proved in Appendix F.1.1).

Lemma F.3. For $t \leq \frac{1}{2}$, we have $e^t - 1 \leq t + t^2$.

By choosing $\lambda_0 \leq 0.2$ we have $2\lambda_0 + \lambda_0^2/2 \leq 1/2$ and so $2\bar{\lambda} + \bar{\lambda}^2/2 \leq 1/2$. Using the above lemma, we have

$$\begin{aligned} e^{2\bar{\lambda}+\bar{\lambda}^2/2} - 1 &\leq 2\bar{\lambda} + \bar{\lambda}^2/2 + (2\bar{\lambda} + \bar{\lambda}^2/2)^2 \\ &= 2\bar{\lambda} + \bar{\lambda}^2/2 + 4\bar{\lambda}^2 + \bar{\lambda}^4/4 + 2\bar{\lambda}^3 \\ &\leq 2\bar{\lambda} + \bar{\lambda} (\lambda_0/2 + 4\lambda_0 + \lambda_0^3/4 + 2\lambda_0^2) \\ &\leq (2 + 6\lambda_0)\bar{\lambda}, \end{aligned}$$

where we used $\lambda_0 \leq 0.2$ in the last step. Next using the inequality $1 + x \leq e^x$ for $x \geq 0$, we get

$$(e^{2\bar{\lambda}+\bar{\lambda}^2/2} - 1)p + 1 \leq (2 + 6\lambda_0)p\bar{\lambda} + 1 \leq e^{(2+6\lambda_0)\bar{\lambda}p}.$$

By using this bound in (F.16) we obtain

$$\mathbb{E} \left[e^{-\lambda(\sum_i \tilde{\xi}_i)^2} \right] \leq e^{-\bar{\lambda}n^2} e^{(2+6\lambda_0)\bar{\lambda}pn^2} = e^{-(1-2p')\bar{\lambda}n^2}, \quad (\text{F.17})$$

with $p' := (1 + 3\lambda_0)p$.

Using the above in (F.14) and then taking expectation with respect to \mathbf{z}_ℓ we get

$$\begin{aligned} \mathbb{P}(\tilde{y}_\ell(\mathbf{x}_\ell, \hat{\boldsymbol{\theta}}_0) < 0) &\leq \frac{1}{2} \mathbb{E} \left[\exp(-(1-2p')\bar{\lambda}n^2) \exp\left(\frac{\|\mathbf{z}_\ell\|_{\ell_2}^2}{2n}\right) \right] \\ &\leq \frac{1}{2} \mathbb{E} \left[\exp(-2(1-2p')\bar{\lambda}n^2) \right]^{1/2} \mathbb{E} \left[\exp\left(\frac{\|\mathbf{z}_\ell\|_{\ell_2}^2}{n}\right) \right]^{1/2}, \end{aligned} \quad (\text{F.18})$$

by applying the Cauchy–Schwarz inequality.

Using the moment generating function of χ_d^2 distribution, we have

$$\mathbb{E} \left[\exp\left(\frac{\|\mathbf{z}_\ell\|_{\ell_2}^2}{n}\right) \right] = \left(1 - \frac{2}{n}\right)^{-d/2} \leq e^{2d/n}, \quad (\text{F.19})$$

for $n \geq 4$. Here, we use the inequality $(1-x)^{-1} \leq e^{2x}$ for $x \in [0, 1/2]$.

We continue by bounding the first term on the right-hand side of (F.18). Define the event \mathcal{E}_1 as follows:

$$\mathcal{E}_1 := \left\{ \|z_\ell\|_{\ell_2} \leq \sqrt{2d} \right\}. \quad (\text{F.20})$$

We then have

$$\begin{aligned} \mathbb{E} \left[\exp \left(-2(1-2p')\bar{\lambda}n^2 \right) \right] &= \mathbb{E} \left[\exp \left(-2(1-2p')\bar{\lambda}n^2 \right) (\mathbb{1}_{\mathcal{E}_1} + \mathbb{1}_{\mathcal{E}_1^c}) \right] \\ &\leq \mathbb{E} \left[\exp \left(-2(1-2p')\bar{\lambda}n^2 \right) \mathbb{1}_{\mathcal{E}_1} \right] + \mathbb{E} \left[\mathbb{1}_{\mathcal{E}_1^c} \right] \\ &\stackrel{(a)}{=} \mathbb{E} \left[\exp \left(-\frac{\gamma^4}{4\lambda_{\max}^2 d} (1-2p')n \right) \mathbb{1}_{\mathcal{E}_1} \right] + \mathbb{P}(\mathcal{E}_1^c) \\ &\leq \exp \left(-\frac{\gamma^4}{4\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + \mathbb{P}(\mathcal{E}_1^c) \end{aligned} \quad (\text{F.21})$$

Note that in step (a), we used the fact that on the event \mathcal{E}_1 , we have $\lambda \geq \gamma^4/(8\lambda_{\max}^2 nd)$ and we choose $\lambda_0 = \gamma^4/(8\lambda_{\max}^2 nd)$. This way, we have $\bar{\lambda} = \min(\lambda, \lambda_0) = \gamma^4/(8\lambda_{\max}^2 nd)$. Also note that by our assumption in the statement of the lemma, we have $\lambda_0 \leq 0.2$ which is the condition assumed in deriving (F.17).

We next bound $\mathbb{P}(\mathcal{E}_1^c)$. Since $z_\ell \sim \mathcal{N}(0, \mathbf{I}_d)$, $\|z_\ell\|_{\ell_2}^2$ has χ^2 distribution with d -degrees of freedom. Using the tail bound of χ^2 distribution (see Theorem 3 of Ghosh (2021)) we have $\mathbb{P}(\|z_\ell\|_{\ell_2}^2 \geq 2d) \leq e^{-d/8}$ and so

$$\mathbb{P}(\mathcal{E}_1^c) \leq \mathbb{P}(\|z_\ell\|_{\ell_2}^2 \geq 2d) \leq e^{-d/8}.$$

Using the above in (F.21) we obtain

$$\mathbb{E} \left[\exp \left(-2(1-2p')\bar{\lambda}n^2 \right) \right] \leq \exp \left(-\frac{\gamma^4}{4\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + e^{-d/8}. \quad (\text{F.22})$$

By combining (F.22) and (F.19) with (F.18) we arrive at

$$\begin{aligned} \mathbb{P}(\tilde{y}'_\ell \langle \mathbf{x}_\ell, \hat{\boldsymbol{\theta}}_0 \rangle < 0) &\leq \frac{1}{2} \left(\exp \left(-\frac{\gamma^4}{4\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + e^{-d/8} \right)^{1/2} e^{d/n} \\ &\leq \frac{1}{2} \left(\exp \left(-\frac{\gamma^4}{8\lambda_{\max}^2} (1-2p') \frac{n}{d} \right) + e^{-d/16} \right) e^{d/n}, \end{aligned}$$

which completes the proof of lemma.

F.1.1 PROOF OF LEMMA F.3

Using the Taylor series expansion of e^t , we have:

$$e^t - 1 = t + \frac{t^2}{2} + \sum_{j=3}^{\infty} \frac{t^j}{j!} \quad (\text{F.23})$$

$$\leq t + \frac{t^2}{2} + \frac{t^2}{2} \sum_{j=3}^{\infty} \frac{1}{j!} \quad (\text{F.24})$$

$$= t + \frac{t^2}{2} + \frac{t^2}{2} \left(e - \frac{5}{2} \right) \quad (\text{F.25})$$

(F.24) follows by using the fact that for $j \geq 3$ and $t \leq \frac{1}{2}$, $t^j \leq t^3 \leq \frac{t^2}{2}$. (F.25) follows by using the fact that $e^1 - 1 = 1 + \frac{1}{2} + \sum_{j=3}^{\infty} \frac{1}{j!}$ after plugging in $t = 1$ in Equation (F.23).

Next, using the fact that $e < 3$ in (F.25), we get:

$$e^t - 1 < t + \frac{3t^2}{4} < t + t^2. \quad (\text{F.26})$$

This completes the proof.

F.2 Proof of Lemma F.2

Under our data model we have $\mathbf{x}_\ell = y_\ell(1 + u_\ell)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell$. Substituting for \mathbf{x}_ℓ in the expression of T_ℓ , we get

$$\begin{aligned} T_\ell &= \text{sign}\left(\langle y_\ell(1 + u_\ell)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell, \beta\boldsymbol{\mu} + \xi_\ell y_\ell \boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell \rangle\right) \langle y\mathbf{x}, y_\ell(1 + u_\ell)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell \rangle \\ &= \text{sign}\left(y_\ell(1 + u_\ell)\beta\gamma^2 + \xi_\ell y_\ell \|\boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell\|_{\ell_2}^2\right) \langle y\mathbf{x}, y_\ell(1 + u_\ell)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell \rangle, \end{aligned}$$

using the fact that $\boldsymbol{\Sigma}^{1/2}\boldsymbol{\mu} = 0$. Taking expectation we get

$$\mathbb{E}[T_\ell|\beta] = \mathbb{E}\left[\text{sign}\left(y_\ell(1 + u_\ell)\beta\gamma^2 + \xi_\ell y_\ell \|\boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell\|_{\ell_2}^2\right) \langle y\mathbf{x}, y_\ell(1 + u_\ell)\boldsymbol{\mu} \rangle \middle| \beta\right]$$

since the other term will be an odd function of \mathbf{z}_ℓ . Letting

$$q_0 := \mathbb{P}\left((1 + u_\ell)\beta\gamma^2 + \xi_\ell \|\boldsymbol{\Sigma}^{1/2}\mathbf{z}_\ell\|_{\ell_2}^2 < 0\right),$$

we get

$$\mathbb{E}[T_\ell|\beta] = (1 - 2q_0)\langle y\mathbf{x}, (1 + u_\ell)\boldsymbol{\mu} \rangle \geq (1 - 2q_0)\langle y\mathbf{x}, \boldsymbol{\mu} \rangle, \quad (\text{F.27})$$

because $u_\ell > 0$ and $\langle y\mathbf{x}, \boldsymbol{\mu} \rangle > 0$, given the positive margin in our data model, and if $q_0 \leq \frac{1}{2}$. So what remains is to show that $q_0 \leq q \leq 1/2$.

We write

$$\begin{aligned} q_0 &\leq \mathbb{P}\left((1 + u_\ell)\beta\gamma^2 - \lambda_{\max} \|\mathbf{z}_\ell\|_{\ell_2}^2 < 0\right) \\ &\leq \mathbb{P}\left(\frac{\beta\gamma^2}{\lambda_{\max}} < \|\mathbf{z}_\ell\|_{\ell_2}^2\right) \quad (\text{since } u_\ell > 0) \\ &< \exp\left(-\frac{\beta\gamma^2}{20\lambda_{\max}}\right), \end{aligned}$$

if $\frac{\beta\gamma^2}{2d\lambda_{\max}} > 1$, where the last step follows from the observation that $\|\mathbf{z}_\ell\|_{\ell_2}^2 \sim \chi_d^2$ and using the tail bound of χ_d^2 (Theorem 3 of Ghosh (2021)). So when $\frac{\beta\gamma^2}{2d\lambda_{\max}} > 1$ and $d \geq 7$,

$$q_0 \leq q := \exp\left(-\frac{\beta\gamma^2}{20\lambda_{\max}}\right) \leq e^{-d/10} < \frac{1}{2}.$$

Plugging this into (F.27) completes the proof.

G Proof of Theorem 4.9

Using Theorem 4.8 (and because the training set \mathcal{T} has measure zero), we have $\text{err}(\widehat{\boldsymbol{\theta}}_1) \leq \mathbb{E}[\alpha_1(\mathbf{x})]$. Based on our data model, $\mathbf{x} = y(1 + u)\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\mathbf{z}$. Consider the probabilistic event $\mathcal{E}_2 := \{\mathbf{z} : \|\mathbf{z}\|_{\ell_2}^2 \leq 2d\}$. For every $\eta > 0$, we have

$$\begin{aligned} \mathbb{E}\left[\exp\left(-\frac{\eta\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2}\right)\right] &= \mathbb{E}\left[\exp\left(-\frac{\eta\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2}\right) (\mathbb{1}_{\mathcal{E}_2} + \mathbb{1}_{\mathcal{E}_2^c})\right] \\ &\leq \exp\left(-\frac{\eta\gamma^4}{\gamma^4 + 2\lambda_{\max}^2 d}\right) + \mathbb{P}(\mathcal{E}_2^c), \end{aligned}$$

where the inequality holds because $\langle \mathbf{x}, \boldsymbol{\mu} \rangle = y(1 + u)\gamma^2$ and so $\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 \geq \gamma^4$ (since u is a non-negative random variable). In addition,

$$\|\boldsymbol{\Sigma}^{1/2}\mathbf{x}\|_{\ell_2}^2 = \|\boldsymbol{\Sigma}\mathbf{z}\|_{\ell_2}^2 \leq \lambda_{\max}^2 \|\mathbf{z}\|_{\ell_2}^2 \leq 2\lambda_{\max}^2 d,$$

under the event \mathcal{E}_2 . Further, $\|\mathbf{z}\|_{\ell_2}^2 \sim \chi_d^2$ and so using tail bounds for χ_d^2 (see Theorem 3 of [Ghosh \(2021\)](#)), we have $\mathbb{P}(\mathcal{E}_2^c) \leq e^{-d/8}$. Putting things together, we obtain

$$\mathbb{E} \left[\exp \left(- \frac{\eta \langle \mathbf{x}, \boldsymbol{\mu} \rangle^2}{\langle \mathbf{x}, \boldsymbol{\mu} \rangle^2 + \|\boldsymbol{\Sigma}^{1/2} \mathbf{x}\|_{\ell_2}^2} \right) \right] \leq \exp \left(- \frac{\eta \gamma^4}{\gamma^4 + 2\lambda_{\max}^2 d} \right) + e^{-d/8}.$$

The claim of theorem follows by setting $\eta = \frac{n}{64}(1 - 2q')^2$.

H Remaining Experimental Details

Here we provide the remaining details about the experiments in Section 5. Our experiments were done using TensorFlow and JAX. In all the cases, we retrain starting from random initialization rather than the previous checkpoint we converged to before RT; the former worked better than the latter. We list training details for each individual dataset next.

CIFAR-10. Optimizer is SGD with momentum = 0.9, batch size = 32, number of epochs in each stage of training (i.e., both stages of baseline, full RT and consensus-based RT) = 30. We use the cosine one-cycle learning rate schedule with initial learning rate = 0.1 for each stage of training. The number of epochs and initial learning rate were chosen based on the performance of the baseline method and *not* based on the performance of full or consensus-based RT. Standard augmentations such as random cropping, flipping and brightness/contrast change were used.

CIFAR-100. Details are the same as CIFAR-10 except that here the number of epochs in each stage of training = 40 and initial learning rate = 0.005.

DomainNet. Details are the same as CIFAR-10 except that here the number of epochs in each stage of training = 15, initial learning rates for the first stage of the baseline in Tables 7, 8, and 9 are $1e-3$, $5e-2$, and $5e-4$, respectively, whereas the initial learning rate for the second stage of the baseline as well as full RT and consensus-based RT is $1e-3$.

AG News Subset. Small BERT model link: <https://www.kaggle.com/models/tensorflow/bert/frameworks/tensorFlow2/versions/bert-en-uncased-1-4-h-512-a-8/versions/2?tfhub-redirect=true>, BERT English uncased preprocessor link: <https://www.kaggle.com/models/tensorflow/bert/frameworks/tensorFlow2/versions/en-uncased-preprocess/versions/3?tfhub-redirect=true>. Optimizer is Adam with fixed learning rate = $1e-5$, batch size = 32, number of epochs in each training stage = 5.

I Accuracy over the Entire Dataset and over the Consensus Set in the Case of AG News Subset

Here we list the accuracies of the predicted labels and the given labels over the entire dataset and the accuracy of the predicted labels (= given labels) over the consensus set for AG News Subset. Just like in Table 3, the accuracy of the predicted labels over the consensus set is significantly more than the accuracy of the predicted and given labels over the entire dataset. This explains why consensus-based RT performs the best, even though the consensus set is much smaller than the full dataset ($\sim 28\%$, 32% and 38% of the entire training set for the three values of ϵ).

Table 11. **AG News Subset.** Accuracies of predicted labels and given labels over the entire dataset and accuracies of predicted labels (= given labels) over the consensus set. Conclusions are the same as in Table 3.

ϵ	Acc. of <i>predicted</i> labels on <i>full</i> dataset	Acc. of <i>given</i> labels on <i>full</i> dataset	Acc. of <i>predicted</i> labels on <i>consensus</i> set
0.3	53.20 ± 2.82	32.52 ± 2.05	61.81 ± 2.66
0.5	66.78 ± 1.31	35.5 ± 0.14	76.48 ± 0.93
0.8	79.98 ± 0.80	42.53 ± 0.13	89.59 ± 0.43

J Consensus-Based Retraining Does Better than Confidence-Based Retraining

Here we compare full and consensus-based RT against another strategy for retraining which we call **confidence-based retraining (RT)**. Specifically, we propose to retrain with the predicted labels of the samples with the top 50% margin (i.e., highest predicted probability - second highest predicted probability); margin is a measure of the model’s confidence. This idea is similar to self-training’s method of sample selection in the semi-supervised setting (Amini et al., 2022). In Tables 12 and 13, we show results for CIFAR-10 and CIFAR-100 (in the same setting as Section 5 and Appendix H) with the smallest value of ϵ from Tables 1 and 2, respectively. Notice that *consensus-based RT is clearly better than confidence-based RT*.

Table 12. **CIFAR-10**. Test set accuracies (mean \pm standard deviation). *Consensus-based RT performs the best.*

ϵ	Baseline	Full RT	Consensus-based RT	Confidence-based RT
1	57.78 \pm 1.13	60.07 \pm 0.63	63.84 \pm 0.56	62.09 \pm 0.55

 Table 13. **CIFAR-100**. Test set accuracies (mean \pm standard deviation). Again, *consensus-based RT performs the best.*

ϵ	Baseline	Full RT	Consensus-based RT	Confidence-based RT
3	23.53 \pm 1.01	24.42 \pm 1.22	29.98 \pm 1.11	24.99 \pm 1.25

K Retraining is Beneficial Even Without a Validation Set

In all our previous experiments, we assumed access to a small clean validation set. Here we show the results of training ResNet-34 on CIFAR-100 when we do *not* have access to a validation set and train for 100 epochs, and also compare them against the corresponding results with a validation set where we trained for 40 epochs (Table 4 in Section 5). The point of these experiments is to show that retraining can offer gains even without a validation set, in which case we train for too many epochs and expect overfitting. Please see the results and discussion in Table 14.

Table 14. **CIFAR-100 w/ ResNet-34 without (top) and with (bottom) a validation set**. Test set accuracies (mean \pm standard deviation). Note that retraining (in particular, consensus-based RT) leads to improvement even without a validation set. However, the performance and amount of improvement obtained with retraining are worse in the absence of a validation set due to a higher degree of overfitting. This is not surprising.

Without Val. Set	ϵ	Baseline	Full RT	Consensus-based RT
	4	26.03 \pm 1.75	28.33 \pm 1.76	30.47 \pm 1.31
	5	42.50 \pm 1.14	44.43 \pm 1.27	46.27 \pm 1.72
With Val. Set (Same as Table 4)	ϵ	Baseline	Full RT	Consensus-based RT
	4	37.53 \pm 1.58	39.33 \pm 1.37	43.87 \pm 1.62
	5	51.13 \pm 0.69	52.33 \pm 0.38	55.43 \pm 0.42

L Beyond Label DP: Evaluating Retraining in the Presence of Human Annotation Errors

Even though our empirical focus in this paper has been label DP training, retraining (RT) can be employed for general problems with label noise. Here we evaluate RT in a setting with “real” label noise due to *human annotation*. Specifically, we focus on training a ResNet-18 model (*without* label DP to be clear) on the CIFAR-100N dataset introduced by Wei et al. (2022) and available on the TensorFlow website. CIFAR-100N is just CIFAR-100 labeled by humans; thus, it has real human annotation errors. The experimental setup and details are the same as CIFAR-100 (as stated in Section 5 and Appendix H); the only difference is that here we use initial learning rate = 0.01.

In Table 15, we list the test accuracies of the baseline which is just standard training with the given labels, full RT and consensus-based RT, respectively. Even here *with human annotation errors*, consensus-based RT is beneficial.

Table 15. **CIFAR-100N**. Test set accuracies (mean \pm standard deviation). *So even with real human annotation errors, consensus-based RT improves performance.*

Baseline	Full RT	Consensus-based RT
55.47 \pm 0.18	56.88 \pm 0.35	57.68 \pm 0.35