
Self-training Avoids Using Spurious Features Under Domain Shift

Yining Chen*, Colin Wei*, Ananya Kumar, Tengyu Ma

Department of Computer Science

Stanford University

{cynnjjs, colinwei, ananya1, tengyuma}@stanford.edu

Abstract

In unsupervised domain adaptation, existing theory focuses on situations where the source and target domains are close. In practice, conditional entropy minimization and pseudo-labeling work even when the domain shifts are much larger than those analyzed by existing theory. We identify and analyze one particular setting where the domain shift can be large, but these algorithms provably work: certain spurious features correlate with the label in the source domain but are independent of the label in the target. Our analysis considers linear classification where the spurious features are Gaussian and the non-spurious features are a mixture of log-concave distributions. For this setting, we prove that entropy minimization on unlabeled target data will avoid using the spurious feature if initialized with a decently accurate source classifier, even though the objective is non-convex and contains multiple bad local minima using the spurious features. We verify our theory for spurious domain shift tasks on semi-synthetic Celeb-A and MNIST datasets. Our results suggest that practitioners collect and self-train on large, diverse datasets to reduce biases in classifiers even if labeling is impractical.

1 Introduction

Reliable machine learning systems need to generalize to test distributions that are different from the training distribution. However, the test performance of machine learning models often significantly degrades as the test domain drifts away from the training domain. Various approaches have been proposed to adapt the models to new domains [37, 10, 38] but theoretical understanding of these algorithms is limited. Prior theoretical works focus on settings where the target domain is sufficiently close to the source domain [5, 29, 17, 34]. To theoretically study realistic scenarios where domain shifts can be much larger, we need to leverage additional structure of the shifts.

Towards this goal, we propose to study a particular structured domain shift for which unsupervised domain adaptation is provably feasible: in the source domain, a subset of “spurious” features correlate with the label, whereas in the unlabeled target data, these features are independent of the label. In real-world training data, these spurious correlations can occur due to biased sampling or artifacts in crowd-sourcing [14]. For example, we may have a labeled dataset for recidivism prediction where race correlates with recurrence of crime due to sample selection bias, but this correlation does not hold on the population. Models which learn spurious correlations can generalize poorly on population data which does not have these biases [23]. In these settings, it could be impractical to acquire labels for an unbiased sub-sample of the population, but unlabeled data is often available.

We prove that in certain settings, perhaps surprisingly, self-training on *unlabeled* target data can avoid using these spurious features. Our theoretical results apply to two closely-related popular algorithms:

*Equal Contribution

self-training [20] and conditional entropy minimization [13]. In practice, self-training has achieved competitive or state-of-the-art results in unsupervised domain adaptation [22, 44, 30], but there are few theoretical analyses of self-training when there is domain shift.

Our theoretical setting and analysis are consistent with recent large-scale empirical results by Xie et al. [42], which suggest that self-training on a more diverse unlabeled dataset can improve the robustness of a model, potentially by avoiding using spurious correlations. These results and our theory help emphasize the value of a large and diverse unlabeled dataset, even if labeled data is scarce.

Formally, we assume that each input consists of two subsets of features, denoted by x_1 and x_2 . x_1 is the “signal” feature that determines the label y in the target distribution. x_2 is the spurious feature that correlates with the label y in the source domain, but x_2 is independent of (x_1, y) in the target domain. For a first-cut result, we consider binary classification and linear models on the features (x_1, x_2) , where the spurious feature x_2 is a multivariate Gaussian and x_1 is a mixture of log-concave distributions. We aim to show that, initialized with some classifier trained on the source data, self-training on the unlabeled target will remove usage of the spurious feature x_2 .

A challenge in the analysis is that self-training on an unlabeled loss can possibly harm, rather than help, target accuracy by amplifying the mistakes of source classifier (see Section 3.1). The classical idea of co-training [8] deals with this by assuming the mistakes of the classifier are independent of x , reducing the problem to learning from noisy labels. However, in our setting the source classifier makes biased mistakes which depend on x , and self-training potentially reinforces these biases if there are no additional assumptions. For example, we require initialization with a decently accurate source classifier, and we empirically verify the necessity of this assumption in Section 5.

Our main contribution (Theorem 3.1) is to prove that self-training and conditional min entropy using finite unlabeled data converge to a solution that has 0 coefficients on the spurious feature x_2 , assuming the following: 1. the signal x_1 is a mixture of well-separated log-concave distributions and 2. the initial source classifier is decently accurate on target data and avoids relying too heavily on the spurious feature. In a simpler setting where x_1 is a univariate Gaussian, we show that self-training using a decently accurate source classifier converges to the Bayes optimal solution (Theorem 3.2).

We run simulations on semi-synthetic colored MNIST [19] and celebA [21] datasets to verify the insights from our theory and show that they apply to multi-layer neural networks and datasets where the spurious features are not necessarily a subset of the input coordinates (Section 5).

1.1 Related Work

Self-training methods have achieved state-of-the-art results for semi-supervised learning [42, 32, 20], adversarial robustness [22, 44, 30], and unsupervised domain adaptation [22, 44, 30], but there is little understanding of when and why these methods work under domain shifts. Two popular forms of self-training are pseudolabeling [20] and conditional entropy minimization [13], which have been observed to be closely connected [2, 20, 30, 7]. We show that our analysis applies to both entropy minimization and a version of pseudo-labeling where we initialize the student model with the teacher model and re-label after each gradient step (Proposition D.1).

Kumar et al. [18] examine self-training for domain adaptation, but under strong assumptions: that $P(X|Y)$ is an isotropic Gaussian, that entropy minimization converges to the nearest local minima, and infinite unlabeled data. They use a symmetry argument that requires all these assumptions. In our setting, the signal x_1 can be a mixture of many log-concave, log-smooth distributions, and we show that self-training does in fact converge with only finite unlabeled data, even though the loss landscape is non-convex. These require new, more general, proof techniques.

Domain adaptation and semi-supervised learning theory: Importance weighting [29, 17, 34] is a popular way to deal with *covariate shift* but these methods assume that $P(Y | X)$ is the same for the source and target, which may not hold when there are spurious correlations in the source but not target. Additionally, sample complexity bounds for importance weighting depend on the expected density ratios between the source and target, which can often scale exponentially in the dimension. Our finite sample guarantees only depend on properties of the target distribution (assuming a decently accurate source classifier) and are agnostic to this density ratio. The theory of $H\Delta H$ -divergence lower bounds target accuracy of a classifier in terms of source accuracy if some distance between the domains is small [6]; Zhang et al. [43] extend this distance measure to multiclass classification. In

contrast, we show self-training can improve accuracy under our structured domain shift, even when the shift is potentially large. Other theoretical papers on semi-supervised learning focus on analyzing when unlabeled data can help, but do not analyze domain shift [27, 31, 5, 4].

Co-training [8] is an algorithm that can leverage unlabeled data when the input features can be split into (x_1, x_2) that are conditionally independent given the label. Co-training assumes this grouping is known a-priori, and that either group can be used to predict the label accurately. In our setting, the spurious feature cannot be used to predict the label accurately in the target domain, and the algorithm does not have access to the grouping between spurious and signal features.

Spurious and non-robust features. Many works seek to identify causal features invariant across domains [3, 26, 15]. Spurious features are also related to adversarial examples, which can possibly be attributed to non-robust features that can predict the label but are brittle under domain shift [16].

A number of papers theoretically analyze the connection between adversarial robustness and accuracy or generalization for linear classifiers in simple Gaussian settings [36, 28, 39]. Carmon et al. [9] show that self-training on unlabeled data can improve adversarially robust generalization for linear models in a Gaussian setting. Though these research questions are orthogonal to ours, one technical contribution of our work is that our analysis extends to more general distributions than Gaussians.

Fairness. Spurious correlations in datasets can lead to unfair predictions when protected attributes are involved. Our work shows that self-training can potentially employ unlabeled population samples to overcome bias in labeled data [12, 35].

2 Setup

Model. We consider a linear model $\hat{y} = w^\top x$ where $w = (w_1, w_2)$ and $x = (x_1, x_2)$ with $w_1, x_1 \in \mathbb{R}^{d_1}$ and $w_2, x_2 \in \mathbb{R}^{d_2}$. We assume that the spurious features x_2 have Gaussian distribution with covariance $\Sigma_2 \succ 0$, so the target data $(x, y) \sim \mathcal{D}_{\text{tg}}$ is generated by

$$\begin{aligned} y &\stackrel{\text{unif}}{\sim} \{\pm 1\}, \text{ and } x_1 \sim \mathcal{D}_{\text{tg},1}(\cdot|y) \\ x_2 &\sim \mathcal{N}(\vec{0}, \Sigma_2), \Sigma_2 \in \mathbb{R}^{d_2 \times d_2} \end{aligned} \quad (2.1)$$

for some distribution $\mathcal{D}_{\text{tg},1}$ over \mathbb{R}^{d_1} . Note that x_2 is a spurious feature because it is independent of the label y . Our results and analysis also transfer to a “scrambled setup” [3] where we observe $z = \mathcal{S}x$ for some rotation matrix $\mathcal{S} \in \mathbb{R}^{(d_1+d_2) \times (d_1+d_2)}$. This follows as a direct consequence of the rotational invariance of the algorithm (2.3) and our assumptions.

Min-entropy objective. The min-entropy objective on a target unlabeled example is defined as $\ell_{\text{ent}}(w^\top x)$ where $\ell_{\text{ent}}(t) = H((1 + \exp(-t))^{-1})$ and H is the binary entropy function.

For mathematical convenience, we consider an approximation $\ell_{\text{exp}}(t) = \exp(-|t|)$, which is commonly used in the literature for studying the logistic loss [33]. ℓ_{exp} approximates ℓ_{ent} up to a constant factor and exhibits the same tail behavior (Figure 10). We experimentally validate in Section E.5 that training using ℓ_{exp} achieves the same effect for the algorithms we analyze. The population unlabeled objective on the target distribution that we consider is

$$L(w) \triangleq \mathbb{E}_{x \sim \mathcal{D}_{\text{tg}}} \ell_{\text{exp}}(w^\top x) \quad (2.2)$$

where \mathcal{D}_{tg} denotes the distribution in the target domain. We mainly focus on analyzing the population loss for simplicity, but in our main results (Theorems 3.1 and 3.2) we also give finite-sample guarantees. We analyze the following equivalent algorithms for self-training.

Entropy minimization. We initialize w from a source classifier w^{S} and run projected gradient descent on the entropy objective:²

$$w^0 = w^{\text{S}} \text{ and } w^{t+1} = \frac{w^t - \eta \nabla L(w^t)}{\|w^t - \eta \nabla L(w^t)\|_2} \quad (2.3)$$

²We project to the unit ball for simplicity, as the loss ℓ_{exp} is not scale-invariant.




Figure 1: Cases where entropy minimization fails to remove w_2 . **Left (Example 3.1):** When $w_1 = 0$, $w^\top x$ is distributed as a Gaussian. Increasing w_2 , which increases this variance (e.g, going from the purple to green curve) decreases $L(w)$ by forcing every prediction further from 0. This means that entropy minimization causes reliance on the spurious feature to *increase*. **Right (Example 3.2):** Distribution of $w_1^\top x_1$ in a hard case for general distributions. If there is a lot of mass of $w_1^\top x_1$ concentrated near the boundaries (i.e, $\pm\mu$ for some large $\mu \rightarrow \infty$) and a small amount of mass near 0, the loss could be small but the classifier will not want to shrink $\|w_2\|_2$.

Pseudo-labeling. We consider a variant of pseudo-labeling where we label the target data using the classifier from the previous iteration and run projected gradient descent on the supervised loss

$$L_{pseudo}^{t+1}(w) \triangleq \mathbb{E}_{x \sim \mathcal{D}_{tg}} \ell_{exp}(w^\top x, y^t) \quad (2.4)$$

where $y^t = \text{sgn}(w^{t\top} x)$ and $\ell_{exp}(t, y) = \exp(-ty)$. The algorithm is the same as 2.3 with $L(w)$ replaced by $L_{pseudo}^{t+1}(w)$. Note that this is different from some versions of pseudo-labeling, which train for many rounds of gradient descent before re-labeling. We observe that the two algorithms above are equivalent because the iterates are the same (see Section D.3 for the formal proof).

3 Overview of Main Results

We would like to show that entropy minimization (2.3) drives the spurious feature w_2 to 0. However, this is somewhat surprising and challenging to prove because nothing in the loss or algorithm explicitly enforces a decrease in $\|w_2\|_2$. Indeed, without additional assumptions on the target distribution \mathcal{D}_{tg} and the initial source classifier w^S , we show that entropy minimization can actually cause $\|w_2\|_2$ to increase because self-training can reinforce existing biases in the source classifier.

Examples 3.1 and 3.2 highlight cases where entropy minimization can fail, which motivates our assumptions of separation (Assumption 3.1) and that x_1 is a mixture of *sliced log concave* distributions. Under these assumptions, our main Theorem 3.1 shows that entropy minimization (2.3) initialized with a decently accurate source classifier drives the coefficient of the spurious feature, w_2 , to 0. For a simpler Gaussian setting, Theorem 3.2 shows that entropy minimization with a sufficiently accurate source classifier converges to the Bayes optimal classifier.

3.1 Failure cases of self-training

We highlight cases where self-training increases reliance on the spurious features, justifying our assumptions in Section 3.2.

Example 3.1 (No contribution from signal, i.e. $w_1^\top x_1 = 0$). *See Figure 1 (Left). For simplicity, suppose that $d_1 = d_2 = 1$, and suppose that $w_1 = 0$, so the signal feature is not used. In this case, increasing $|w_2|$ drives every prediction further from 0, decreasing the expected loss $L(w)$. Thus, in this example the min-entropy loss actually encourages the weight on the spurious feature, $|w_2|$, to increase. Note that this is not trivially true when w_1 is nonzero.*

In a realistic scenario, it's unlikely that $w_1^\top x_1 = 0$ for all examples because then the source accuracy on the target domain is very poor. So a priori, if we assume the source accuracy is decent (which implies $L(w^S)$ is small), we may avoid the pathological case above. However, this is not sufficient.

Example 3.2 (Initial $L(w^S)$ is small, but self-training still increases $\|w_2\|_2$). *See Figure 1 (Right). Suppose that restricting to the signal feature, we have a mixture of perfectly and extremely confidently predicted examples, and a small amount of unconfident examples as in Example 3.1. The majority group of confident examples is already perfectly predicted with no incentive to remove w_2 (because the loss gradient is near 0), and the minority group encourages $\|w_2\|_2$ to increase as in Example 3.1, so the overall effect is for $\|w_2\|_2$ to increase though $L(w)$ is small.*

For self-training to succeed, the correctly and confidently predicted examples must help remove the spurious features. As demonstrated above, this requires some continuum between confidently and unconfidently predicted examples. This motivates the log-concavity and smoothness assumptions, which guarantees that the sample distribution is not supported on too many extremely isolated clusters.

3.2 Mixtures of log-concave and log-smooth distributions

To avoid the failure cases above, we make realistic assumptions which are plausible in real-world data distributions. We start by defining a variant of log-concave and log-smooth distributions.

Definition 3.1 (sliced log-concavity, log-smoothness). *A distribution over \mathbb{R}^d with density p is α -log-concave for $\alpha > 0$ if $\nabla^2 \log p(t) \preceq -\alpha \cdot I_{d \times d}$, and is β -log-smooth if $\|\nabla^2 \log p(t)\|_{\text{op}} \leq \beta$. A distribution p over \mathbb{R}^d is sliced α -log-concave or sliced β -log-smooth if for any unit vector v , the random variable $v^\top x$ with $x \sim p$ is α -log-concave or β -log-smooth, respectively.*

A 1 dimensional density that is not Gaussian which satisfies these assumptions is $p(x) \propto \exp(-x^2 + \cos x)$. This density is 1-log concave and 3-log smooth. Now we state our main assumption that x_1 consists of a mixture of sliced-log-concave and smooth distributions with sufficient separation.

Assumption 3.1 (Separation assumption on the data). *We assume that the distribution of x_1 in the target domain, denoted by $\mathcal{D}_{\text{tg},1}$, is a mixture of K sliced α -log-concave and β -log-smooth distributions. (The reader can think of α , β and K as absolute constants for simplicity.) Let τ_1, \dots, τ_K denote the probability of each mixture and $\tau = \min_i \tau_i$. We assume that these mixtures are sufficiently separated in the sense that for scalar κ (that depends on α and β), there exists $(w_1, 0) \in \mathbb{R}^{d_1+d_2}$ such that $L((w_1, 0)) \leq \tau\kappa$.*

We formally define κ in Section B.1. When α and β are of constant scale, κ is also a constant. Assumption 3.1 is always satisfied if the means of each mixture distribution in $\mathcal{D}_{\text{tg},1}$ are sufficiently bounded from 0. We can see why Assumption 3.1 is a separation condition by considering the case when there exists $(w_1, 0)$ with good classification accuracy on x_1 . Obtaining good classification accuracy is only possible if the means of different classes are sufficiently far from 0 and also on opposite sides of 0, resulting in separation between the two classes.

The sliced log-concavity ensures that each mixture component of $w_1^\top x_1$ is uni-modal, with upper bound α on its “width”. Likewise, the sliced log-smoothness condition ensures that each component is not too narrow. These conditions rule out the hard distribution in Figure 1 (right), as each of the three components change quite sharply, violating log-smoothness. Next, we assume the source classifier is decently accurate and has bounded usage of the spurious feature.

Assumption 3.2 (Source classifier is decently accurate, doesn’t rely too much on spurious features.). *We assume that the source classifier has a significant mass in the space of the signal x_1 in the sense that $\|w^S_1\|_2 \geq 1/2$. We further assume that either Σ_2 is sufficiently small or w^S_2 , the initialization in the spurious feature space, is sufficiently small, in the sense that*

$$\sigma^2 = w^S_2^\top \Sigma_2 w^S_2 \leq c \cdot \min\{1, \alpha/\beta^2, \beta^{-1}, (\beta/|\log \beta|)^{-1}\}$$

for some sufficiently small universal constant c (e.g., $c = 0.03$ can work.) Furthermore, we assume the source classifier w^S has small entropy bound $L(w^S) \leq \tau\kappa$, where κ is the constant in Assumption 3.1.

The conditions on $\|w^S_1\|_2$ and σ can be satisfied if w^S_2 is not too large. The following theorem shows that under our assumptions, entropy minimization succeeds in removing the spurious w_2 .

Theorem 3.1 (Main result). *In the setting above, suppose Assumptions 3.1 and 3.2 hold and L is smooth.³ If we run Algorithm (2.3) initialized with w^S with sufficiently small step size η , after*

³As there is a discontinuity in $\frac{d}{dt} \ell_{\text{exp}}(t)$ at $t = 0$, we need to assume smoothness. This regularity condition is easy to satisfy; for example, it holds if $\mathcal{D}_{\text{tg},1}$ is a mixture of Gaussians.

$O(\log \frac{1}{\epsilon})$ iterations, we will obtain \hat{w} with very small usage of spurious features, i.e. $\|\hat{w}_2\|_2 \leq \epsilon$. The same conclusion holds with probability $1 - \delta$ in the finite sample setting with $O(\frac{1}{\epsilon^4} \log \frac{1}{\delta})$ unlabeled samples from \mathcal{D}_{tg} .

Above, the notation $O(\cdot)$ hides dependencies on α, β, Σ_2 , and other distribution-dependent parameters. The interpretation of Theorem 3.1 is that self-training can take a source classifier that has decent accuracy on the target and de-noise it completely, improving target accuracy by removing spurious extrapolations. The main proof of Theorem 3.1 is given in Section B. We provide proof intuitions in Section 4. In Section B.4, we show we can ensure convergence to an approximate local minimum of the objective $\min_{\|w_1\|_2 \leq 1} L((w_1, 0))$ by adding Gaussian noise to the gradient updates.

Special case: mixtures of Gaussians. We provide a slightly stronger analysis of Algorithm 2.3 when the signal x_1 is a one-dimensional Gaussian mixture, i.e. we set $\mathcal{D}_{\text{tg},1}(\cdot|y) = \mathcal{N}(y\gamma, \sigma_1^2)$. Let

$\tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}$ denote the minimum and maximum eigenvalues of $\tilde{\Sigma} \triangleq \begin{pmatrix} \sigma_1^2 & 0 \\ 0 & \Sigma_2 \end{pmatrix}$.

We analyze a slightly more general variant of Algorithm 2.3 which projects to the R -norm ball rather than unit ball. We now show that starting from an initial source classifier with sufficiently high accuracy on the target domain, self-training will avoid using the spurious feature and converge to the Bayes optimal classifier. Note that this is a stronger statement than Theorem 3.1, which does not bound the final target accuracy of the classifier.

Theorem 3.2. *In the setting above, suppose we are given a classifier (trained on a source distribution) w^S with $\|w^S\| \leq R$ and 0-1 error on the target domain at most $\rho = \frac{1}{2} \operatorname{erfc}\left(\frac{r(R\tilde{\sigma}_{\max})}{\sqrt{2R\tilde{\sigma}_{\min}}}\right)$.⁴ (r is a function as defined in Section A). Then Algorithm 2.3 converges to w^K satisfying $w_1^K \geq \sqrt{R^2 - \epsilon^2}$ and $\|w_2^K\|_2 \leq \epsilon$ within $K = O(\log \frac{1}{\epsilon})$ iterations. For the finite sample setting, the same conclusion holds with probability $1 - \delta$ using $O(\frac{1}{\epsilon^4} \log \frac{1}{\delta})$ samples.*

As above, $O(\cdot)$ hides dependencies in $R, \tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}, \rho, d_2$. In particular, w converges to $(R, 0)$, the classifier in $\{w : \|w\|_2 \leq R\}$ with the best possible accuracy. The full proof is in Section A.

4 Overview of Analysis

We will summarize the key intuitions for proving Theorems 3.1 and Theorem 3.2. The main ingredient is to show that the min entropy objective encourages a decrease in $\|w_2\|_2$, as stated below:

Lemma 4.1. *In the setting of Theorem 3.2, suppose that the classifier w has at most ρ error on the target. Then $\langle \nabla_{w_2} L(w), w_2 \rangle \geq 0$. This same conclusion holds in the setting of Theorem 3.1 for any w satisfying the conditions in Assumption 3.2.*

The consequence of Lemma 4.1 is that one step of gradient descent on the loss function $L(w)$ shrinks the norm of w_2 . This leads to the conclusion of Theorems 3.1 and 3.2, modulo a few other nuances such as showing that the conditions of Lemma 4.1 hold for all the iterates, which is done inductively. We also show that $\|w_1\|_2$ increases after one gradient step (Lemma A.2), so the norm of w_2 still decreases after re-normalization. To prove Lemma 4.1, we first express the objective as follows:

$$L(w) = \mathbb{E}_{x_1} \left[\mathbb{E}_{x_2} \ell_{\text{exp}}(w_1^\top x_1 + w_2^\top x_2) \right] \quad (4.1)$$

Note that $w_2^\top x_2$ has Gaussian distribution with mean zero and variance $\sigma^2 \triangleq w_2^\top \Sigma_2 w_2$. Let $g_\sigma(t) = \mathbb{E}_{z \sim \mathcal{N}(0, \sigma^2)} [\ell(t + z)]$ denote the convolution of ℓ_{exp} with $\mathcal{N}(0, \sigma^2)$. Then we can rewrite the loss as $L(w) = \mathbb{E}_{w_1^\top x_1} [g_\sigma(w_1^\top x_1)]$, so $\nabla_{w_2} L(w) = \frac{\partial L(w)}{\partial \sigma} \cdot \frac{\partial \sigma}{\partial w_2} = \frac{\partial L(w)}{\partial \sigma} \cdot 2\Sigma_2 w_2$, which implies $\langle \nabla_{w_2} L(w), w_2 \rangle = 2w_2^\top \Sigma_2 w_2 \frac{\partial}{\partial \sigma} L(w)$. As $\Sigma_2 \succ 0$, proving Lemma 4.1 is equivalent to proving $\frac{\partial}{\partial \sigma} L(w) \geq 0$. Letting $\mu \triangleq w_1^\top x_1$, we have $\frac{\partial}{\partial \sigma} L(w) = \mathbb{E}_\mu [q_\sigma(\mu)]$ where $q_\sigma(\mu) = \frac{\partial}{\partial \sigma} g_\sigma(\mu)$. We now investigate when $q_\sigma(\mu) \geq 0$. As visualized in Figure 2, $q_\sigma(\mu) < 0$ for μ near 0.

Case when $\mu \gg \sigma$: Recall that $g_\sigma(\mu) = \mathbb{E}_{z \sim \mathcal{N}(\mu, \sigma^2)} [\ell_{\text{exp}}(z)]$ is the average of the entropy function over a Gaussian distribution. When μ is sufficiently large, most of the mass of the Gaussian

⁴ $\operatorname{erfc}(t) = \frac{2}{\pi} \int_t^\infty \exp(-x^2) dx$.

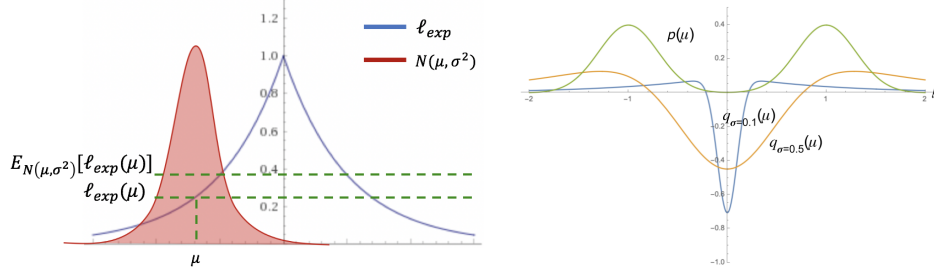


Figure 2: Analyzing dependence of L on σ . **Left:** $\mu \gg \sigma$. A visual depiction of why $q_\sigma(\mu) > 0$ when $\mu \gg \sigma$. Conditioned on $\mu = w_1^\top x_1$, $w^\top x$ is Gaussian with mean μ . As $\mu \gg \sigma$, most of its mass is in the region where ℓ_{exp} is convex. By Jensen’s inequality, driving σ to 0 decreases the loss in this region. **Right: plot of $q_\sigma(\mu)$.** The function $q_\sigma(\mu)$ will be convolved with p , the distribution over μ . To guarantee $\mathbb{E}[q_\sigma(\mu)] \geq 0$, we would like to μ has as large amount of mass right to the positive root and left to the negative root of $q_\sigma(\cdot)$ as possible.

distribution $\mathcal{N}(\mu, \sigma^2)$ is on the positive side of the real line, where the function ℓ_{exp} is convex. For convex functions f , Jensen’s inequality tells us $\mathbb{E}_{z \sim \mathcal{N}(\mu, \sigma^2)}[f(\mu)] > f(\mu)$ if $\sigma > 0$. As decreasing σ decreases the expected loss, we can see that $q_\sigma(\mu) > 0$. This is visualized in Figure 2 (left).

General case: We plot $q_\sigma(\mu)$ as a function of μ for various choices of σ in Figure 2 (right). We can see from the figure that for any $\sigma > 0$, there is a threshold $r(\sigma)$ (defined in (A.1)) such that for any $|\mu| > r(\sigma)$, $q_\sigma(\mu) > 0$. In Lemma A.4, we bound this value $r(\sigma)$ in terms of σ .

For the Gaussian setting, we can compute $\mathbb{E}_\mu[q_\sigma(\mu)]$ exactly and show it is positive for sufficiently accurate classifiers. For the general case (Theorem 3.1), it is difficult to bound this expectation because the expression for $q_\sigma(\mu)$ is complicated. Intuitively, our argument for why $\mathbb{E}_\mu[q_\sigma(\mu)] > 0$ is as follows: log-concavity and smoothness of each mixture component in μ ensures that the densities are uni-modal and do not change too fast. Thus, when $L(w)$ is sufficiently small, the mass of each component is spread over the real line, with most of the mass in the middle where $q_\sigma(\mu)$ is significantly positive, guaranteeing $\mathbb{E}_\mu[q_\sigma(\mu)] > 0$. To formalize this, we use a second order Taylor expansion of the log density of μ and bound the error incurred by the expansion using smoothness and concavity. This analysis is presented in Section B.

In Section C, we prove finite sample guarantees by showing that the gradient updates on the population and sample loss are similar ($\nabla \hat{L}(w) \approx \nabla L(w)$ for all w , where $\hat{L}(w)$ is the empirical loss).

5 Experiments

We validate our theory in a variety of empirical settings. We study a more general setting with nonlinear models where the signal x_1 and spurious feature x_2 are not distinct dimensions of the data. Using a semi-synthetic colored MNIST dataset, we verify that 1. self-training avoids using spurious features in a manner consistent with our theory and 2. as our theory predicts, self-training can harm performance when the source classifier is not sufficiently accurate. We also confirm our theoretical conclusions on a celebA dataset modified to have spurious correlations in training data but not in test.

Next, we investigate the connection between entropy minimization (2.3) and a variant of pseudo-labeling (2.4). We demonstrate that entropy minimization can converge to better target accuracy within a fixed wall clock time-budget, suggesting that practitioners may benefit from pseudo-labeling with more rounds and fewer epochs per round (Section E.3). In Section E.4 we verify that our conclusions hold for more common variants of pseudo-labeling in a toy Gaussian setting.

Colored MNIST. We create colored variants of the MNIST dataset [19] where the shape of the digit is the signal feature and the color is the spurious feature. In the first variant, denoted CMNIST10, there are 10 classes. Color correlates with the label in the source with probability $p = 0.95$, but is uncorrelated with the label in the target. In the second variant, denoted CMNIST2, we group digits into two classes: 0-4 and 5-9, which allows detailed investigation of our theory. Color correlates with the class label in the source but not in the target using a construction described in detail in Section E.1. We train 3-layer feed-forward network on the source, and use this to initialize entropy minimization (Algorithm 2.3) on unlabeled target data. Evaluation is performed on held-out target samples.

Table 1: Accuracy of models on the target before/after self-training, demonstrating that self-training can boost target accuracy under our structured domain shift. The exception is CMNIST10 with 0.97 probability of correlation between color and class. Here self-training decreases accuracy because initial accuracy is poor (only 72%), justifying our assumption of a decently accurate source classifier.

	CELEBA	CMNIST10 (P = 0.95)	CMNIST2	CMNIST10 (P = 0.97)
TRAINED ON SOURCE	81%	82%	94%	72%
AFTER SELF-TRAINING	88%	91%	96%	67%

Table 2: Number of test examples explainable by our theory. See text for definitions and interpretation.

	-/+	+/-	+/+	-/-	TOTAL
EXPLAINABLE	271	45	8785	150	9251
TOTAL	349	86	9286	279	10000

CelebA dataset. Inspired by [15], we partition the celebA dataset [21] so that gender correlates perfectly with hair color in source data (Figure 7a) but not in the target (Figure 7b). We train a neural net to predict gender by first training on source data alone and then performing self-training with unlabeled target data. During self-training, we add the labeled source loss to the min-entropy loss on target data. (Section E.2 has more details.)

Self-training improves target accuracy. Table 1 shows that with a decently accurate source classifier, self-training on unlabeled target data leads to substantial improvements in the target domain. For example, on celebA the classifier achieves 81% accuracy before self-training and 88% after. This suggests that practitioners can potentially avoid overfitting to spurious correlations by self-training on large unlabeled datasets in the target domain.

Self-training requires decent source classifier accuracy to succeed. We test whether self-training is effective when the source classifier is bad by increasing the correlation between label and spurious color feature from 0.95 to 0.97 for CMNIST10. The resulting source classifier only obtains 72% initial accuracy on target data, which *drops* to 67% after self-training (see Table 1, last column, and plots in Section E.1). This shows that our assumption that the source classifier has to obtain non-trivial target accuracy (with bounded usage of the spurious feature) is also necessary in practice.

Self-training reduces reliance on the spurious features. In the CelebA experiment, test predictions corrected by self-training were mostly mistaken due to the spurious correlation. Figure 7c, a random sample of the corrected examples, consists of mostly blond females, non-blond males, and subjects with hats or irregular hairstyles.

For 2-class colored MNIST, let $\mu_S(x_1)$, $\sigma_S(x_1)$ denote the mean and standard deviation of the source classifier conditioned on grayscale image x_1 , with color distributed independently of x_1 . Define $\mu_T(x_1)$, $\sigma_T(x_1)$ similarly for the classifier after self-training. Our theory suggests that $\text{sgn } \mu_S(x_1) = \text{sgn } \mu_T(x_1)$, $|\mu_S(x_1)| < |\mu_T(x_1)|$, and $\sigma_S(x_1) > \sigma_T(x_1)$, and we say a test example is *explainable* by our theory if this holds. We divide the test examples into four categories: “-/+”, “+/-”, “+/+”, “-/-”, where, for example, “-/+” indicates source classifier was wrong but corrected by self-training. Table 2 summarizes the number of explainable examples in each category, showing that for the majority (> 90%) of examples, entropy minimization works due to the reason we hypothesized. In Section E.1, we provide additional detailed analyses of the influence of the spurious color feature on the prediction before and after self-training.

6 Conclusion

We study the impact of self-training under domain shift. We show that when there are spurious correlations in the source domain which are not present in the target, self-training leverages the unlabeled target data to avoid relying on these spurious correlations. Our analysis highlights several conditions for self-training to work in theory, such as good separation between classes and a decently accurate source classifier. Our experiments support that 1) these theoretical requirements can capture the initial conditions needed for self-training to work and 2) under these initial conditions, self-training can indeed prevent the model from using spurious features in ways predicted by our theory. It is an interesting question for future work to explore other settings we can analyze with our framework.

Broader Impact

Our work promotes robustness and fairness in machine learning. First, we study algorithms that make machine learning models robust when deployed in the real world. Second, our work addresses the scenario where the target domain is under-resourced and hence collecting labels is difficult. Third, our theoretical work guides efforts to mitigate dataset bias. We demonstrate that curating a diverse pool of unlabeled data from the true population can help combating existing bias in labeled datasets. We give conditions for when bias will be mitigated and when it will be reinforced or amplified by popular algorithms used in practice. We take a first step towards understanding and preventing the adverse effects of self-training.

Acknowledgments and Disclosure of Funding

We are grateful to Rui Shu, Shiori Sagawa, and Pang Wei Koh for insightful discussions. The authors would like to thank the Stanford Graduate Fellowship program for funding. CW acknowledges support from a NSF Graduate Research Fellowship. TM is also partially supported by the Google Faculty Award, Stanford Data Science Initiative, and the Stanford Artificial Intelligence Laboratory.

References

- [1] N. Agarwal, Z. Allen-Zhu, B. Bullins, E. Hazan, and T. Ma. Finding approximate local minima faster than gradient descent. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1195–1199, 2017.
- [2] M. Amini and P. Gallinari. Semi-supervised learning with explicit misclassification modeling. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2003.
- [3] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz. Invariant risk minimization, 2019.
- [4] M.-F. Balcan and A. Blum. A discriminative model for semi-supervised learning. *Journal of the ACM (JACM)*, 57(3):1–46, 2010.
- [5] S. Ben-David, T. Lu, and D. Pal. Does unlabeled data provably help? worst-case analysis of the sample complexity of semi-supervised learning. In *Conference on Learning Theory (COLT)*, 2008.
- [6] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [7] D. Berthelot, N. Carlini, E. D. Cubuk, A. Kurakin, K. Sohn, H. Zhang, and C. Raffel. Remix-match: Semi-supervised learning with distribution matching and augmentation anchoring. In *International Conference on Learning Representations*, 2020.
- [8] A. Blum and T. Mitchell. Combining labeled and unlabeled data with co-training. In *Conference on Learning Theory (COLT)*, 1998.
- [9] Y. Carmon, A. Raghunathan, L. Schmidt, P. Liang, and J. C. Duchi. Unlabeled data improves adversarial robustness. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [10] Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. In *International Conference on Machine Learning (ICML)*, pages 1180–1189, 2015.
- [11] R. Ge, F. Huang, C. Jin, and Y. Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842, 2015.
- [12] B. Gong, F. Sha, and K. Grauman. Overcoming dataset bias: An unsupervised domain adaptation approach. In *In NIPS Workshop on Large Scale Visual Recognition and Retrieval*, 2012.
- [13] Y. Grandvalet and Y. Bengio. Entropy regularization. In *Semi-Supervised Learning*, 2005.
- [14] S. Gururangan, S. Swayamdipta, O. Levy, R. Schwartz, S. Bowman, and N. A. Smith. Annotation artifacts in natural language inference data. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, June 2018.
- [15] C. Heinze-Deml and N. Meinshausen. Conditional variance penalties and domain shift robustness. *arXiv preprint arXiv:1710.11469*, 2017.

- [16] A. Ilyas, S. Santurkar, D. Tsipras, L. Engstrom, B. Tran, and A. Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [17] H. Jiayuan, S. A. J., G. Arthur, B. K. M., and S. Bernhard. Correcting sample selection bias by unlabeled data. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2006.
- [18] A. Kumar, T. Ma, and P. Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning (ICML)*, 2020.
- [19] Y. LeCun, C. Cortes, and C. Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [20] D. Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *International Conference on Machine Learning (ICML) Workshop*, 2013.
- [21] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [22] M. Long, J. Wang, G. Ding, J. Sun, and P. S. Yu. Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE international conference on computer vision*, pages 2200–2207, 2013.
- [23] T. McCoy, E. Pavlick, and T. Linzen. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, July 2019.
- [24] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii. Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 41(8):1979–1993, 2018.
- [25] Y. Nesterov and B. T. Polyak. Cubic regularization of newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [26] J. Peters, P. Bühlmann, and N. Meinshausen. Causal inference using invariant prediction: identification and confidence intervals, 2015.
- [27] P. Rigollet. Generalization error bounds in semi-supervised classification under the cluster assumption. *Journal of Machine Learning Research (JMLR)*, 8:1369–1392, 2007.
- [28] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry. Adversarially robust generalization requires more data. In *Advances in Neural Information Processing Systems*, pages 5014–5026, 2018.
- [29] H. Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90:227–244, 2000.
- [30] R. Shu, H. H. Bui, H. Narui, and S. Ermon. A DIRT-T approach to unsupervised domain adaptation. In *International Conference on Learning Representations (ICLR)*, 2018.
- [31] A. Singh, R. Nowak, and J. Zhu. Unlabeled data: Now it helps, now it doesn’t. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2008.
- [32] K. Sohn, D. Berthelot, C. Li, Z. Zhang, N. Carlini, E. D. Cubuk, A. Kurakin, H. Zhang, and C. Raffel. Fixmatch: Simplifying semi-supervised learning with consistency and confidence. *arXiv*, 2020.
- [33] D. Soudry, E. Hoffer, and N. Srebro. The implicit bias of gradient descent on separable data. In *International Conference on Learning Representations*, 2018.
- [34] M. Sugiyama, M. Krauledat, and K. Muller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research (JMLR)*, 8:985–1005, 2007.
- [35] T. Tommasi, N. Patricia, B. Caputo, and T. Tuytelaars. A deeper look at dataset bias. In *Domain adaptation in computer vision applications*, pages 37–55. Springer, 2017.
- [36] D. Tsipras, S. Santurkar, L. Engstrom, A. Turner, and A. Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018.
- [37] E. Tzeng, J. Hoffman, N. Zhang, K. Saenko, and T. Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- [38] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell. Adversarial discriminative domain adaptation. In *Computer Vision and Pattern Recognition (CVPR)*, 2017.

- [39] J. Uesato, J. Alayrac, P. Huang, R. Stanforth, A. Fawzi, and P. Kohli. Are labels required for improving adversarial robustness? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [40] R. Vershynin. *High-Dimensional Probability*. Cambridge University Press, 2018.
- [41] E. W. Weisstein. *Erfc*. From MathWorld—A Wolfram Web Resource., 2020. <http://mathworld.wolfram.com/Erfc.html>.
- [42] Q. Xie, M. Luong, E. Hovy, and Q. V. Le. Self-training with noisy student improves imagenet classification. *arXiv*, 2020.
- [43] Y. Zhang, T. Liu, M. Long, and M. I. Jordan. Bridging theory and algorithm for domain adaptation. *arXiv preprint arXiv:1904.05801*, 2019.
- [44] Y. Zou, Z. Yu, X. Liu, B. Kumar, and J. Wang. Confidence regularized self-training. *arXiv preprint arXiv:1908.09822*, 2019.

A Warmup: Proofs for Gaussian Setting (Theorem 3.2)

We will define the function r in Theorem 3.2 as follows:

$$r(\sigma) = \begin{cases} \sigma^2 + \sigma \sqrt{2 \log \frac{4\sqrt{2}}{\sqrt{\pi}\sigma}}, & \text{if } 0 < \sigma \leq \frac{4\sqrt{2}}{\sqrt{\pi}}. \\ 2\sigma^2, & \text{if } \sigma > \frac{4\sqrt{2}}{\sqrt{\pi}}. \end{cases} \quad (\text{A.1})$$

The algorithm that we consider is a variant of Algorithm 2.3 which more generally projects to the R -norm ball rather than unit ball: $w^{t+1} = R_{\frac{w^t - \eta \nabla L(w^t)}{\|w^t - \eta \nabla L(w^t)\|}}$. Now define

$$\begin{aligned} a &= \sqrt{2} R \tilde{\sigma}_{\min} \operatorname{erfc}^{-1}(2\rho) = r(R \tilde{\sigma}_{\max}) \\ S &= \{w : w_1 \gamma \geq a, \|w\|_2 \leq R\}. \end{aligned} \quad (\text{A.2})$$

where the function r is defined in Lemma A.4. We first observe that the condition that classifier w has at most $1 - \rho$ error corresponds directly to $w \in S$.

Lemma A.1. *In the setting of Theorem 3.2, suppose some classifier w has at least $1 - \rho$ accuracy in the sense that $\Pr_{x,y} [\operatorname{sgn}(w^\top x) = y] \geq 1 - \rho$ and $\|w\|_2 \leq R$. Then $w_1 \gamma \geq a$.*

Proof of Lemma A.1. Let $\tilde{\sigma}^2 = w^\top \tilde{\Sigma} w$. We have $R \tilde{\sigma}_{\min} \leq \tilde{\sigma} \leq R \tilde{\sigma}_{\max}$.

$$\begin{aligned} & \Pr_{y \sim \{\pm 1\}, z \sim \mathcal{N}(\tilde{0}, \tilde{\sigma}^2)} [\operatorname{sgn}(y w_1 \gamma + z) = y] \geq 1 - \rho \\ \iff & \Pr_{z \sim \mathcal{N}(\tilde{0}, \tilde{\sigma}^2)} [w_1 \gamma + z \geq 0] \geq 1 - \rho \\ \iff & \operatorname{erfc}\left(\frac{w_1 \gamma}{\sqrt{2} \tilde{\sigma}}\right) \leq 2\rho \\ \implies & w_1 \gamma \geq a \end{aligned}$$

□

Next, our proof of Theorem 3.2 will be based on the following two lemmas. The first lemma shows that w_1 is increasing, and the second shows that $\|w_2\|_2$ is decreasing at a fast enough rate.

Lemma A.2. *In the setting of Theorems 3.1 and 3.2, for any $w \in S$,*

$$\langle \nabla_{w_1} L(w), w_1 \rangle < 0.$$

Proof. Recall the definition

$$g_\sigma(\mu) = \mathbb{E}_{z \sim \mathcal{N}(\mu, \sigma^2)} [\ell_{\text{exp}}(z)]$$

Now we can compute

$$\begin{aligned} \langle \nabla_{w_1} L(w), w_1 \rangle &= \mathbb{E}_{x_1} [g'_\sigma(w_1^\top x_1) x_1^\top w_1] \\ &= \mathbb{E}_{x_1} [g'_\sigma(w_1^\top x_1) w_1^\top x_1] \\ &< 0 \end{aligned}$$

since $g'_\sigma(\mu)$ and μ always have opposite signs. □

Lemma A.3. *In the same setting as in Lemma A.2, we have that for any $w \in S$,*

$$\langle \nabla_{w_2} L(w), w_2 \rangle \geq c \|w_2\|_2^2$$

for some constant c dependent only on $R, \gamma, \tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}$.

This lemma relies on the following bound stating that for $|\mu| > r(\sigma)$, $q_\sigma(\mu) > 0$.

Lemma A.4. *Define $r(\sigma)$ as in (A.1). Then for $|\mu| \geq r(\sigma)$, $q_\sigma(\mu) \geq \frac{1}{4} \sigma \ell_{\text{exp}}(\mu) > 0$.*

We prove this lemma in Section D.1. We also require the following claim that $q_{\tilde{\sigma}}(w_1 \gamma)$ is lower bounded by some positive constant for all $w \in S$.

Claim A.1. Define the function $r(\sigma)$ as in Lemma A.4. Suppose $a \leq \mu \leq R\gamma$. Then for all $w \in S$,

$$\frac{\partial g_{\tilde{\sigma}}(w_1\gamma)}{\partial \tilde{\sigma}} \geq c_1$$

for some constant c_1 dependent only on $R, \gamma, \tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}$.

Proof of Claim A.1. We note that r and q satisfy the following properties:

1. $r(\sigma)$ is a monotonically increasing function.
2. $q_{\sigma}(\mu) = \frac{\partial g_{\tilde{\sigma}}(w_1\gamma)}{\partial \tilde{\sigma}} > 0$ for all $\mu \geq r(\sigma)$. (See Lemma A.4 for proof.)

For arbitrary $\tilde{\sigma} \in [R\tilde{\sigma}_{\min}, R\tilde{\sigma}_{\max}]$, $\mu \geq a = r(R\tilde{\sigma}_{\max})$ ensures that $\mu \geq r(\tilde{\sigma})$ by property 1. By property 2, $q(\tilde{\sigma}, \mu) > 0$. Setting $c_1 = \min_{\tilde{\sigma} \in [R\tilde{\sigma}_{\min}, R\tilde{\sigma}_{\max}], \mu \in [a, R\gamma]} q_{\sigma}(\mu)$ finishes the proof. c_1 is dependent only on $R, \gamma, \tilde{\sigma}_{\min}, \tilde{\sigma}_{\max}$. \square

Proof of Lemma A.3. Using the fact that x_1 is a uni-variate mixture of Gaussians, and therefore $w^\top x$ is itself a mixture of two Gaussians with variance $\tilde{\sigma}^2$, we have

$$\begin{aligned} L(w) &= \mathbb{E}_{y \sim \{\pm 1\}, z \sim \mathcal{N}(0, \tilde{\sigma}^2)} [l_{\exp}(yw_1\gamma + z)] \\ &= \frac{1}{2} (g_{\tilde{\sigma}}(w_1\gamma) + g_{\tilde{\sigma}}(-w_1\gamma)) \\ &= g_{\tilde{\sigma}}(w_1\gamma) \end{aligned}$$

Now we differentiate with respect to w_2 to obtain

$$\begin{aligned} \nabla_{w_2} L(w) &= \frac{\partial L(w)}{\partial \tilde{\sigma}} \cdot \frac{\partial \tilde{\sigma}}{\partial w_2} \\ &= \frac{\partial g_{\tilde{\sigma}}(w_1\gamma)}{\partial \tilde{\sigma}} \cdot 2\Sigma_2 w_2 \end{aligned}$$

We now use the lower bound on $\frac{\partial g_{\tilde{\sigma}}(w_1\gamma)}{\partial \tilde{\sigma}}$ given by Claim A.1. This gives

$$\langle \nabla_{w_2} L(w), w_2 \rangle \geq 2c_1 \tilde{\sigma}_{\min} \|w_2\|_2^2.$$

Setting $c = 2c_1 \tilde{\sigma}_{\min}$ finishes the proof. \square

We can now complete the proof of Theorem 3.2.

Proof of Theorem 3.2. Define $\tilde{w}^{t+1} = w^t - \eta \nabla L(w^t)$. By Lemma A.1, $w_1^0 \gamma \geq a$. Note that by assumption $a > r(R\tilde{\sigma}_{\max})$. By Lemma A.2 and A.3, at iteration $t \geq 0$, taking constant step size η ,

$$\begin{aligned} |\tilde{w}_1^{t+1}| &> |w_1^t| \\ \|\tilde{w}_2^{t+1}\|_2^2 &= \|w_2^t - \eta \nabla_{w_2} L(w)|_{w=w^t}\|_2^2 \\ &= \|w_2^t\|_2^2 + \eta^2 \|\nabla_{w_2} L(w)|_{w=w^t}\|_2^2 \\ &\quad - 2\eta \langle \nabla_{w_2} L(w)|_{w=w^t}, w_2^t \rangle \end{aligned}$$

By Lemma A.3,

$$\nabla_{w_2} L(w) = q_{\sigma(w)}(\mu(w)) \cdot 2\Sigma_2 w_2$$

for some continuous function $q_{\sigma}(\mu)$ where $\sigma(w)^2 = w_2^\top \Sigma_2 w_2$, $\mu(w) = w_1\gamma$ over compact set S . Therefore q is bounded. Suppose $|q| \leq c_2$ for all $w \in S$, then

$$\|\nabla_{w_2} L(w)|_{w=w^t}\|_2^2 \leq 4c_2^2 \tilde{\sigma}_{\max}^2 \|w_2^t\|_2^2.$$

Therefore $\|\tilde{w}_2^{t+1}\|_2^2 \leq c' \|w_2^t\|_2^2$ for some constant $c' < 1$ for appropriate choice of η . As $|\tilde{w}_1^{t+1}| > |w_1^t|$, renormalization results in some constant factor decrease in $\|w_2^t\|_2$. Therefore $\|w_2^t\|_2^2 \leq \epsilon$ when $t \geq K = O(\log(\frac{1}{\epsilon}))$. \square

B Missing Proofs for Theorem 3.1

Define the following function depending on parameters ρ, ν defined later:

$$\tilde{\kappa}(\rho, \nu) \triangleq \min \left\{ \frac{\sqrt{\pi}}{4\sqrt{\rho}} (p^*(\rho, \nu))^{1-\frac{\nu}{4\rho}} \left(\frac{\nu}{2\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}}, \frac{\sqrt{\nu}}{8\sqrt{2\pi}(\sqrt{\rho} + \sqrt{2})} \exp \left(- \left(\frac{\sqrt{\rho} + 4}{2\sqrt{\nu}} \right)^2 \right) \right\} \quad (\text{B.1})$$

for

$$p^*(\rho, \nu) \triangleq \frac{\sqrt{\nu}}{2\sqrt{\pi}} \min \left\{ 1, \frac{\sqrt{\nu}}{\sqrt{\rho}} \left(\frac{\sqrt{\pi}}{44\sqrt{2\rho}} \right)^{\frac{8\rho}{\nu}} \right\}$$

Now we choose the constant κ in Assumptions 3.1 and 3.2 to be $\kappa(\beta, \alpha) \triangleq \min_{a \in [1, 4]} \tilde{\kappa}(a\beta, a\alpha)$.

Throughout the proof, we will use p to refer to the density of $\mu = w_1^\top x_1$, and define $s(\mu) \triangleq \frac{\partial}{\partial \mu} \log p(\mu) = \frac{p'(\mu)}{p(\mu)}$. Throughout our proofs, we use ν, ρ to refer to parameters such that p is ν -log-concave and ρ -log-smooth. By Assumption 3.1, we have that the density of $\frac{w_1^\top x_1}{\|w_1\|_2}$ is α -log-concave and β -log-smooth, so we can choose $\nu = \alpha/\|w_1\|_2^2$ and $\rho = \beta/\|w_1\|_2^2$ by the linear transformation formula of a probability density. We use $\sigma = w_2^\top \Sigma_2 w_2$ to be the variance of the output of the current classifier restricted to the spurious coordinates.

B.1 Proof overview

We will argue that if the initial conditions in Assumption 3.2 hold, then they will continue to hold throughout training. Furthermore, under these initial conditions, the loss gradient will force $\|w_2\|_2$ to decrease.

The following three lemmas which analyze a single update of the algorithm will form the main technical core of our proof. They will be used to show that when the loss is sufficiently small, the norm of w_2 is always decreasing. The first lemma states that if the loss is small, then p cannot have a large density at 0.

Lemma B.1. *In the setting of Theorem 3.1, suppose that $K = 1$. When $L(w) \leq \tilde{\kappa}(\rho, \nu)$, we must have*

$$L(w) \geq \frac{\sqrt{\pi}}{4\sqrt{\rho}} p(0)^{1-\frac{\nu}{4\rho}} \left(\frac{\sqrt{\nu}}{2\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} \quad (\text{B.2})$$

As a consequence, when $L(w) \leq \tilde{\kappa}(\rho, \nu)$, we have:

$$p(0) \leq p^*(\rho, \nu)$$

Next, observe that w_2 will decrease if $\frac{\partial}{\partial \sigma} L(w) > 0$. The following lemma lower bounds $\frac{\partial}{\partial \sigma} L(w) > 0$ in terms of $p(0)$, showing that if $p(0)$ is small, $\frac{\partial}{\partial \sigma} L(w)$ will be positive.

Lemma B.2. *In the setting of Lemma B.1, we have*

$$\frac{\partial}{\partial \sigma} L(w) \geq p(0)\sigma \left(\frac{\sqrt{\pi}}{11\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} - 2\sqrt{2} \max \left\{ 1, \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{8\rho}} \right\} \right) \quad (\text{B.3})$$

Finally, this next lemma combines the two lemmas above to show that when the loss is sufficiently small, w_2 is always shrinking.

Lemma B.3. *In the setting of Lemma B.1, when $L(w) \leq \tilde{\kappa}(\rho, \nu)$ for $\tilde{\kappa}(\rho, \nu)$ defined in (B.1), we have*

$$\frac{\partial}{\partial \sigma} L(w) \geq \sigma p(0)^{1-\frac{\nu}{4\rho}} \frac{\sqrt{\pi}}{22\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} > 0$$

Proof of Lemma B.3. Define $a_1 \triangleq \frac{\sqrt{\pi}}{11\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}}$, $a_2 \triangleq 2\sqrt{2} \max \left\{ 1, \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{8\rho}} \right\}$, so that the right hand side of (B.3) becomes $p(0)\sigma(a_1 - a_2)$.

Now we apply Lemma B.1 to conclude that when $L(w) \leq \tilde{\kappa}(\rho, \nu)$, $p(0) \leq p^*(\rho, \nu)$. Note that when $p(0) \leq \frac{\sqrt{\nu}}{2\sqrt{\pi}} \min \left\{ \left(\frac{\sqrt{\pi}}{44\sqrt{2\rho}} \right)^{\frac{4\rho}{\nu}}, \frac{\sqrt{\nu}}{\sqrt{\rho}} \left(\frac{\sqrt{\pi}}{44\sqrt{2\rho}} \right)^{\frac{8\rho}{\nu}} \right\}$, we must have $a_1 \geq 2a_2$ by the definitions of a_1, a_2 . Furthermore, the r.h.s. of this bound is lower-bounded by $p^*(\rho, \nu)$. As a result, when $p(0) \leq p^*$, by Lemma B.2, we have $\frac{\partial}{\partial \sigma} L(w) \geq p(0)\sigma^{\frac{a_1}{2}}$. Applying the definition of a_1 gives the desired result. \square

B.2 Proof of Lemmas B.1

The following claim will be useful for proving both Lemma B.1 and Lemma B.2.

Claim B.1. Recall that we defined $s(\mu) = \frac{\partial}{\partial \mu} \log p(\mu) = \frac{p'(\mu)}{p(\mu)}$. The following bound holds:

$$\int_0^\infty \exp \left(\frac{|s(0)|}{2} \delta - \frac{\rho}{2} \delta^2 \right) d\delta \geq \frac{\sqrt{\pi}}{\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}}$$

Proof of Claim B.1. From Lemma D.3, we start with

$$\int_0^\infty \exp \left(\frac{|s(0)|}{2} \delta - \frac{\rho}{2} \delta^2 \right) d\delta \geq \frac{\sqrt{\pi} \exp \left(\frac{s(0)^2}{4\rho} \right)}{\sqrt{\rho}}$$

Now we apply the lower bound $s(0)^2 \geq \nu \log \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)$ and obtain

$$\begin{aligned} \int_0^\infty \exp \left(\frac{|s(0)|}{2} \delta - \frac{\rho}{2} \delta^2 \right) d\delta &\geq \frac{\sqrt{\pi}}{\sqrt{\rho}} \exp \left(\frac{\nu}{4\rho} \log \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right) \right) \\ &\geq \frac{\sqrt{\pi}}{\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} \end{aligned}$$

\square

Our starting point is to first lower bound $L(w)$ in terms of $p(0)$ and $s(0)$.

Claim B.2. The following lower bound on the loss $L(w)$ holds:

$$L(w) \geq 0.25p(0) \int_0^\infty \exp \left((|s(0)| - 1)\delta - \frac{\rho}{2} \delta^2 \right) d\delta$$

Proof of Claim B.2. Without loss of generality, assume that $s(0) \geq 0$ (otherwise, by symmetry of ℓ_{exp} the same arguments hold). Then we have

$$\begin{aligned} L(w) &= \int_{-\infty}^\infty p(\delta) g_\sigma(\delta) d\delta \\ &\geq p(0) \int_{-\infty}^\infty \exp \left(s(0)\delta - \frac{\rho}{2} \delta^2 \right) g_\sigma(\delta) d\delta && \text{(by log-smoothness)} \\ &\geq 0.25p(0) \int_{-\infty}^\infty \exp \left(s(0)\delta - \frac{\rho}{2} \delta^2 \right) \ell_{exp}(\delta) d\delta && \text{(by Lemma D.1)} \\ &\geq 0.25p(0) \int_0^\infty \exp \left((|s(0)| - 1)\delta - \frac{\rho}{2} \delta^2 \right) d\delta \quad \text{(substituting } \ell_{exp}(\delta) = \exp(-\delta) \text{ for } \delta \geq 0) \end{aligned}$$

Now the loss is symmetric around 0, so the same argument would also work for $s(0) < 0$. Thus, we obtain the desired result. \square

Next, we argue that if $L(w)$ is bounded above by some threshold, then $s(0)$ will be large in absolute value.

Claim B.3. Suppose that our classifier w satisfies the following loss bound:

$$L(w) \leq \frac{\sqrt{\nu}}{8\sqrt{2\pi}(\sqrt{\rho} + \sqrt{2})} \exp\left(-\left(\frac{\sqrt{\rho} + 4}{2\sqrt{\nu}}\right)^2\right) \quad (\text{B.4})$$

Then $|s(0)| \geq \frac{\sqrt{\rho}}{2} + 2$.

Proof of Claim B.3. Assume for the sake of contradiction that $|s(0)| \leq \frac{\sqrt{\rho}}{2} + 2$. First, we consider the case when $s(0) \in [1, \frac{\sqrt{\rho}}{2} + 2]$. In this case, by Lemma D.3, we have

$$p(0) \geq \frac{\sqrt{\nu}}{2\sqrt{\pi}} \exp\left(-\left(\frac{\sqrt{\rho} + 4}{2\sqrt{\nu}}\right)^2\right)$$

Furthermore, in this case we also have $|s(0)| - 1 > 0$, so we can apply (D.7) from Claim D.2. Plugging into Claim B.2, we obtain

$$\begin{aligned} L(w) &\geq 0.25p(0) \frac{\sqrt{\pi} \exp\left(\frac{(|s(0)|-1)^2}{\rho}\right)}{\sqrt{\rho}} \\ &\geq \frac{\sqrt{\nu}}{8\sqrt{\rho}} \exp\left(-\left(\frac{\sqrt{\rho} + 4}{2\sqrt{\nu}}\right)^2\right) \end{aligned}$$

In the other case where $0 \leq |s(0)| \leq 1$, by Claim D.2 and Claim B.2, we first have

$$L(w) \geq 0.25 \frac{\sqrt{\pi}}{\sqrt{\rho}} p(0) \exp\left(\frac{(|s(0)|-1)^2}{\rho}\right) \left(\operatorname{erf}\left(\frac{s(0)-1}{\sqrt{\rho}}\right) + 1\right)$$

Now applying the lower bound on $p(0)$ from Lemma D.3, we have

$$L(w) \geq \frac{\sqrt{\nu}}{8\sqrt{\rho}} \exp(-s(0)^2/\nu) \exp\left(\frac{(|s(0)|-1)^2}{\rho}\right) \left(\operatorname{erf}\left(\frac{s(0)-1}{\sqrt{\rho}}\right) + 1\right)$$

Now by Claim D.3, we have

$$\operatorname{erf}\left(\frac{|s(0)|-1}{\sqrt{\rho}}\right) + 1 \geq \frac{1}{\sqrt{\pi}} \frac{\exp\left(-\frac{(|s(0)|-1)^2}{\rho}\right)}{\sqrt{2} + 2(1 - |s(0)|)/\sqrt{\rho}}$$

Thus, we have

$$\begin{aligned} L(w) &\geq \frac{\sqrt{\nu}}{8\sqrt{\pi}(\sqrt{2\rho} + 2(1 - |s(0)|))} \exp\left(-\frac{s(0)^2}{\nu}\right) \\ &\geq \frac{\sqrt{\nu}}{8(\sqrt{2\pi\rho} + 2\sqrt{\pi})} \exp(-1/\nu) \end{aligned}$$

Combining the two cases allows us to conclude that if $|s(0)| < \frac{\sqrt{\rho}}{2} + 2$, the loss must satisfy

$$L(w) \geq \min\left\{\frac{\sqrt{\nu}}{8\sqrt{\rho}} \exp\left(-\left(\frac{\sqrt{\rho} + 4}{2\sqrt{\nu}}\right)^2\right), \frac{\sqrt{\nu}}{8(\sqrt{2\pi\rho} + 2\sqrt{\pi})} \exp(-1/\nu)\right\}$$

Now we note that the r.h.s. of the above equation is lower bounded by the r.h.s of (B.4). Thus, the loss would violate (B.4), a contradiction. \square

Proof of Lemma B.1. First, by Claim B.3, when $L(w) \leq \tilde{\kappa}(\rho, \nu)$, we must have $|s(0)| \geq \frac{\sqrt{\rho}}{2} + 2$. Now we lower bound the loss in terms of $p(0)$. Starting from Claim B.2, we have

$$\begin{aligned} L(w) &\geq 0.25p(0) \int_0^\infty \exp\left((|s(0)|-1)\delta - \frac{\rho}{2}\delta^2\right) d\delta \\ &\geq 0.25p(0) \int_0^\infty \exp\left(\frac{|s(0)|}{2}\delta - \frac{\rho}{2}\delta^2\right) d\delta \quad (\text{since } |s(0)| \geq 2) \end{aligned}$$

Now applying Claim B.1, we obtain

$$\begin{aligned} L(w) &\geq 0.25 \frac{\sqrt{\pi}}{\sqrt{\rho}} p(0) \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} \\ &\geq 0.25 \frac{\sqrt{\pi}}{\sqrt{\rho}} p(0)^{1-\frac{\nu}{4\rho}} \left(\frac{\sqrt{\nu}}{2\sqrt{\pi}} \right)^{\frac{\nu}{4\rho}} \end{aligned}$$

This completes the first part of the lemma. For the second part, we note that if $L(w) \leq \tilde{\kappa}(\rho, \nu)$, then $L(w)$ is bounded above by the r.h.s. of (B.2) with $p^*(\rho, \nu)$ substituted for $p(0)$ by the definition of $\tilde{\kappa}(\rho, \nu)$. Combined with the first part of the lemma, this immediately gives $p(0) \leq p^*(\rho, \nu)$. \square

B.3 Proof of Lemma B.2

We rely on the following lemma which lower bounds $\frac{\partial}{\partial \sigma} L(w)$.

Lemma B.4. *Suppose $\sigma \leq \frac{4\sqrt{2}}{\sqrt{\pi}}$ satisfies $\gamma_\sigma^* \leq \min \left\{ 1, \frac{1}{4\sqrt{\rho}} \right\}$ for $\gamma_\sigma^* \triangleq \sigma^2 + \sigma \sqrt{2 \log \frac{4\sqrt{2}}{\sqrt{\pi}\sigma}}$. Then the following lower bound on the derivative $\frac{\partial}{\partial \sigma} L(w)$ holds:*

$$\begin{aligned} \frac{\partial}{\partial \sigma} L(w) &\geq \\ p(0) &\left(\frac{\sigma}{11} \int_0^\infty \exp \left(\left(|s(0)| - \frac{\sqrt{\rho}}{4} - 1 \right) \delta - \frac{\rho}{2} \delta^2 \right) d\delta - 2\sqrt{2} \exp \left(\frac{s(0)^2}{\nu + \sigma^{-2}} \right) \frac{1}{\sqrt{\nu + \sigma^{-2}}} \right) \end{aligned} \quad (\text{B.5})$$

Proof of Lemma B.4. We compute $\frac{\partial}{\partial \sigma} L(w)$ in two parts:

$$\frac{\partial}{\partial \sigma} L(w) = \int_{|\mu| \leq \gamma_\sigma^*} p(\mu) \frac{\partial}{\partial \sigma} g_\sigma(\mu) + \int_{|\mu| > \gamma_\sigma^*} p(\mu) \frac{\partial}{\partial \sigma} g_\sigma(\mu) \quad (\text{B.6})$$

For $|\mu| \leq \gamma_\sigma^*$, we lower bound the integral using (D.2) in Lemma D.1. For $|\mu| > \gamma_\sigma^*$, we lower bound the integral using Lemma A.4. By (D.2), we have

$$\begin{aligned} \int_{|\mu| \leq \gamma_\sigma^*} p(\mu) \frac{\partial}{\partial \sigma} g_\sigma(\mu) &\geq -\sqrt{\frac{2}{\pi}} \int_{|\mu| \leq \gamma_\sigma^*} p(\mu) \exp \left(-\frac{\mu^2}{2\sigma^2} \right) d\mu \\ &\geq -\sqrt{\frac{2}{\pi}} \int_{|\delta| \leq \gamma_\sigma^*} p(0) \exp \left(s(0)\delta - \frac{\nu}{2} \delta^2 \right) \exp \left(-\frac{\delta^2}{2\sigma^2} \right) d\delta \\ &\quad \text{(by log-strong concavity)} \\ &\geq -p(0) \sqrt{\frac{2}{\pi}} \int_{-\infty}^\infty \exp \left(s(0)\delta - \left(\frac{\nu}{2} + \frac{1}{2\sigma^2} \right) \delta^2 \right) d\delta \\ &= -p(0) 2\sqrt{2} \exp \left(\frac{s(0)^2}{\nu + \sigma^{-2}} \right) \frac{1}{\sqrt{\nu + \sigma^{-2}}} \end{aligned} \quad (\text{B.7})$$

We obtained the last equation via Claim D.2. Now we lower bound the second integral in (B.6). By Lemma A.4, $\frac{\partial}{\partial \sigma} g_\sigma(\mu) \geq \frac{\sigma}{4} \ell_{\exp}(\mu) > 0$ for $|\mu| > \gamma_\sigma^*$. Assume without loss of generality that $s(0) > 0$ (so we restrict our attention to $\mu > \gamma_\sigma^* > 0$). By symmetry, our arguments still hold if $s(0) < 0$. Now we have

$$\begin{aligned} \int_{|\mu| > \gamma_\sigma^*} p(\mu) \frac{\partial}{\partial \sigma} g_\sigma(\mu) &> \int_{\delta > 0} p(\gamma_\sigma^* + \delta) \frac{\partial}{\partial \sigma} g_\sigma(\gamma_\sigma^* + \delta) \\ &\geq \frac{\sigma}{4} \int_{\delta > 0} p(\gamma_\sigma^* + \delta) \ell_{\exp}(\gamma_\sigma^* + \delta) \\ &\geq \frac{p(0)\sigma}{4} \int_{\delta > 0} \exp \left(s(0)(\gamma_\sigma^* + \delta) - \frac{\rho}{2} (\delta + \gamma_\sigma^*)^2 \right) \exp(-\delta - \gamma_\sigma^*) d\delta \end{aligned} \quad (\text{B.8})$$

Now we note that for γ_σ^* satisfying $\sqrt{\rho}\gamma_\sigma^* \leq \frac{1}{4}$ and $\delta > 0$, we have

$$s(0)(\gamma_\sigma^* + \delta) - \frac{\rho}{2}(\delta + \gamma_\sigma^*)^2 > s(0)\delta - \frac{\rho}{2}\delta^2 - \rho\delta\gamma_\sigma^* - \frac{\rho}{2}\gamma_\sigma^{*2} \geq \left(s(0) - \frac{\sqrt{\rho}}{4}\right)\delta - \frac{\rho}{2}\delta^2 - \frac{1}{32}$$

As a result, plugging this back into (B.8) gives

$$\int_{|\mu| > \gamma_\sigma^*} p(\mu) \frac{\partial}{\partial \sigma} g_\sigma(\mu) > \exp\left(-\gamma_\sigma^* - \frac{1}{32}\right) \frac{p(0)\sigma}{4} \int_0^\infty \exp\left(\left(s(0) - \frac{\sqrt{\rho}}{4} - 1\right)\delta - \frac{\rho}{2}\delta^2\right) d\delta \quad (\text{B.9})$$

Now we use the fact that $\gamma_\sigma^* \leq 1$ to lower bound $\exp(-\gamma_\sigma^*)$. Finally, we obtain (B.5) by combining (B.7) and (B.9). \square

Now we complete the proof of Lemma B.2.

Proof of Lemma B.2. Now we proceed to lower bound $\frac{\partial}{\partial \sigma} L(w)$. Our starting point is Lemma B.4. We will lower bound the first integral:

$$\begin{aligned} \int_0^\infty \exp\left(\left(|s(0)| - \frac{\sqrt{\rho}}{4} - 1\right)\delta - \frac{\rho}{2}\delta^2\right) d\delta &\geq \int_0^\infty \exp\left(\left(\frac{|s(0)|}{2}\right)\delta - \frac{\rho}{2}\delta^2\right) d\delta \\ &\quad \text{(using } |s(0)| \geq \frac{\sqrt{\rho}}{2} + 2\text{)} \\ &\geq \frac{\sqrt{\pi}}{\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{4\rho}} \quad \text{(from Claim B.1)} \end{aligned}$$

Applying this with equation (B.5) in Lemma B.4, we obtain

$$\frac{\partial}{\partial \sigma} L(w) \geq p(0) \left(\frac{\sigma\sqrt{\pi}}{11\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{4\rho}} - 2\sqrt{2} \exp\left(\frac{s(0)^2}{\nu + \sigma^{-2}}\right) \frac{1}{\sqrt{\nu + \sigma^{-2}}} \right) \quad (\text{B.10})$$

Now we lower bound the second term in (B.10). By applying the upper bound on $s(0)$ in Lemma D.3, we obtain

$$\begin{aligned} \exp\left(\frac{s(0)^2}{\nu + \sigma^{-2}}\right) &\leq \exp\left(\frac{\rho}{\nu + \sigma^{-2}} \log\left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)\right) \\ &\leq \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)^{\frac{\rho}{\nu + \sigma^{-2}}} \end{aligned}$$

Plugging this back into (B.10) and observing that $\frac{1}{\sqrt{\nu + \sigma^{-2}}} \leq \sigma$, we obtain

$$\frac{\partial}{\partial \sigma} L(w) \geq p(0)\sigma \left(\frac{\sqrt{\pi}}{11\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{4\rho}} - 2\sqrt{2} \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)^{\frac{\rho}{\nu + \sigma^{-2}}} \right)$$

Now suppose that the condition $\sigma^2 \rho^2 / \nu \leq 1/8$ holds. Then $\frac{\rho}{\nu + \sigma^{-2}} < \frac{\nu}{8\rho}$, so $\left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)^{\frac{\rho}{\nu + \sigma^{-2}}} \leq \max\left\{1, \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{8\rho}}\right\}$. It follows that

$$\frac{\partial}{\partial \sigma} L(w) \geq p(0)\sigma \left(\frac{\sqrt{\pi}}{11\sqrt{\rho}} \left(\frac{\sqrt{\nu}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{4\rho}} - 2\sqrt{2} \max\left\{1, \left(\frac{\sqrt{\rho}}{2p(0)\sqrt{\pi}}\right)^{\frac{\nu}{8\rho}}\right\} \right)$$

\square

B.4 Proof of Theorem 3.1

We first show that the loss must be lower-bounded by some constant depending only on the distribution over x_1 if w_1 is bounded.

Lemma B.5. For any w with $\|w_1\|_2 \leq R$, as long as $\sigma \leq 1$ the following holds:

$$L(w) \geq 0.25 \exp(-R \mathbb{E}_{x_1} [\|x_1\|_2])$$

Proof. We have

$$\begin{aligned} L(w) &= \mathbb{E}_{x_1} [g_\sigma(w_1^\top x_1)] \\ &\geq \mathbb{E}_{x_1} [g_\sigma(\|w_1\|_2 \|x_1\|_2)] \\ &\geq 0.25 \mathbb{E}_{x_1} [\ell_{\exp}(\|w_1\|_2 \|x_1\|_2)] \quad (\text{by Lemma D.1}) \\ &\geq 0.25 \exp(-R \mathbb{E}_{x_1} [\|x_1\|_2]) \end{aligned} \tag{B.11}$$

The last line followed because $\ell_{\exp}(\mu) = \exp(-\mu)$ for positive μ . Since $\exp(-\mu)$ is convex, we applied Jensen's inequality. \square

Next, we argue that $p(0)$ must be lower-bounded by some constant depending only on the distribution over x_1 if w_1 is bounded.

Lemma B.6. There exists some constant c_1 which only depends on $\mathbb{E}_{x_1} [\|x_1\|_2]$, α , β such that for all w satisfying $\sigma \leq 1/2$ and $1/2 \leq \|w_1\|_2 \leq 1$, we have

$$p(0) \geq c_1$$

Proof. Fix $\bar{\mu} = \log(4/L(w))$. Then note that we must have

$$\int_{-\bar{\mu}}^{\bar{\mu}} p(\mu) d\mu + \int_{|\mu| > \bar{\mu}} p(\mu) \max_{|\mu| > \bar{\mu}} g_\sigma(\mu) \geq L(w).$$

Now note that since $\sigma \leq 1/2$, Lemma B.6 tells us that $g_\sigma(\mu) \leq 2 \exp(-\mu)$. Thus, $\max_{|\mu| > \bar{\mu}} g_\sigma(\mu) \leq L(w)/2$. Thus, we obtain

$$\int_{-\bar{\mu}}^{\bar{\mu}} p(\mu) d\mu + \left(1 - \int_{-\bar{\mu}}^{\bar{\mu}} p(\mu) d\mu\right) \frac{L(w)}{2} \geq L(w)$$

This gives

$$\int_{-\bar{\mu}}^{\bar{\mu}} p(\mu) d\mu \geq \frac{L(w)/2}{1 - L(w)/2}$$

Thus, we can conclude that there exists $\mu' \in [-\bar{\mu}, \bar{\mu}]$ such that $p(\mu') \geq \frac{L(w)}{2\bar{\mu}(2 - L(w))}$.

Now we apply Lemma D.3 to obtain

$$|s(\mu')| \leq \sqrt{\rho \log \left(\frac{\sqrt{\rho} \bar{\mu} (2 - L(w))}{L(w) \sqrt{\pi}} \right)}$$

Now we apply Claim D.1 to conclude that

$$p(0) \geq p(\mu') \exp \left(-|s(\mu')| \bar{\mu} - \frac{\rho}{2} \bar{\mu}^2 \right)$$

Now note that $s(\mu')$, $p(\mu')$, $\bar{\mu}$ depend only on $L(w)$ which is upper bounded by 1 and lower bounded by some function of $\mathbb{E}_{x_1} [\|x_1\|_2]$ by Lemma B.5. Furthermore, $\rho \in [\beta, 2\beta]$. Thus, $s(\mu')$, $p(\mu')$, $\bar{\mu}$ are all upper and lower bounded by some function of $\mathbb{E}_{x_1} [\|x_1\|_2]$. As a result, the same applies to $p(0)$, giving us the desired statement. \square

Proof of Theorem 3.1. We start with proving the case when $K = 1$. First, we note that $w_1^\top \nabla_{w_1} L(w) < 0$ by using the same argument as Lemma A.2. Furthermore, for all $\|w_1\|_2 \in [1/2, 1]$, the upper bound on σ required for Lemmas B.1, B.2, and B.3 are all satisfied by Assumption 3.2. Thus, if $L(w^s) \leq \kappa(\beta, \alpha) = \min_{a \in [1, 4]} \tilde{\kappa}(a\beta, a\alpha)$, then initially the loss upper bound for Lemmas B.1, B.2, and B.3 are satisfied. Combining this with Lemma B.6, we get that $\frac{\partial}{\partial \sigma} L(w) \geq c_1 \sigma$ for the constant c_1 defined in Lemma B.6. Furthermore, the loss $L(w)$ is also always decreasing

for sufficiently small choice of step size. As a result, the following invariants hold throughout the optimization algorithm: $\|w_1\|_2$ is non-decreasing, σ is non-increasing, and $L(w) \leq \kappa(\beta, \alpha)$. Thus, the initial conditions Lemmas B.1, B.2, and B.3 will always hold, so we can conclude using the same argument as in Lemma A.3 that $\|w_2\|_2$ is always decreasing with rate $c_2\|w_2\|_2$, where c_2 is some value depending on α, β , and the data distribution. This implies that w_2 converges to 0, providing the first statement in Theorem 3.1.

Finally, in the case that $K > 1$, we observe that when Assumption 3.2 is satisfied, we must have $L_i(w) \leq \kappa(\beta, \alpha)$ for all i , where $L_i(w)$ is the expectation of the loss conditioned on the i -th mixture component. Thus, this immediately reduces to the $K = 1$ case. \square

To prove convergence of noisy gradient descent to an approximate local minimum of the objective

$$\min_{\|w_1\|_2 \leq 1} L((w_1, 0)) \quad (\text{B.12})$$

we also assume that $L((w_1, w_2))$ is twice-differentiable, and furthermore there exists C such that $\nabla_{w_1} L((w_1, w_2)), \nabla_{w_1}^2 L((w_1, w_2))$ are Lipschitz in w_2 for $\|w_2\|_2 \leq C, \|w_1\|_2 \leq 1$.

We will first formally define an (ϵ, γ) -approximate local minimum of (B.12). Define $P_{w_1^\perp} \triangleq I - \frac{w_1 w_1^\top}{\|w_1\|_2^2}$ to be the projection onto the space orthogonal to w_1 . Then an (ϵ, γ) -approximate local minimum of (B.12) is a point $w_1 : \|w_1\|_2 \leq 1$ satisfying:

1. $\|w_1\|_2 \geq 1 - \epsilon$.
2. $\|P_{w_1^\perp} \nabla_{w_1} L((w_1, 0))\|_2 \leq \epsilon$.
3. $P_{w_1^\perp} \nabla_{w_1}^2 L((w_1, 0)) P_{w_1^\perp} - (w_1^\top \nabla_{w_1} L((w_1, 0))) P_{w_1^\perp} \succeq -\gamma I$.

Note that the first condition simply reflects the fact that all true local minima of (B.12) will satisfy $\|w_1\|_2 = 1$ and therefore lie on the unit sphere \mathbb{S}^{d_1-1} , as scaling up the weights only decreases the objective. The second two conditions essentially adapt the classical conditions for approximate local minima (see [25, 1]) to the setting where the domain is a Riemannian manifold (in our case, the unit sphere \mathbb{S}^{d_1-1}). In particular, they replace the standard gradient and Hessian with the gradient and Hessian on a Riemannian manifold, in the special case when the manifold is the unit sphere. In other words, they capture the intuition in order for w_1 to be a local minimizer of the constrained objective, the only local change to w_1 which decreases the loss should be increasing its norm.

To conclude convergence to an approximate local minimum, we note that by the argument of [11], there are sufficiently small step size and additive noise such that for any choice of ϵ , the algorithm converges to an (ϵ, γ) -approximate local minimizer of the objective $\min_{\|w\|_2 \leq 1} L(w_1, w_2)$ (defined in the same manner) satisfying $\|w_2\|_2 \leq \epsilon$. By the regularity conditions on L , this is also a $(C'\epsilon, C'\gamma)$ -approximate local minimizer of the purified objective for some C' depending on the regularity conditions.

C Proofs for finite sample setting

Given n samples X_1, \dots, X_n define the empirical loss \hat{L} on *unlabeled* data as:

$$\hat{L}(w) = \frac{1}{n} \sum_{i=1}^n \ell_{exp}(w^T X_i)$$

We analyze self-training on the empirical loss, which begins with a classifier w^S , and does projected gradient descent on \hat{L} with learning rate η .

$$w^0 = w^S$$

$$w^{t+1} = R \frac{w^t - \eta \nabla \hat{L}(w^t)}{\|w^t - \eta \nabla \hat{L}(w^t)\|_2}$$

Recap of analysis in infinite setting: In the infinite sample case, gradient descent moves in direction $-\nabla L(w)$, and we show that $\|w_2\|_2 \rightarrow 0$ as we self-train. If the loss were convex, we could just analyze the minima and show it had the desired property that $\|w\|_2 = 0$. Standard results in convex analysis would then show convergence. Since the loss is non-convex, the core of the proof is to bound certain directional gradients. In particular, we showed that $\langle \nabla_{w_1} L(w), w_1 \rangle < 0$, $\langle \nabla_{w_2} L(w), w_2 \rangle \geq c_1^2 \|w_2\|_2^2$, and $\|\nabla L(w)\|_2^2 \leq c_2^2 \|w\|_2^2$. Using this, we analyzed the gradient descent iterates and showed that $\|w_2\|_2$ decreased by a multiplicative factor at each step of self-training.

Finite sample proof overview: With finite samples, gradient descent instead moves in direction $-\nabla \hat{L}(w)$ where $\hat{L}(w)$ is the empirical loss on n samples. Here w_2 won't go to exactly 0, but to a very small value: we will show $\|w_2\|_2 \rightarrow \tau$ with high probability if we use $\tilde{O}(1/\tau^4)$ samples. At a high level, we will show a uniform concentration bound which shows that the empirical gradient $\nabla \hat{L}(w)$ and population gradient $\nabla L(w)$ are close for all w (Lemma C.1). In Theorem C.1, this lets us show that $\langle \nabla_{w_1} \hat{L}(w), w_1 \rangle$, $\langle \nabla_{w_2} \hat{L}(w), w_2 \rangle$, and $\|\nabla \hat{L}(w)\|_2^2$ are similar to the population versions above with $L(w)$ instead of $\hat{L}(w)$. We use this to show that $\|w\|_2$ will keep decreasing until $\|w\|_2 \leq \tau$, and will then stabilize and stay below $\|\tau\|$ forever.

Notation: To avoid defining too many constants, we use big-O notation in the following sense that is different from the standard computer science usage but common in learning theory proofs. When we use $O(e)$ in an expression, we mean that expression e can be replaced by $c_e e$ for some universal constant c_e that does not depend on *any* problem parameters (like $\delta, \sigma_{\min}, \gamma$, etc)—it is literally just some number like say $67/32$, but explicitly putting the numbers everywhere makes expressions messy. For sub-Gaussian, sub-exponential, we will use notation and standard results from [40], in particular the *norm* of a sub-Gaussian random variable (Definition 2.5.6), equivalent properties of sub-Gaussian random variables (Proposition 2.5.2), the relationship between sub-Gaussian and sub-exponential random variables (Lemma 2.7.6), and Bernstein's inequality for sub-exponential random variables (Theorem 2.8.1).

We define the empirical expectation \hat{E} for any function f over the n samples X_1, \dots, X_n :

$$\hat{E}[f(X)] = \frac{1}{n} \sum_{i=1}^n f(X_i)$$

C.1 Results

Our first Lemma shows that the empirical gradients $\nabla \hat{L}(w)$ and population gradient $\nabla L(w)$ are close for all w , if the distribution is sub-Gaussian. We will show later that Gaussian distributions and mixtures of K log-concave distributions are indeed sub-Gaussian. Data that is normalized will also satisfy the sub-Gaussian assumption.

Lemma C.1. *Let $\pi : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ be a projection operator with $d' \leq d$, that is π is a d' -by- d matrix where each row of π is orthonormal, and suppose the distribution $X \sim p(x)$ satisfies that $\|X\|_2$ is sub-Gaussian with norm B (equivalently, variance parameter B^2). Suppose we choose:*

$$n = \tilde{O}\left(\frac{d}{\epsilon^2} R^2 B^2 \log 1/\delta\right)$$

Where we hide terms that are logarithmic in $\frac{1}{\delta}$, B , R , and d in the big- O here to highlight the prominent terms, but give the full version in the proof. Then, with probability $\geq 1 - \delta$, for all w with $\|w\|_2 \leq R$, we have:

$$|\widehat{E}[l'(w^\top X)\pi(w)^\top \pi(X)] - E[l'(w^\top X)\pi(w)^\top \pi(X)]| \leq \epsilon \quad (\text{C.1})$$

Proof. We will use a discretization argument. We first show the concentration in Equation C.1 for fixed w using the fact that the distribution is sub-Gaussian and then applying Hoeffding's inequality. We will then construct an ϵ -cover of the R -ball (in ℓ_2 norm) and use union bound so that the concentration holds for each member of the ϵ -cover. Finally, we will show that the concentration holds for all w with $\|w\|_2 \leq R$. Let $h(w, X) = l'(w^\top X)\pi(w)^\top \pi(X)$.

Step 1: Concentration for single w : First, we have:

$$\begin{aligned} |h(w, X)| &= |l'(w^\top X)\pi(w)^\top \pi(X)| \\ &\leq |\pi(w)^\top \pi(X)| \\ &\leq \|\pi(w)\|_2 \|\pi(X)\|_2 \\ &\leq \|w\|_2 \|X\|_2 \\ &\leq R \|X\|_2 \end{aligned}$$

We want to bound $|\widehat{E}[h(w, X)] - E[h(w, X)]|$. Since $\|X\|_2$ is sub-Gaussian with norm B , $h(w, X)$ is sub-Gaussian with norm RB , it then follows that $h(w, X) - E[h(w, X)]$ is a mean 0 sub-Gaussian random variable with norm $2RB$. So the average, $\widehat{E}[h(w, X)] - E[h(w, X)]$ is mean 0 and sub-Gaussian with norm $2RB/\sqrt{n}$. By the sub-Gaussian tail bound, we then get that with probability at least $1 - \delta/5$:

$$|\widehat{E}[h(w, X)] - E[h(w, X)]| \leq O\left(\frac{1}{\sqrt{n}} RB \sqrt{\log(1/\delta)}\right)$$

To control the RHS to be less than $\epsilon/3$ with probability at least $1 - \delta/5$, it suffices to choose:

$$n = O\left(\frac{1}{\epsilon^2} R^2 B^2 \log(1/\delta)\right)$$

Note that this is only for a single w .

Step 2: κ -covering: We will now construct a κ -covering consisting of M vectors w . We will want the above inequality to hold for all M vectors—to do this we will apply union bound. More precisely, a standard covering argument tells us that we can choose M with $\log M \leq d \log(1 + (2R)/\kappa)$ and M vectors w_1, \dots, w_M s.t. for any w with $\|w\|_2 \leq R$, there exists w_i with $\|w_i\|_2 \leq R$ and $\|w_i - w\|_2 \leq \kappa$. By union bound, we have that if we choose:

$$n = O\left(\frac{1}{\epsilon^2} R^2 B^2 (\log M + \log(1/\delta))\right)$$

Then for all w_i the empirical concentration holds, that is with probability at least $1 - \delta/5$, for all $w = w_i$:

$$|\widehat{E}[h(w, X)] - E[h(w, X)]| \leq \epsilon/3$$

It now remains to choose κ so that we can show this for all w (not just $w = w_i$).

Step 3: Handling all w by choosing κ small: To extend the result to all w (not just $w = w_i$), we consider arbitrary w, w' with $\|w\|_2, \|w'\|_2 \leq R$ and $\|w - w'\|_2 \leq \kappa$. We want to show that the difference in their directional derivatives is not too large. More precisely, we would like to show that with probability $\geq 1 - \delta/5$, for all such w, w' :

$$|\widehat{E}[h(w', X)] - \widehat{E}[h(w, X)]| \leq \epsilon/3 \quad (\text{C.2})$$

And similarly for its empirical counterpart, \widehat{E} . This then proves the main claim, because for any w with $\|w\|_2 \leq R$, we can choose some w_i in the κ -cover above. We then have:

$$\begin{aligned} |\widehat{E}[h(w, X)] - E[h(w, X)]| &\leq |\widehat{E}[h(w, X)] - \widehat{E}[h(w_i, X)]| \\ &\quad + |\widehat{E}[h(w_i, X)] - E[h(w_i, X)]| \\ &\quad + |E[h(w_i, X)] - E[h(w, X)]| \end{aligned}$$

And each of the terms in the RHS will be bounded above by $\epsilon/3$, so the LHS will be bounded above by ϵ . To show Equation C.2, we first write:

$$\begin{aligned} h(w', X) - h(w, X) &= (l'(w'^\top X)\pi(w')^\top \pi(X) - l'(w^\top X)\pi(w')^\top \pi(X)) \\ &\quad + (l'(w^\top X)\pi(w')^\top \pi(X) - l'(w^\top X)\pi(w)^\top \pi(X)) \end{aligned}$$

Using Cauchy-Schwarz, and using the fact that $l'(r) \leq 1$ for all r and that l' is 1-Lipschitz, we can show for the first term in the RHS:

$$|l'(w'^\top X)\pi(w')^\top \pi(X) - l'(w^\top X)\pi(w')^\top \pi(X)| \leq \kappa R \|X\|_2^2$$

And for the second term in the RHS:

$$|l'(w^\top X)\pi(w')^\top \pi(X) - l'(w^\top X)\pi(w)^\top \pi(X)| \leq \kappa \|X\|_2$$

Combining the above 2 inequalities:

$$|h(w', X) - h(w, X)| \leq \kappa R \|X\|_2^2 + \kappa \|X\|_2$$

We will show below that $\mathbb{E}[\|X\|_2] = O(B)$, $\widehat{\mathbb{E}}[\|X\|_2] = O(B)$, $\mathbb{E}[\|X\|_2^2] = O(B^2)$, $\widehat{\mathbb{E}}[\|X\|_2^2] = O(B^2)$ (for the empirical expectations, this will hold with probability at least $1 - \delta/5$). Assuming this for now, this gives us that it suffices to choose κ such that:

$$\frac{1}{\kappa} \geq \frac{1}{\epsilon} [RB^2 + B]$$

In which case, Equation C.2 and its empirical counterpart hold. In total, this means we require n to be:

$$n = O\left(\frac{1}{\epsilon^2} R^2 B^2 \left(d \log \left[1 + \frac{R^2 B^2 + RB}{\epsilon}\right] + \log(1/\delta)\right)\right)$$

Or omitting log terms except in $1/\delta$ (we keep $1/\delta$ to make the dependence on the probability explicit):

$$n = \tilde{O}\left(\frac{d}{\epsilon^2} R^2 B^2 \log(1/\delta)\right)$$

Bounding the norm and norm-squared: Finally, we bound the expectations of the norm and norm-squared of X , which we used above. By taking integrals, since X is sub-Gaussian with norm B , we can show that:

$$\mathbb{E}[\|X\|_2] \leq O(B)$$

$$\mathbb{E}[\|X\|_2^2] \leq O(B^2)$$

Next, we will like to bound the empirical means of these quantities. Since $\|X\|_2$ is sub-Gaussian with norm B . This means that $\|X\|_2 - \mathbb{E}[\|X\|_2]$ is mean 0 and sub-Gaussian with norm $2B$. So for the average of n iid samples, we have that with probability $\geq 1 - \delta$:

$$\widehat{\mathbb{E}}[\|X\|_2] \leq \mathbb{E}[\|X\|_2] + O\left(\frac{B}{\sqrt{n}} \sqrt{\log \frac{1}{\delta}}\right)$$

As long as we choose $n \geq O(\log(1/\delta))$, we have with probability at least $1 - \delta/5$:

$$\widehat{\mathbb{E}}[\|X\|_2] \leq O(B)$$

Squares of sub-Gaussian random variables are sub-exponential, so $\|X\|_2^2$ is sub-exponential with sub-exponential norm $O(B^2)$. Then, $\|X\|_2^2 - \mathbb{E}[\|X\|_2^2]$ is sub-exponential with sub-exponential norm $O(B^2)$. Then by Bernstein's inequality for sub-exponentials, as long as we choose $n \geq O(\log(1/\delta))$, we have with probability at least $1 - \delta/5$:

$$\widehat{\mathbb{E}}[\|X\|_2^2] \leq \mathbb{E}[\|X\|_2^2] + O(B^2) \leq O(B^2)$$

□

The main theorem of this section shows that for a sub-Gaussian distribution, if we have bounds on $\langle \nabla_{w_1} L(w), w_1 \rangle$, $\langle \nabla_{w_2} L(w), w_2 \rangle$, and $\|\nabla L(w)\|_2^2$ for the *population*, but do entropy minimization on the *empirical samples*, we will still converge with $\|w_2\| \leq \tau$. We will later instantiate these bounds for the Gaussian setting and the more general log-concave setting.

Theorem C.1. Suppose that for all w , $\langle \nabla_{w_1} L(w), w_1 \rangle < 0$, $\langle \nabla_{w_2} L(w), w_2 \rangle \geq c_1 \|w_2\|_2^2$, and $\|\nabla L(w)\|_2^2 \leq c_2^2 \|w\|_2^2$, for some $c_1, c_2 > 0$ where c_1, c_2 are not a function of w . Let $\tau < 0.5$ be the desired norm for the spurious feature w_2 , that is, we want $\|w_2\|_2 \leq \tau$ after running self-training. Let $\epsilon = O(c_1 \tau^2)$, and choose $n = \tilde{O}\left(\frac{1}{\epsilon^2} R^2 B^2 \log(1/\delta)\right)$ according to the Lemma C.1 such that with probability $\geq 1 - \delta$, for all w with $\|w\|_2 \leq \max(R, 1)$ the empirical gradients along both w_1 and w_2 are near the true gradient:

$$|\hat{E}[l'(w^\top X)w_1^\top x_1] - E[l'(w^\top X)w_1^\top x_1]| \leq \epsilon$$

$$|\hat{E}[l'(w^\top X)w_2^\top x_2] - E[l'(w^\top X)w_2^\top x_2]| \leq \epsilon$$

Then if initially $\|w_2^0\|_2 \leq 0.5$, self-training with step size $\eta = O(\frac{c_1}{c_2^2})$, will converge to $\|w_2\|_2 \leq \tau$. Specifically, if at step t , $\|w_2^t\|_2 \geq \tau/2$, then the norm of w_2 shrinks by a multiplicative factor and rapidly reduces to less than $\tau/2$:

$$\|w_2^{t+1}\|_2^2 < \left(1 - O\left(\frac{c_1}{c_2^2}\right)^2\right) \|w_2^t\|_2^2$$

Furthermore, once this has happened, the norm stabilizes: if $\|w_2^t\|_2 < \tau/2$, then $\|w_2^{t+1}\|_2 \leq \tau$.

Proof. We note that $\langle \nabla_{w_2} L(w), w_2 \rangle = \mathbb{E}[l'(w^\top X)w_2^\top x_2]$, $\langle \nabla_{w_2} \hat{L}(w), w_2 \rangle = \hat{E}[l'(w^\top X)w_2^\top x_2]$, and similarly for w_1 . So we have for all w with $\|w\|_2 \leq \max(R, 1)$:

$$|\langle \nabla_{w_1} L(w), w_1 \rangle - \langle \nabla_{w_1} \hat{L}(w), w_1 \rangle| \leq \epsilon \quad (\text{C.3})$$

$$|\langle \nabla_{w_2} L(w), w_2 \rangle - \langle \nabla_{w_2} \hat{L}(w), w_2 \rangle| \leq \epsilon \quad (\text{C.4})$$

Step 1: Bounding empirical gradients: The main optimization analysis requires us to bound 3 quantities: $\langle \nabla_{w_2} \hat{L}(w), w_2 \rangle$, $\langle \nabla_{w_1} \hat{L}(w), w_1 \rangle$, and $\|\nabla_{w_2} \hat{L}(w)\|_2^2$, which we first do.

Equation C.4 gives us a bound on the empirical gradient along w_2 :

$$\langle \nabla_{w_2} \hat{L}(w), w_2 \rangle \geq c_1 \|w_2\|_2^2 - \epsilon$$

Equation C.3 this gives us a bound on the empirical gradient along w_1 :

$$\langle \nabla_{w_1} \hat{L}(w), w_1 \rangle < \epsilon$$

Finally, we bound $\|\nabla_{w_2} \hat{L}(w)\|_2^2$:

$$\begin{aligned} \|\nabla_{w_2} \hat{L}(w)\|_2 &= \left(\max_{\|v\|_2 \leq 1} \langle v, \nabla_{w_2} \hat{L}(w) \rangle \right)^2 \\ &\leq \left(\max_{\|v\|_2 \leq 1} \langle v, \nabla_{w_2} L(w) \rangle + \epsilon \right)^2 \\ &= (\|\nabla_{w_2} L(w)\|_2 + \epsilon)^2 \\ &= 2\|\nabla_{w_2} L(w)\|_2^2 + 2\epsilon^2 \\ &\leq 2c_2^2 \|w_2\|_2^2 + 2\epsilon^2 \end{aligned}$$

Where in the first line we used the variational form of the 2-norm, second line we used Equation C.4, in the third line we used the variational form of the 2-norm again, fourth line we used the fact that $(a + b)^2 \leq 2a^2 + 2b^2$, and in the fifth line we used the bound on $\|\nabla_{w_2} L(w)\|_2^2$ in the assumption of the theorem.

Step 2: Show w_2 decreases and stabilizes: Our updates involve taking a gradient descent step, and then projecting back to the sphere, $\|w\|_2 \leq R$. Define $\tilde{w}^{t+1} = w^t - \eta \nabla L(w^t)$ to be the iterate before projecting. Then, we have:

$$\begin{aligned} \|\tilde{w}_2^{t+1}\|_2^2 &= \|w_2^t - \eta \nabla_{w_2} \hat{L}(w)\|_{w=w^t}^2 \\ &= \|w_2^t\|_2^2 + \eta^2 \|\nabla_{w_2} \hat{L}(w)\|_{w=w^t}^2 \\ &\quad - 2\eta \langle \nabla_{w_2} \hat{L}(w)\|_{w=w^t}, w_2^t \rangle \\ &\leq (1 + 2\eta^2 c_2^2 - 2\eta c_1) \|w_2^t\|_2^2 + (2\eta^2 \epsilon^2 + 2\eta \epsilon) \end{aligned}$$

We choose η as:

$$\eta = \frac{c_1}{c_2^2}$$

Which gives us:

$$\|\tilde{w}_2^{t+1}\|_2^2 \leq \left(1 - \frac{1}{2} \frac{c_1^2}{c_2^2}\right) \|w_2^t\|_2^2 + (2\eta^2 \epsilon^2 + 2\eta\epsilon)$$

Since the norm is always non-negative, we note that $c_1^2/c_2^2 \leq 2$. To control the error terms, we choose ϵ as:

$$\epsilon = \frac{1}{48} c_1 \tau^2$$

Then, we get,

$$2\eta^2 \epsilon^2 + 2\eta\epsilon \leq \frac{1}{4} \frac{c_1^2}{c_2^2} (\tau/2)^2$$

In other words, if $\|w_2^t\|_2 \geq \tau/2$, then the norm decreases:

$$\|\tilde{w}_2^{t+1}\|_2^2 \leq \left(1 - \frac{1}{4} \frac{c_1^2}{c_2^2}\right) \|w_2^t\|_2^2$$

And if $\|w_2^t\|_2 < \tau/2$, then the norm stabilizes:

$$\|\tilde{w}_2^{t+1}\|_2^2 \leq \|w_2^t\|_2^2 + \frac{1}{2} (\tau/2)^2 \leq \frac{3}{2} (\tau/2)^2$$

Step 3: Show w_1 does not decrease much: We have shown that \tilde{w}_2^{t+1} is smaller than w_2^t . Next, we need to deal with the renormalization step to show that w_2^{t+1} is also smaller than w_2^t . We will show that \tilde{w}_1 cannot decrease by too much, so that after renormalization, the norm of w_2 is still decreasing sufficiently. We have:

$$\begin{aligned} \|\tilde{w}_1^{t+1}\|_2^2 &= \|w_1^t - \eta \nabla_{w_1} \hat{L}(w)|_{w=w^t}\|_2^2 \\ &\geq \|w_1^t\|_2^2 - 2\eta \langle \nabla_{w_1} \hat{L}(w)|_{w=w^t}, w_1^t \rangle \\ &\geq \|w_1^t\|_2^2 - 2\eta\epsilon \end{aligned}$$

From our choice of η and ϵ , we can show that $\eta\epsilon$ is actually very small. In particular, with some algebra, we can show that

$$2\eta\epsilon < \frac{1}{24} \frac{c_1^2}{c_2^2} \leq \frac{1}{12}$$

In effect, the decrease in the norm of w_1 is at least 10 times smaller than the decrease in the norm of w_2 . Now we note that since at all times t , $\|w_2^t\|_2 \leq 0.5$, we have $\|w_1^t\|_2 \geq \|w_2^t\|_2$. So w_1 is larger and decreases by a much smaller amount, which means that after renormalizing w_2 still decreases by around the same amount. Formally, with a bit of algebra, we get that if $\|w_2^t\|_2 \geq \tau/2$, then *after renormalizing*,

$$\|w_2^{t+1}\|_2^2 \leq \left(1 - \frac{1}{10} \frac{c_1^2}{c_2^2}\right) \|w_2^t\|_2^2$$

And on the other hand, if $\|w_2^t\|_2 < \tau/2$ then *after renormalizing*, $\|w_2^{t+1}\|_2^2 < \tau^2$. This completes the proof. □

C.2 Applying the finite sample results

We now instantiate the above Theorem C.1 for the Gaussian setting.

Proof of Theorem 3.2 finite sample guarantee. By Lemma A.3, $\langle \nabla_{w_2} L(w), w_2 \rangle \geq O(\tilde{\sigma}_{\min} \|w_2\|_2^2)$. By Lemma A.2, $\langle \nabla_{w_1} L(w), w_1 \rangle < 0$. In the proof of Theorem 3.2 we showed $\|\nabla L(w)\|_2^2 \leq O(\tilde{\sigma}_{\max}^2 \|w\|_2^2)$. Furthermore, $\|X\|_2$ is sub-Gaussian with norm $O(\gamma + \tilde{\sigma}_{\min} \sqrt{d})$. Applying Theorem C.1, we get that if we choose

$$n = \tilde{O}\left(\frac{1}{\tau^4} \frac{R^2(\gamma + \tilde{\sigma}_{\max}^2 d)}{\tilde{\sigma}_{\min}^2} \log(1/\delta)\right)$$

then after t iterations, if for all $t \geq T$, where

$$T = O\left(\frac{\tilde{\sigma}_{\max}^2}{\tilde{\sigma}_{\min}^2} \log \frac{R}{\tau}\right)$$

We will have $\|w_2^t\| \leq \tau$. Furthermore, since w_1 is 1-dimensional in this case and is non-negative initially, from the proof of Theorem C.1, w_1 after re-normalizing always increases and stays non-negative. As such, $w_1 \geq \sqrt{R^2 - \tau^2}$. \square

Next, we instantiate the theorem for the mixture of K sliced log-concave setting:

Proof of Theorem 3.1 finite sample guarantee. In the population case proof of Theorem 3.1, we showed that $\langle \nabla_{w_2} L(w), w_2 \rangle \geq O(c_1 \|w_2\|_2^2)$, $\langle \nabla_{w_1} L(w), w_1 \rangle < 0$, and $\|\nabla L(w)\|_2^2 \leq O(c_2^2 \|w\|_2^2)$. Additionally, we note that a mixture of K sliced log-concave distributions is sub-Gaussian. So we get that there exists some c, c' that depends on the distribution, such that if we choose $n \geq c/\tau^4$ samples then after t iterations, for all $T \geq t$, where $T = O(\log(R/\tau))$, we will have $\|w_2^t\| \leq \tau$. \square

D Additional Missing Proofs

D.1 Bounds on ℓ_{exp} and derivatives

Proof of Lemma A.4. We can exactly compute the expression

$$\begin{aligned} q_\sigma(\mu) &= \frac{\partial g_\sigma(\mu)}{\partial \sigma} \\ &= \frac{1}{2} \exp\left(\frac{\sigma^2}{2}\right) \sigma \left[\exp(\mu) \operatorname{erfc}\left(\frac{\sigma}{\sqrt{2}} + \frac{\mu}{\sqrt{2}\sigma}\right) + \exp(-\mu) \operatorname{erfc}\left(\frac{\sigma}{\sqrt{2}} - \frac{\mu}{\sqrt{2}\sigma}\right) \right] \\ &\quad - \sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \end{aligned} \quad (\text{D.1})$$

As $q_\sigma(\mu)$ and ℓ_{exp} are both symmetric around 0, we assume w.l.o.g. that $\mu \geq 0$. We first consider $\sigma \leq \frac{4\sqrt{2}}{\sqrt{\pi}}$. Since $\mu \geq \sigma^2$, $\sigma - \mu/\sigma \leq 0$ so $\operatorname{erfc}((\sigma - \mu/\sigma)/\sqrt{2}) \geq 1$. Thus, we have

$$q_\sigma(\mu) \geq -\sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) + \frac{1}{2} \sigma \exp\left(-\mu + \frac{\sigma^2}{2}\right)$$

Note that for $\mu \geq \gamma_\sigma^*$, we have

$$\frac{1}{4} \sigma \exp\left(-\mu + \frac{\sigma^2}{2}\right) \geq \sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right)$$

in which case we can obtain $q_\sigma(\mu) \geq \frac{1}{4} \sigma \exp\left(-\mu + \frac{\sigma^2}{2}\right)$ by rearranging.

Now we consider $\sigma > \frac{4\sqrt{2}}{\sqrt{\pi}}$. Since $\mu \geq 2\sigma^2$, we have $\operatorname{erfc}((\sigma - \mu/\sigma)/\sqrt{2}) \geq 1$ and $-\sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \geq -\sqrt{\frac{2}{\pi}} \ell_{exp}(\mu)$. Therefore

$$q_\sigma(\mu) \geq -\sqrt{\frac{2}{\pi}} \ell_{exp}(\mu) + \frac{1}{2} \sigma \ell_{exp}(\mu) \geq \frac{1}{4} \sigma \ell_{exp}(\mu)$$

□

Lemma D.1. For all μ , the following holds:

$$q_\sigma(\mu) = \frac{\partial}{\partial \sigma} g_\sigma(\mu) \geq -\sqrt{\frac{2}{\pi}} \exp\left(-\frac{\mu^2}{2\sigma^2}\right) \quad (\text{D.2})$$

Furthermore, for $\sigma \leq 1$, we also have

$$g_\sigma(\mu) \geq 0.25 \ell_{exp}(\mu) \quad (\text{D.3})$$

Proof. To conclude (D.2), we simply use the fact that erfc is always positive, so only the last term in (D.1) can be negative.

To conclude the second statement, assume without loss of generality that $\mu > 0$. We first note that with probability at least 0.68, $1 \geq Z \geq -1$, and additionally, $\ell_{exp}(\mu + \sigma Z) \geq \exp(-\sigma) \ell_{exp}(\mu)$ for $1 \geq Z \geq -1$. When $\sigma \leq 1$, we thus have $g_\sigma \geq 0.25 \ell_{exp}(\mu)$. □

Lemma D.2. For all μ and $\sigma \leq 1/2$, the following holds:

$$g_\sigma(\mu) \leq 2 \ell_{exp}(\mu)$$

Proof. Without loss of generality, assume that $\mu > 0$. We note that we can upper bound $\ell_{exp}(\mu)$ by the loss function $\exp(-\mu)$. It follows that

$$\begin{aligned} g_\sigma(\mu) &\leq \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left(-\mu - \sigma Z - \frac{Z^2}{2}\right) dZ \\ &= \exp(-\mu) \exp(\sigma^2) \sqrt{2} \end{aligned} \quad (\text{D.4})$$

As ℓ_{exp} is symmetric around 0, we could also apply the same argument to $\mu < 0$ using $\exp(\mu)$ as the loss upper bound. This gives the desired result. □

D.2 Log-concave and smooth densities

Claim D.1. Let $p : \mathbb{R} \rightarrow \mathbb{R}$ be any density such that $\log p$ is differentiable, ν -strongly concave, and ρ -smooth. Define $s(\mu) \triangleq \frac{\partial}{\partial \mu} \log p$. Then for all $\mu \in \mathbb{R}$, the following hold:

$$\begin{aligned} p(\mu + \delta) &\geq p(\mu) \exp\left(s(\mu)\delta - \frac{\rho}{2}\delta^2\right) \\ p(\mu + \delta) &\leq p(\mu) \exp\left(s(\mu)\delta - \frac{\nu}{2}\delta^2\right) \end{aligned}$$

Proof. By strong concavity and smoothness, we have

$$\begin{aligned} \log p(\mu + \delta) &\geq \log p(\mu) + \frac{\partial}{\partial \mu} \log p(\mu)\delta - \frac{\rho}{2}\delta^2 \\ \log p(\mu + \delta) &\leq \log p(\mu) + \frac{\partial}{\partial \mu} \log p(\mu)\delta - \frac{\nu}{2}\delta^2 \end{aligned}$$

Exponentiating both sides and using the definition of s gives the desired result. \square

Lemma D.3. In the setting of Claim D.1, we have the following upper and lower bounds for s in terms of p :

$$\begin{aligned} s^2(\mu) &\geq \nu \log\left(\frac{\sqrt{\nu}}{2p(\mu)\sqrt{\pi}}\right) \\ s^2(\mu) &\leq \rho \log\left(\frac{\sqrt{\rho}}{2p(\mu)\sqrt{\pi}}\right) \end{aligned}$$

In other words, by rearranging,

$$\begin{aligned} p(\mu) &\geq \frac{\sqrt{\nu}}{2\sqrt{\pi}} \exp(-s(\mu)^2/\nu) \\ p(\mu) &\leq \frac{\sqrt{\rho}}{2\sqrt{\pi}} \exp(-s(\mu)^2/\rho) \end{aligned}$$

Proof. First, from Claim D.1, we have

$$\begin{aligned} 1 &= \int_{-\infty}^{\infty} p(\mu + \delta) d\delta \leq \int_{-\infty}^{\infty} p(\mu) \exp\left(s(\mu)\delta - \frac{\nu}{2}\delta^2\right) d\delta \\ &= \frac{2\sqrt{\pi}p(\mu) \exp(s(\mu)^2/\nu)}{\sqrt{\nu}} \end{aligned}$$

Solving, we obtain

$$s^2(\mu) \geq \nu \log\left(\frac{\sqrt{\nu}}{2p(\mu)\sqrt{\pi}}\right)$$

Likewise, we can use the same reasoning to obtain

$$s^2(\mu) \leq \rho \log\left(\frac{\sqrt{\rho}}{2p(\mu)\sqrt{\pi}}\right)$$

\square

Claim D.2. For any $a \in \mathbb{R}$, $b > 0$, we have

$$\int_{-\infty}^{\infty} \exp(ax - bx^2) = \frac{\sqrt{2\pi} \exp(a^2/2b)}{\sqrt{b}} \quad (\text{D.5})$$

$$\int_0^{\infty} \exp(ax - bx^2) = \frac{\sqrt{\pi} \exp(a^2/2b) \left(\operatorname{erf}\left(\frac{a}{\sqrt{2b}}\right) + 1\right)}{\sqrt{2b}} \quad (\text{D.6})$$

Furthermore, when $a \geq 0$, we additionally have

$$\int_0^{\infty} \exp(ax - bx^2) \in \left[\frac{\sqrt{\pi} \exp(a^2/2b)}{\sqrt{2b}}, \frac{\sqrt{2\pi} \exp(a^2/2b)}{\sqrt{b}} \right] \quad (\text{D.7})$$

Proof. Equations (D.5) and (D.6) follow from direct computation. Equation (D.7) follows because for $a \geq 0$, $1 \geq \operatorname{erf}(a/\sqrt{2b}) \geq 0$. \square

Claim D.3. For $a < 0$,

$$1 + \operatorname{erf}(a) > \frac{2}{\sqrt{\pi}} \frac{\exp(-a^2)}{-a + \sqrt{a^2 + 2}} \geq \frac{1}{\sqrt{\pi}(\sqrt{2} - 2a)} \exp(-a^2)$$

Proof. We have $1 + \operatorname{erf}(a) = 1 - \operatorname{erf}(-a) = 1 - (1 - \operatorname{erfc}(-a)) = \operatorname{erfc}(-a)$. Now as $-a > 0$, we can apply the lower bound on $\operatorname{erfc}(-a)$ in [41] to obtain the desired result. \square

D.3 Equivalence between pseudo-labeling variant and entropy minimization

Proposition D.1. The pseudo-labeling algorithm above converges to the same solution as the entropy minimization algorithm in (2.3).

Proof of Proposition D.1. We compute

$$\begin{aligned} \nabla_w L_{pseudo}^{t+1}(w)|_{w=w^t} &= \nabla_w \mathbb{E}_{x \sim \mathcal{D}_{\text{lg}}} \ell_{exp}(w^\top x, \operatorname{sgn}(w^{t\top} x))|_{w=w^t} \\ &= - \mathbb{E}_{x \sim \mathcal{D}_{\text{lg}}} \exp(-w^\top x \cdot \operatorname{sgn}(w^{t\top} x)) \cdot \operatorname{sgn}(w^{t\top} x) x|_{w=w^t} \\ &= - \mathbb{E}_{x \sim \mathcal{D}_{\text{lg}}} \exp(-w^\top x \cdot \operatorname{sgn}(w^\top x)) \cdot \operatorname{sgn}(w^\top x) x|_{w=w^t} \\ &= \nabla_w L(w)|_{w=w^t} \end{aligned}$$

Therefore for all $t \geq 0$, pseudo-labeling algorithm has the same iterate as entropy minimization (2.3). \square

E Additional experiments and details

E.1 Colored MNIST

Among 70K MNIST images, we split the source training / source test / target training / target test into 2:1:3:1. The model architecture is 3-layer feed-forward network with hidden layer sizes 128 and 64. For training on source, we use SGD optimizer with learning rate 0.03, momentum 0.9, weight decay 0.002, and always train until convergence.

Additional construction and training details for 10-way MNIST. For each source image, with probability p , we assign it a weight $w \sim \text{unif}[0.1k, 0.1k + 0.1)$ when image is digit k ; with probability $1 - p$, we assign $w \sim \text{unif}[0, 1)$. Each target image is assigned $w \sim \text{unif}[0, 1)$. We create two color channels by scaling the gray-scale image with weights w and $1 - w$.

In entropy minimization phase, we perform full gradient descent on target training set, with learning rate 0.03, momentum 0.9, weight decay 0.002, and train for 300 epochs when $p = 0.95$ and 50 epochs when $p = 0.97$.

Detailed construction of binary colored MNIST. In this setup, we assign digits 0-4 label 0 and 5-9 label 1. For each gray-scale image, we first draw a Gaussian random variable $\tilde{w} \sim \mathcal{N}(0, (0.5/3)^2)$. In source domain, with probability $p = 0.8$, example with label k is assigned with $w = 0.5 + (2k - 1)|\tilde{w}|$; with probability 0.2, $w = 0.5 + \tilde{w}$. In the target domain, we always have $w = 0.5 + \tilde{w}$. We create two color channels by rescaling the original image with weights w and $1 - w$.

For training, we keep all other hyper-parameters the same as the 10-way setting and only reduce the learning rate to 0.003.

Distribution of predictions conditioned on gray-scale image. We examine the effect of entropy minimization on each test example in binary MNIST experiment. For each gray-scale test image x_1 , we draw 1000 x_2 , i.e., $\tilde{w} \sim \mathcal{N}(0, (0.5/3)^2)$, and plot the distribution of logits $f(x_1, x_2)_y - f(x_1, x_2)_{1-y}$ where y is the true label of x_1 . According to our theory, if the distribution is concentrated around the positive side, entropy minimization would push the distribution to be more concentrated and positive.

Examples classified wrongly by source classifier ⁵ (because they were on the negative tail of the distribution) can be corrected due to this effect. Figure 3 is an example image where source classifier was wrong before training on target but corrected due to the explanation we provide. Conversely, examples classified right by source classifier (because they happen to be on the positive tail of the distribution) can turn wrong due to entropy minimization (see Figure 4). The success of entropy minimization relies on more examples concentrated on the positive than the negative side, i.e., source classifier has non-trivial target accuracy.

Distribution of mean activation. Figure 5 shows the distribution of $f(x_1, \bar{x}_2)$ before and after self-training for the binary MNIST experiment, where \bar{x}_2 indicates neutral color ($w = 0.5$). We see that qualitatively, the empirical distribution of μ has increasing mass far away from 0 throughout self-training, even for a multi-layer network. This is necessary for our theory, as seen in Figure 2.

Importance of non-trivial source classifier accuracy. We provide additional details on our study of 10-way colored MNIST when the spurious correlation probability is $p = 0.97$. In this setting, the source classifier has 98% test accuracy on source but only 72% on target. Entropy minimization initialized at \tilde{f} causes target accuracy to drop to 67% (see Figure 6 right).

E.2 CelebA dataset

We partition the celebA dataset [21] so that the source domain has a perfect correlation between gender and hair color: 1250 blond males, 1749 non-blond females. The target domain has 57K unlabeled examples with the same correlation between gender and hair color as in the original dataset (Figure 7).

⁵A particular \widehat{x}_2 is drawn for each x_1 in the target test set. If $f(x_1, \widehat{x}_2)_y - f(x_1, \widehat{x}_2)_{1-y} < 0$, f makes a wrong prediction.




Figure 3: Distribution of $f(x_1, x_2)_y - f(x_1, x_2)_{1-y}$ before (**left**) and after (**right**) self-training for a test image whose prediction turned from wrong to correct. Green line shows $f(x_1, \hat{x}_2)$ turning positive for the particular \hat{x}_2 in test set.

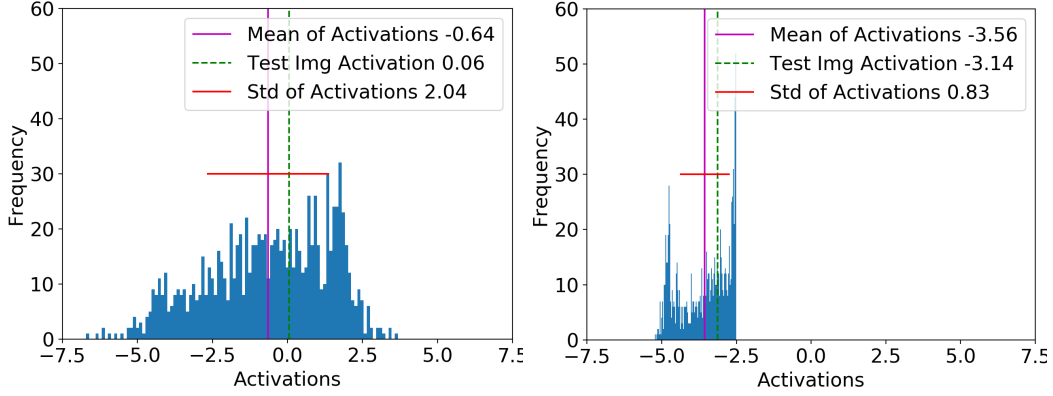


Figure 4: Distribution of $f(x_1, x_2)_y - f(x_1, x_2)_{1-y}$ before (**left**) and after (**right**) self-training for a test image whose prediction turned from right to wrong. Green line shows $f(x_1, \hat{x}_2)$ turning negative for the particular \hat{x}_2 in test set.




Figure 5: Distribution of $f(x_1, \bar{x}_2)$ before (**left**) and after (**right**) self-training across all test images x_1 for neutral color \bar{x}_2 . Qualitatively the empirical distribution of μ has more mass far away from 0 after self-training (which is the desired case for our theory, as seen in Figure 2).




Figure 6: In the 10-way MNIST experiment, entropy minimization raises target test accuracy by 9% when we initialize with a good source classifier (**left**) and decreases target accuracy when we initialize with a bad source classifier (**right**). **Left**: Spurious correlation in source is $p = 0.95$ so source classifier obtains high target accuracy; **Right**: Spurious correlation in source is $p = 0.97$ so source classifier does not learn the right features.



(a) Synthetic source data: blondness perfectly correlates with the male gender. (b) Synthetic target data: each gender has a variety of hair colors. (c) Predictions corrected by self-training were mostly mistaken due to the spurious correlation.

Figure 7: In the synthetic CelebA experiment, the source has perfect correlation between hair color and gender, and the target does not. A classifier trained only on the source domain uses the spurious correlation. However, continuing to self-train on the unlabeled target domain reduces reliance on the spurious feature.

We use entropy minimization on this dataset with the Conv-Small model in [24]. The source classifier has 94% accuracy on source data and 81% on target. After training on the sum of the source labeled loss and target entropy loss, the target accuracy increases to 88%.

E.3 Connection between entropy minimization and stochastic pseudo-labeling

In Equation ?? we point out that entropy minimization is equivalent to a stochastic version of pseudo-labeling where we update the pseudo-labels after every SGD step. In practice, pseudo-labels are often updated for only a few rounds, and the student model is usually trained to convergence between rounds [42]. In the 10-way MNIST experiment, we perform 3, 6, 30 rounds of pseudo-labeling with 100, 50, and 10 epochs of training per round, interpolating between more common versions of pseudo-labeling and entropy minimization. Figure 8 shows that entropy minimization converges to better target accuracy within the same clock-time, suggesting that practitioners may benefit from pseudo-labeling with more rounds and fewer epochs per round.

E.4 Toy Gaussian mixture setting

Generating data. We generate source examples in the following fashion: For each example $(x_1, x_2) \in \mathbb{R}^4$, we first sample y uniformly from $\{-1, 1\}$, and then $x_1 \in \mathbb{R}^2 \sim \mathcal{N}(\gamma y, I)$, where γ is a random 2-dimensional vector. For the source examples, we then sample $\tilde{x}_2 \sim N(\vec{0}, I)$. For each coordinate i of x_2 ($i \in \{1, 2\}$), with probability 0.8, we set $(x_2)_i = y|(\tilde{x}_2)_i|$ (correlated); with probability 0.2, we set $(x_2)_i = (\tilde{x}_2)_i$ (uncorrelated). For target examples, we sample $x_2 \sim N(\vec{0}, I)$.

The source training dataset, source test set, and target test set all have 10K examples.

Algorithms. We use entropy minimization as well as the following version of pseudo-labeling: starting with the source classifier, we perform 200 rounds of pseudo-labeling with 50 epochs of training in each round. We also set up a threshold $\tau = 0.1$ where we throw out least-confident target




Figure 8: In the 10-way MNIST experiment, 3 rounds of pseudo-labeling with 100 epochs per round (**left**), 6 rounds of pseudo-labeling with 50 epochs per round (**middle**), and 30 rounds of pseudo-labeling with 10 epochs per round (**right**) increase in target accuracy.

example x with $|w^\top x| < \tau$ in each round to mimic most popular pseudo-labeling algorithms used in practice [42]. We experiment on this version of pseudo-labeling algorithm because the version in equation ?? is equivalent to entropy minimization.

For entropy minimization, we use a new batch of 10K target training examples in each epoch. We use SGD optimizer with learning rate $1e-3$ and normalize the linear model after each gradient step.

For pseudo-labeling, we use a new batch of 10K target examples in each round. Optimizer choices are the same as entropy minimization.

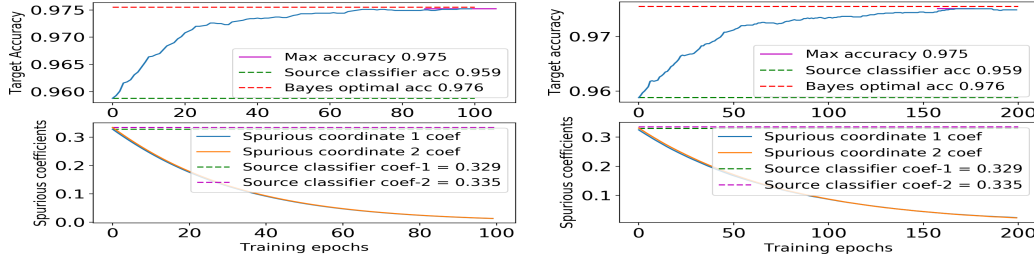


Figure 9: Entropy minimization (**left**) and pseudo-labeling (**right**) increase target test accuracy from 95.9% to 97.5%, and reduce coefficients on two spurious coordinates from 0.33 to 0 in the Gaussian mixture experiment.

Improvement of target test accuracy and de-emphasis of spurious features. In the Gaussian mixture experiment, the source classifier gets an accuracy of 95.9% on the target domain. Both entropy minimization and pseudo-labeling algorithms raise the target accuracy to Bayes-optimal while driving the coefficients w_2 on spurious features x_2 to 0 (Figure 9). Notably, even though we use confidence thresholding and train for 50 epochs in each round, the model behavior still closely tracks that of entropy minimization, as predicted by our theory.

E.5 Justification of approximation $l(t) = \exp(-|t|)$

Self-training on target using $l(t) = \exp(-|t|)$ as an approximation for $\ell_{ent}(t)$ produces the same effect for binary MNIST (see figure 11). We plot the training loss $l(t) = \exp(-|t|)$ and $\ell_{ent}(t)$ (Figure 12) to show that they track each other really well.




Figure 10: Plot of $\exp(-|t|)$ and entropy loss. The losses are within a constant factor of each other and exhibit the same tail behavior.




Figure 11: Target accuracy using $l(t) = \exp(-|t|)$ on binary MNIST dataset.




Figure 12: **Left:** Training loss using $l(t) = \exp(-|t|)$; **Right:** Entropy loss when training using $l(t) = \exp(-|t|)$.