REGULARIZED LEARNING FOR DOMAIN ADAPTATION UNDER LABEL SHIFTS

Anonymous authors
Paper under double-blind review

ABSTRACT

We propose Regularized Learning under Label shifts (RLLS), a principled and practical domain adaptation approach to correct for shifts in the label distribution between a source and a target domain. We first estimate importance weights using both labeled source and unlabeled target data before training a classifier on the weighted source samples. Depending on the number of source and target samples, we thereby regularize the influence of these estimated weights. We derive a generalization bound for the classifier on the target domain which is independent of the (ambient) data dimensions and instead depends on the complexity of the function class. To the best of our knowledge, this is the first generalization bound for label-shift problems where the labels in the target domain are unknown. Based on this bound, we propose a regularized estimator for the small-sample regime which accounts for the uncertainty in the estimated weights. Experiments on the CIFAR-10 and MNIST datasets show that RLLS improves classification accuracy in the low sample and large-shift regimes compared to previous methods.

1 INTRODUCTION

In supervised learning, we usually make predictions on an unlabeled *target* set by training a predictive model on a labeled *source* set. This is a sensible approach especially when the source and target sets are expected to have i.i.d. samples from the same distribution. However, this assumption does not hold in many real-world applications. For example, often times predictions for medical diagnostics are based on data from a particular population of patients and machines due to the limited variety of data that is available in the training phase.

Let us consider the following scenario. A classifier is trained to predict whether a patient has contracted a severe disease in country A based on the physical measurements and diagnosis (labeled source data). The disease is potentially even more prevalent in country B, where no experts are available to make a reliable diagnosis. The following questions could be relevant in this scenario:

- If we send a group of volunteers to country B with the preliminary equipment to record measurements of patients, how can the data in country A and the new unlabeled data from country B be used to update the classifier to give good diagnostic predictions for patients in country B?
- If our group members visit the patients in country B in online and daily manner rather than batch and the disease is deadly, can the classifier be updated after only few patients have been measured, so as to maximize the accuracy of the diagnosis for the following patients?

Similar problems arise in many other domains besides health care, but we will use this medical example throughout the paper for ease of presentation.

If we view the data points (X_i, Y_i) in the source and target set as i.i.d. samples from a joint distributions $\mathbb{P}$ and $\mathbb{Q}$ taking values in $\mathcal{X} \times \mathcal{Y}$, we are dealing with the *distributional shift* problem when $\mathbb{P} \neq \mathbb{Q}$. There are various ways to model the difference between $\mathbb{P}$ and $\mathbb{Q}$, depending on the use case at hand.Traditionally, the two following perspectives are considered: (i) the causal setting where the input x explains the label y, and (ii) the anti-causal setting, where the label y causes the input x.

For the causal setting, one usually assumes that the marginal distribution of the covariates x shifts between $\mathbb{P}$ and $\mathbb{Q}$, whereas the conditional probability mass function of y given x stays constant,

which we denote as $p(y|x) = q(y|x)$ for all $y \in \mathcal{Y}$ and $x \in \mathcal{X}$. Although this is valid in some situations, in many data science applications an inherent ground truth Y causes an observation X. For example, reported physical symptoms might be approximately similar across all patients that have contracted a particular disease. This simplified relation can be modeled e.g. by a directed graphical model $Y \to X$. It is now possible that the percentage of ill patients in the two countries A and B are very different, i.e. $\mathbb{P}(Y = y) \neq \mathbb{Q}(Y = y)$ ($p(y) \neq q(y)$ for short) for all $y \in \mathcal{Y}$. Similarly, within one country, the number of ill patients may suddenly change because of an irregular event such as a natural catastrophe.

Covariate Shift	Label Shift
$p(x) \neq q(x)$	$p(y) \neq q(y)$
$p(y\|x) = q(y\|x)$	$p(x\|y) = q(x\|y)$

In all of the above cases, there is a shift in the distribution of the label y, while given a fixed label, the conditional distribution of x given y is constant, i.e., $p(x|y) = q(x|y)$. By the Bayes rule, *label shift* causes the covariate shift condition to be violated, i.e. $p(y|x) \neq q(y|x)$. Despite being an intuitively natural scenario in some cases, even this simplified model has not yet been sufficiently studied in the literature. This paper aims to fill the gap in terms of both theoretical understanding and practical methods for the label shift setting and thereby move a step closer towards more complicated and realistic distributional shift models. In particular, our goal was to find an efficient method which is scalable to large-scale data where conventional approaches such as ones based on kernel mean embeddings, may become computationally feasible.

Lipton et al. (2018) introduced an importance-weight estimator, which has access to the labeled source data set and a large batch of unlabeled target data, and demonstrated promising performance on real data. Using a black-box classifier which can be biased, uncalibrated and inaccurate, they first estimate importance weights $q(y)/p(y)$ for the source samples and train a classifier on the weighted data. In the following we refer to the procedure as *black box shift learning* (BBSL) which the authors proved to be good for large enough sample size. There are two relevant questions which remain unanswered by their work: What are the generalization guarantees for the final predictor which uses the weighted samples? How do we deal with the uncertainty of the weight estimation when only few samples are available?

In this work, our first contribution is to propose an efficient weight estimator for which we can obtain good statistical guarantees without a requirement on the problem-dependent minimum sample complexity as necessary for BBSL, which in their case can be arbitrarily large for small number of samples. We then derive a dimension-independent generalization bound for the final Regularized Learning under Label Shift (RLLS) classifier trained on the reweighted samples. Our method improves the weight estimation error and the excess risk of the classifier on reweighted samples by a factor of $k \log(k)$, where k is the number of classes, i.e. the cardinality of $\mathcal{Y}$. It also follows from our analysis that explicitly controlling the influence of our regularized weight estimates could help in practice. In particular, it suggests that when the number of target samples is below a certain threshold, the classifier would obtain higher accuracy if the weight estimates are not trusted because of high uncertainty and the samples stay unweighted.

We empirically study the performance of RLLS and show weight estimation as well as prediction accuracy for a variety of shifts on the CIFAR-10 and MNIST datasets, including shifting source data and shifting target data with various parameters, different sample sizes and regularization parameters. We achieve an order of magnitude smaller weight estimation error than baseline methods in large shift cases and enjoy 20% higher accuracy and F-1 score in corresponding predictive tasks. We also see at least 10% accuracy improvement using our regularized estimator when faced with small target sample sizes.

2 Regularized learning of label shifts (RLLS)

Formally let us the short hand for the marginal probability mass functions of Y on finite $\mathcal{Y}$ with respect to $\mathbb{P}, \mathbb{Q}$ as $p, q : [k] \to [0, 1]$ with $p(i) = \mathbb{P}(Y = i)$, and $q(i) = \mathbb{Q}(Y = i)$ for all $i \in [k]$, representable by vectors in $\mathbb{R}^k_+$ which sum to one. In the label shift setting, we define the importance weight vector $w \in \mathbb{R}^k$ between these two domains as $w(i) = \frac{q(i)}{p(i)}$. We quantify the shift using the exponent of the infinite and second order Renyi divergence as follows

$$d_\infty(q||p) := \sup_i \frac{q(i)}{p(i)} \quad , \quad \textit{and} \quad d(q||p) := \mathbb{E}_{Y\sim\mathbb{Q}}\left[w(Y)^2\right] = \sum_i^k q(i)\frac{q(i)}{p(i)}.$$

Given a hypothesis class $\mathcal{H}$ and a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \to [0, 1]$, our aim is to find the hypothesis $h \in \mathcal{H}$ which minimizes

$$\mathcal{L}(h) = \mathbb{E}_{X,Y\sim\mathbb{Q}}\left[\ell(Y, h(X))\right] = \mathbb{E}_{X,Y\sim\mathbb{P}}\left[w(Y)\ell(Y, h(X))\right]$$

In the usual finite sample setting however, $\mathcal{L}$ unknown and we observe samples $\{(x_j, y_j)\}_{j=1}^n$ from $\mathbb{P}$ instead. If we are given the vector of importance weights w we could then minimize the empirical loss with importance weighted samples defined as

$$\mathcal{L}_n(h) = \frac{1}{n}\sum_{j=1}^{n} w(y_j)\ell(y_j, h(x_j))$$

where n is the number of available observations drawn from $\mathbb{P}$ used to learn the classifier h. As w is unknown in practice, we have to find the minimizer of the empirical loss with *estimated* importance weights

$$\mathcal{L}_n(h; \widehat{w}) = \frac{1}{n}\sum_{j=1}^{n} \widehat{w}(y_j)\ell(y_j, h(x_j)) \tag{1}$$

where $\widehat{w}$ are estimates of w. Given a set D_p of n_p samples from the source distribution $\mathbb{P}$, we first divide it into two sets where we use $(1-\beta)n_p$ samples in set D_p^{weight} to compute the estimate $\widehat{w}$ and the remaining $n = \beta n_p$ in the set D_p^{class} to find the classifier which minimizes the loss (1), i.e. $\widehat{h}_{\widehat{w}} = \arg\min_{h\in\mathcal{H}} \mathcal{L}_n(h; \widehat{w})$. In the following, we describe how to estimate the weights $\widehat{w}$ and provide guarantees for the resulting estimator $\widehat{h}_{\widehat{w}}$.

Plug-in weight estimation The following simple correlation between the label distributions p, q was noted in Lipton et al. (2018): for a fixed hypothesis h, if for all $y \in \mathcal{Y}$ it holds that $q(y) \geq 0 \implies p(y) \geq 0$, we have

$$\begin{aligned} q_h(Y = i) &= \sum_{j=1}^{k} q(h(X) = i|Y = j)q(j) = \sum_{j=1}^{k} p(h(X) = i|Y = j)q(j) \\ &= \sum_{j=1}^{k} p(h(X) = i, Y = j)\frac{q(j)}{p(j)} = \sum_{j=1}^{k} p(h(X) = i, Y = j)w_j \end{aligned}$$

for all $i, j \in \mathcal{Y}$. This can equivalently be written in matrix vector notation as

$$q_h = C_h w, \tag{2}$$

where C_h is the confusion matrix with $[C_h]_{i,j} = p(h(X) = i, Y = j)$ and q_h is the vector which represents the probability mass function of $h(X)$ under distribution $\mathbb{Q}$. The requirement $q(y) \geq 0 \implies p(y) \geq 0$ is a reasonable condition since without any prior knowledge, there is no way to properly reason about a class in the target domain that is not represented in the source domain.

In reality, both q_h and C_h can only be estimated by the corresponding finite sample averages $\widehat{q}_h, \widehat{C}_h$. Lipton et al. (2018) simply compute the inverse of the estimated confusion matrix $\widehat{C}_h$ in order to estimate the importance weight, i.e. $\widehat{w} = \widehat{C}_h^{-1}\widehat{q}_h$. While $C_h^{-1}\widehat{q}_h$ is a statistically efficient estimator, $\widehat{w}$ with estimated $\widehat{C}_h^{-1}$ can be arbitrarily bad since $\widehat{C}_h^{-1}$ can be arbitrary close to a singular matrix especially for small sample sizes and small minimum singular value of the confusion matrix. Intuitively, when there are very few samples, the weight estimation will have high variance in which case it might be better to avoid importance weighting altogether. Furthermore, even when the sample complexity in Lipton et al. (2018), unknown in practice, is met, the resulting error of this estimator is linear in k which is problematic for large k.

We therefore aim to address these shortcomings by proposing the following two-step procedure to compute importance weights. In the case of no shift we have $w = \mathbf{1}$ so that we define the amount of weight shift as $\theta = w - \mathbf{1}$. Given a "decent" black box estimator which we denote by h_0, we make the final classifier less sensitive to the estimation performance of C (i.e. regularize the weight estimate) by

1. calculating the measurement error adjusted $\widehat{\theta}$ (described in Section 2.1 for h_0) and
2. computing the regularized weight $\widehat{w} = \mathbf{1} + \lambda\widehat{\theta}$ where λ depends on the sample size $(1-\beta)n_p$.

By "decent" we refer to a classifier h_0 which yields a full rank confusion matrix C_{h_0}. A trivial example for a non-"decent" classifier h_0 is one that always outputs a fixed class. As it does not capture any characteristics of the data, there is no hope to gain any statistical information without any prior information.

2.1 Estimator correcting for finite sample errors

Both the confusion matrix C_{h_0} and the label distribution q_{h_0} on the target for the black box hypothesis h_0 are unknown and we are instead only given access to finite sample estimates $\widehat{C}_{h_0}, \widehat{q}_{h_0}$. In what follows all empirical and population confusion matrices, as well as label distributions, are defined with respect to the hypothesis $h = h_0$. For notation simplicity, we thus drop the subscript h_0 in what follows. The reparameterized linear model (2) with respect to θ then reads

$$b := q - C\mathbf{1} = C\theta$$

with corresponding finite sample quantity $\widehat{b} = \widehat{q} - \widehat{C}\mathbf{1}$. When $\widehat{C}$ is near singular, the estimation of θ becomes unstable. Furthermore, large values in the true shift θ result in large variances. We address this problem by adding a regularizing ℓ_2 penalty term to the usual loss and thus push the amount of shift towards 0, a method that has been proposed in (Pires & Szepesvári, 2012). In particular, we compute

$$\widehat{\theta} = \underset{\widehat{\theta}_\gamma : \gamma \in \{2^{i+1}\Delta_C\Delta_b, i \in \mathbb{N}\}}{\arg\min} \|\widehat{C}\widehat{\theta}_\gamma - \widehat{b}\|_2 + 3\Delta_C\|\widehat{\theta}_\gamma\|_2 \tag{3}$$

where $\widehat{\theta}_\gamma$ is defined by

$$\widehat{\theta}_\gamma = \underset{\theta'}{\arg\min}\ \|\widehat{C}\theta' - \widehat{b}\|_2^2 + \gamma\|\theta'\|_2.$$

Here, Δ_C, Δ_b are parameters which will eventually be high probability upper bounds for $\|\widehat{C} - C\|_2$ and $\|\widehat{b} - b\|_2$ respectively, where $\|\cdot\|_2$ denotes the spectral norm. Let $\sigma_{\min}$ be the minimum singular value of the true confusion matrix C. For the estimate $\widehat{\theta}$, we then have the following upper bound

Lemma 1 *For $\widehat{\theta}$ as defined in equation* (3)*, we have with probability at least* $1 - \delta$ *that*[1]

$$\|\widehat{\theta} - \theta\|_2 \leq \epsilon_\theta(n_p, n_q, \|\theta\|_2, \delta) := \mathcal{O}\left(\frac{1}{\sigma_{min}}\big(\|\theta\|_2\sqrt{\frac{\log(k/\delta)}{(1-\beta)n_p}} + \sqrt{\frac{\log(1/\delta)}{(1-\beta)n_p}} + \sqrt{\frac{\log(1/\delta)}{n_q}}\big)\right).$$

The proof of this lemma can be found in Appendix A.1. A couple of remarks are in order at this point. First of all, notice that the weight estimation procedure (3) does not require a minimum sample complexity which is in the order of $\sigma_{\min}^{-2}$ to obtain the guarantees for BBSL. This is due to the fact that errors in the covariates are accounted for. In order to directly see the improvements in the upper bound of Lemma 1 compared to Theorem 3 in Lipton et al. (2018), first observe that in order to obtain their upper bound with a probability of at least $1 - \delta$, it is necessary that $3kn_p^{-10} + 2kn_q^{-10} \leq \delta$. As a consequence, the upper bound in Theorem 3 of Lipton et al. (2018) is bigger than $\frac{1}{3\sigma_{\min}}\big(\|\theta\|_2\sqrt{\frac{\log(3k/\delta)}{n_p}} + \sqrt{\frac{k\log(2k/\delta)}{n_q}}\big)$. Thus Lemma 1 improves upon the previous upper bound by a factor of k.

Furthermore, as in Lipton et al. (2018), this result holds for *any* black box estimator h_0 which enters the bound via $\sigma_{\min}(C_{h_0})$. We can directly see how a good choice of h_0 helps to decrease the upper bound in Lemma 1. In particular, if h_0 is an ideal estimator, and the source set is balanced, C is the unit matrix with $\sigma_{\min} = 1/k$. In contrast, when the model h_0 is uncertain, the singular value $\sigma_{\min}$ is close to zero.

Moreover, for least square problems with Gaussian measurement errors in both input and target variables, it is standard to use regularized total least squares approaches which requires a singular

[1] Throughout the paper, $\mathcal{O}$ hides universal constant factors. Furthermore, we use $\mathcal{O}(\cdot + \cdot)$ for short to denote $\mathcal{O}(\cdot) + \mathcal{O}(\cdot)$.

Algorithm 1 Regularized Learning of Label Shift (RLLS)

1: Input: source set D_p, D_q, $\theta_{\max}$, estimate of $\sigma_{\min}$, black box estimator h_0, model class $\mathcal{H}$
2: Determine optimal split ratio $\beta^\star$ and regularizer $\lambda^\star$ by minimizing the RHS of Eq. (6) using an estimate of $\sigma_{\min}$
3: Randomly partition source set D_p into $D_p^{\text{class}}, D_p^{\text{weight}}$ such that $|D_p^{\text{class}}| = \beta^\star n_p =: n$
4: Compute $\widehat{\theta}$ using Eq. (3) and $\widehat{w} := 1 + \lambda^\star \widehat{\theta}$
5: Minimize the importance weighted empirical loss to obtain the weighted estimator

$$\widehat{h}_{\widehat{w}} = \arg\min_{h \in \mathcal{H}} \mathcal{L}_n(h; \widehat{w}), \qquad \textit{where} \qquad \mathcal{L}_n(h; \widehat{w}) = \frac{1}{n} \sum_{(x,y) \in D_p^{\text{class}}} \widehat{w}(y) \ell(y, h(x))$$

6: Deploy $\widehat{h}_{\widehat{w}}$ if the risk is acceptable

value decomposition. Finally, our choice for the alternative estimator in Eq. 3 with norm instead of norm squared regularization is motivated by the cases with large shifts θ, where using the squared norm may shrink the estimate $\widehat{\theta}$ too much and away from the true θ.

2.2 Regularized estimator and generalization bound

When a few samples from the target set are available or the label shift is mild, the estimated weights might be too uncertain to be applied. We therefore propose a regularized estimator defined as follows

$$\widehat{w} = \mathbf{1} + \lambda \widehat{\theta}. \tag{4}$$

Note that $\widehat{w}$ implicitly depends on λ, and β. By rewriting $\widehat{w} = (1-\lambda)\mathbf{1} + \lambda(\mathbf{1} + \widehat{\theta})$, we see that intuitively λ closer to 1 the more reason there is to believe that $\mathbf{1} + \widehat{\theta}$ is in fact the true weight.

Define the set $\mathcal{G}(\ell, \mathcal{H}) = \{g_h(x,y) = w(y)\ell(h(x), y) : h \in \mathcal{H}\}$ and its Rademacher complexity measure

$$\mathcal{R}_n(\mathcal{G}) := \mathbb{E}_{(X_i,Y_i)\sim\mathbb{P}:i\in[n]} \left[\mathbb{E}_{\xi_i:i\in[n]} \frac{1}{n} \left[\sup_{h\in\mathcal{H}} \sum_{i=1}^n \xi_i g_h(X_i, h(Y_i)) \right] \right]$$

with $\xi_i,\ \forall i$ as the Rademacher random variables (see e.g. Bartlett & Mendelson (2002)). We can now state a generalization bound for the classifier $\widehat{h}_{\widehat{w}}$ in a general hypothesis class $\mathcal{H}$, which is trained on source data with the estimated weights defined in equation (4).

Theorem 1 (Generalization bound for $\widehat{h}_{\widehat{w}}$) *Given n_p samples from the source data set and n_q samples from the target set, a hypothesis class $\mathcal{H}$ and loss function ℓ, for $\epsilon_{\mathcal{G}}(n_p, \delta) := 2\mathcal{R}_n(\mathcal{G}) + d_\infty(q||p)\sqrt{\frac{\log(2/\delta)}{\beta n_p}}$ the following upper bound holds with probability at least $1-\delta$*

$$\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}(h^*) \leq \epsilon_{\mathcal{G}}(n_p, \delta, \beta) + (1-\lambda)\,\|\theta\|_2 + \lambda \epsilon_\theta(n_p, n_q, \|\theta\|_2, \delta, \beta). \tag{5}$$

The proof can be found in Appendix A.4. Additionally, we derive the analysis also for finite hypothesis classes in Appendix A.5 to provide more insight into the proof of general hypothesis classes. The size of $\mathcal{R}_n(\mathcal{G})$ is determined by the structure of the function class $\mathcal{H}$ and the loss ℓ. For example for the $0/1$ loss, the VC dimension of $\mathcal{H}$ can be deployed to upper bound the Rademacher complexity.

The bound (5) in Theorem 1 holds for all choices of λ. In order to exploit the possibility of choosing λ and β to have an improved accuracy depending on the sample sizes, we first let the user define a set of shifts θ against which we want to be robust against, i.e. all shifts with $\|\theta\|_2 \leq \theta_{\max}$. For these shifts, we obtain the following upper bound

$$\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}(h^*) \leq \epsilon_{\mathcal{G}}(n_p, \delta) + (1-\lambda)\,\theta_{\max} + \lambda \epsilon_\theta(n_p, n_q, \theta_{\max}, \delta) \tag{6}$$

The bound in equation (6) suggests using Algorithm 1 as our ultimate label shift correction procedure. where for step 2 of the algorithm, we choose $\lambda^\star = 1$ whenever $n_q \geq \frac{1}{\theta_{\max}^2(\sigma_{\min} - \frac{1}{\sqrt{n_p}})^2}$ (hereby neglecting the log factors and thus dependencies on k) and 0 else. When using this rule, we obtain

$$\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}(h^*) \leq \epsilon_{\mathcal{G}}(n_p, \delta) + \max\{\theta_{\max}, \epsilon_\theta(n_p, n_q, \theta_{\max}, \delta)\}. \tag{7}$$

Notice that in practice, we do not know $\sigma_{\min}$ in advance so that in Algorithm 1 we need to use an estimate of $\sigma_{\min}$, which could e.g. be the minimum eigenvalue of the empirical confusion matrix $\widehat{C}$ with an additional computational complexity of at most $O(k^3)$.

Figure 1 shows how the oracle thresholds vary with n_q and $\sigma_{\min}$ when n_p is kept fix. When the parameters are above the curves for fixed n_p, λ should be chosen as 1 otherwise the samples should be unweighted, i.e. $\lambda = 0$. This figure illustrates that when the confusion matrix has small singular values, the estimated weights should only be trusted for rather high n_q and high believed shifts $\theta_{\max}$. Although the overall statistical rate of the excess risk of the classifier does not change as a function of the sample sizes, $\theta_{\max}$ could be significantly smaller than ϵ_θ when $\sigma_{\min}$ is very small and thus the accuracy in this regime could improve. Indeed we observe this to be the case empirically in Section 3.3.

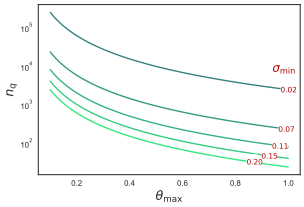

Figure 1: Given a $\sigma_{\min}$ and $\theta_{\max}$, λ switches from 0 to 1 at a particular n_q. n_p and k are fixed.

3 EXPERIMENTS

In this section we illustrate the theoretical analysis by running RLLS on a variety of artificially generated shifts on the MNIST (LeCun & Cortes, 2010) and CIFAR10 (Krizhevsky & Hinton, 2009) datasets. We first randomly separate the entire dataset into two sets (source and target pool) of the same size. Then we sample, unless specified otherwise, the same number of data points from each pool to form the source and target set respectively. We chose to have equal sample sizes to allow for fair comparisons across shifts.

There are various kinds of shifts which we consider in our experiments. In general we assume one of the source or target datasets to have uniform distribution over the labels. Within the non-uniform set, we consider two types of sampling strategies in the main text: the *Minority-Class Shift* refers to the case where we set a fixed number of classes m to have probability $p < 0.1$, while the distribution over the rest of the classes is uniform. For the *Dirichlet shift*, we draw a probability vector p from the Dirichlet distribution with concentration parameter set to α for all classes, before including sample points which correspond to the multinomial label variable according to p. Results for the tweak-one shift strategy as in Lipton et al. (2018) can be found in Section A.6.1.

After artificially shifting the label distribution in one of the source and target sets, we then follow algorithm 1, where we choose the black box predictor h_0 to be a two-layer fully connected neural network trained on the source dataset. Note that any black box predictor could be employed here, though the higher the accuracy, the more likely weight estimation will be precise.

In order to compute $\widehat{\omega} = \mathbf{1} + \widehat{\theta}$ in Eq. (3), performing a grid search over γ on the entire set of natural numbers is not feasible. Instead, we call a built-in solver to directly solve the low dimensional problem $\min_\theta \|\widehat{C}\theta - \widehat{b}\|_2 + \gamma 3\Delta_C \|\theta\|_2$ where the choice of $\gamma = 0.01$ is observed to yield a low loss in the first line of Eq. (3) on various levels of label shift pre-computed beforehand. We thus treat it as a hyperparameter that can be chosen using standard cross validation methods. Finally, we train a classifier on the source samples weighted by $\widehat{\omega}$, where we use a two-layer fully connected neural network for MNIST and a ResNet-18 (He et al., 2016) for CIFAR10.

We sample 20 datasets with the label distributions for each shift parameter. to evaluate the empirical mean square estimation error (MSE) and variance of the estimated weights $\mathbb{E}\|\widehat{w} - w\|_2^2$ and the predictive accuracy on the target set. We use these measures to compare our procedure with the black box shift learning method (BBSL) in Lipton et al. (2018). Notice that although KMM methods (Zhang et al., 2013) would be another standard baseline to compare with, it is not scalable to large sample size regimes for n_p, n_q above $n = 8000$ as mentioned by Lipton et al. (2018).

3.1 WEIGHT ESTIMATION AND PREDICTIVE PERFORMANCE FOR SOURCE SHIFT

In this set of experiments on the MNIST dataset, we illustrate our weight estimation and prediction performance for Minority-Class source shifts and compare it with BBSL. For this set of experiments, we set the number of data points in both source and target set to 1000 and sample from the two pools without replacement.

Figure 2 illustrates the weight estimation alongside final classification performance for Minority-Class source shift of MNIST. We created large shifts of three or more minority classes with $p = 0.005$. We use a fixed black-box classifier that is trained on biased source data, with tweak-one $\rho = 0.5$. Observe that the MSE in weight estimation is relatively large and RLLS outperforms BBSL as the number of minority classes increases. As the shift increases the performance for all methods deteriorates. Furthermore, Figure 2 (b) illustrates how the advantage of RLLS over the unweighted classifier increases as the shift increases. Across all shifts, the RLLS based classifier yields higher accuracy than the one based on BBSL. Results for smaller and larger values of p can be found in Section A.6.2.

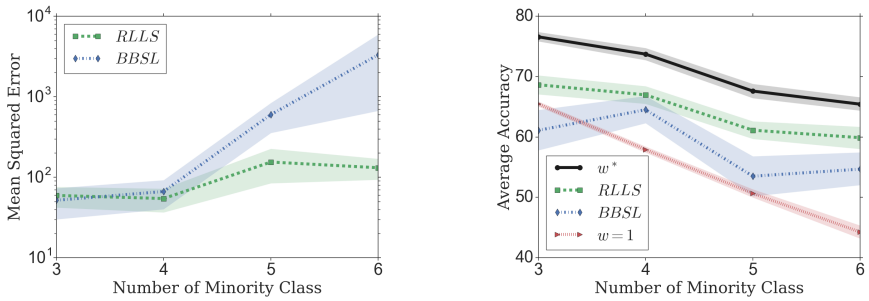

(a) (b)

Figure 2: (a) Mean squared error in estimated weights and (b) accuracy on MNIST for minority-class shifted source and uniform target with p = 0.005.

3.2 WEIGHT ESTIMATION AND PREDICTIVE PERFORMANCE FOR TARGET SHIFT

In this section, we compare the predictive performances between a classifier trained on unweighted source data and the classifiers trained on weighted loss obtained by the RLLS and BBSL procedure on CIFAR10. The target set is shifted using the Dirichlet shift with parameters $\alpha = [0.01, 0.1, 1, 10]$. The number of data points in both source and target set is 10000.

In the case of target shifts, larger shifts actually make the predictive task easier, such that even a constant majority class vote would give high accuracy. However it would have zero accuracy on all but one class. Therefore, in order to allow for a more comprehensive performance performance between the methods, we also compute the macro-averaged *F-1 score* by averaging the per-class quantity $2(\text{precision} \cdot \text{recall})/(\text{precision} + \text{recall})$ over all classes. For a class i, *precision* is the percentage of correct predictions among all samples predicted to have label i, while *recall* is the proportion of correctly predicted labels over the number of samples with true label i. This measure gives higher weight to the accuracies of minority classes which have no effect on the total accuracy.

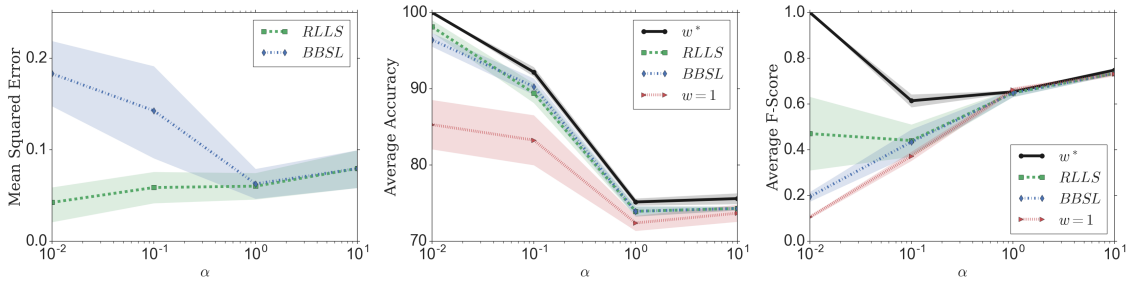

(a) (b) (c)

Figure 3: (a) Mean squared error in estimated weights, (b) accuracy and (c) F-1 score on CIFAR10 for uniform source and Dirichlet shifted target.

Figure 3 depicts the MSE of the weight estimation (a), the corresponding performance comparison on accuracy (b) and F-1 score (c). Recall that the accuracy performance for low shifts is not comparable with standard CIFAR10 benchmark results because of the overall lower sample size chosen for the comparability between shifts. We can see that in the large target shift case for $\alpha = 0.01$, the F-1 score for BBSL and the unweighted classifier is rather low compared to RLLS while the accuracy is

high. As mentioned before, the reason for this observation and why in Figure 3 (b) the accuracy is higher when the shift is larger, is that the predictive task actually becomes easier with higher shift.

In conclusion, our proposed RLLS procedure yields better weight estimation and predictive performance for a whole range of source and target shifts, while the improvement is bigger for bigger shifts.

3.3 Regularized weights in the low sample regime for source shift

In the following, we present the average accuracy of RLLS in Figure 4 as a function of the number of target samples n_q for different values of λ for small n_q. Here we fix the sample size in the source set to $n_p = 1000$ and investigate a Minority-Class source shift with fixed $p = 0.01$ and five minority classes.

A motivation to use intermediate λ is discussed in Section 2.2, as λ in equation (4) may be chosen according to $\theta_{\max}, \sigma_{\min}$. In practice, since $\theta_{\max}$ is just an upper bound on the true amount of shift $\|\theta\|_2$, in some cases λ should in fact ideally be 0 when $\frac{1}{\theta_{\max}^2(\sigma_{\min}-\frac{1}{\sqrt{n_q}})^2} \leq n_q \leq \frac{1}{\|\theta\|_2(\sigma_{\min}-\frac{1}{\sqrt{n_q}})^2}$. Thus for target sample sizes n_q that are a little bit above the threshold (depending on the certainty of the belief how close to $\theta_{\max}$ the norm of the shift is believed to be), it could be sensible to use an intermediate value $\lambda \in (0, 1)$.

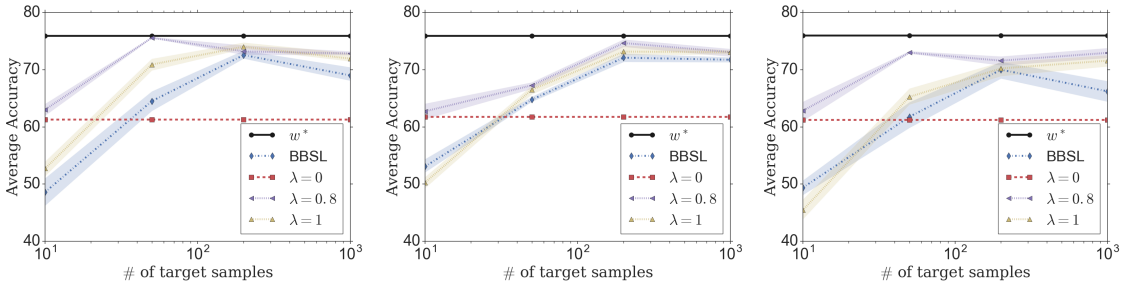

(a) (b) (c)

Figure 4: Performance on MNIST for Minority-Class shifted source and uniform target with various target sample size and λ using (a) better predictor h_0 trained on tweak-one shifted source with $\rho = 0.2$, (b) neutral predictor h_0 with $\rho = 0.5$ and (c) corrupted predictor h_0 with $\rho = 0.8$.

Figure 4 suggests that unweighted samples (red) yield the best classifier for very few samples n_q, while for $10 \leq n_q \leq 500$ an intermediate $\lambda \in (0, 1)$ (purple) has the highest accuracy and for $n_q > 1000$, the weight estimation is certain enough for the fully weighted classifier (yellow) to have the best performance (see also the corresponding data points in Figure 2). The unweighted BBSL classifier is also shown for completeness. We can conclude that regularizing the influence of the estimated weights allows us to adjust to the uncertainty on importance weights and generalize well for a wide range of target sample sizes.

Furthermore, the different plots in Figure 4 correspond to black-box predictors h_0 for weight estimation which are trained on more or less corrupted data, i.e. have a better or worse conditioned confusion matrix. The fully weighted methods with $\lambda = 1$ achieve the best performance faster with a better trained black-box classifier (a), while it takes longer for it to improve with a corrupted one (c). Furthermore, this reflects the relation between eigenvalue of confusion matrix $\sigma_{\min}$ and target sample size n_q in Theorem 1. In other words, we need more samples from the target data to compensate a bad predictor in weight estimation. So the generalization error decreases faster with an increasing number of samples for good predictors.

4 Related Work

The task of being robust against distributional shifts has been widely studied in the literature. One line of work focuses on distributional robustness (see e.g. Esfahani & Kuhn; Namkoong & Duchi (2016)) without explicitly modeling the shift by making assumptions on the distribution. Our theoretical

results adapt to any shift θ. However, these methods achieve robustness only against tiny changes in the distribution and are often impractical for large-scale high dimensional data.

Another line of work has addressed the causal setting where the conditional distribution of Y given X is assumed to be constant. This assumption is sensible when there is reason to believe that there is a true optimal mapping from X to Y which does not change if the distribution of X changes. In our example, if a given set of symptoms correctly predicts a disease for the "right reasons", the predictive function should be the same, no matter which country you are moving to. Among the various methods to correct for covariate shift, the majority uses the concept of importance weights $q(x)/p(x)$ (Zadrozny, 2004; Cortes et al., 2010; Cortes & Mohri, 2014; Shimodaira, 2000), which are unknown but can be estimated for example via kernel embeddings (Huang et al., 2007; Gretton et al., 2009; 2012; Zhang et al., 2013; Zaremba et al., 2013) or by learning a binary discriminative classifier between source and target (Lopez-Paz & Oquab, 2016; Liu et al., 2017). Finally, Sanderson & Scott (2014); Ramaswamy et al. (2016) formulate the label shift problem as a mixture of the class conditional covariate distributions with unknown mixture weights. Under the pairwise mutual irreducibility (Scott et al., 2013) assumption on the class conditional covariate distributions, they deploy the Neyman-Pearson criterion (Blanchard et al., 2010) to estimate the class distribution $q(y)$. Common issues shared by these methods is that they either result in a massive computational burden for large scale problems or cannot be deployed for neural networks.

Despite the wide applicability of label shift, approaches with global guarantees in high dimensional data regimes remain under-explored. The correction of label shift mainly requires to estimate the importance weights $q(y)/p(y)$ over the labels which typically live in a very low-dimensional space. Bayesian approaches are studied when a prior over the marginal label distribution is assumed (Storkey, 2009; Chan & Ng, 2005). However, these methods often require explicitly computing the posterior distribution of y and suffer from the curse of dimensionality. Recently, for the label shift setting, Lipton et al. (2018) proposed Black Box Shift Estimation (BBSL) which can be applied to large scale data. This approach is related to Buck et al. (1966); Forman (2008); Saerens et al. (2002) in the low dimensional setting but lacks guarantees for the excess risk.

On the other hand, generalization theory, usually used to bound the excess loss, has historically been developed for the case when $\mathbb{P} = \mathbb{Q}$ (see e.g. Vapnik (1999); Bartlett & Mendelson (2002); Kakade et al. (2009); Wainwright (2019)). Ben-David et al. (2010) provides theoretical analysis and generalization guarantees for distribution shifts when the H-divergence between joint distributions is considered, whereas Crammer et al. (2008) proves generalization bounds for learning from multiple sources. For the covariate shift setting, Cortes et al. (2010) provides a generalization bound when $q(x)/p(x)$ is known which however does not apply in practice.

5 Discussion

In this work, we establish the first generalization guarantee for the label shift setting and propose an importance weighting procedure for which no prior knowledge of $q(y)/p(y)$ is required. Although RLLS is inspired by BBSL, it leads to a more robust importance weight estimator as well as generalization guarantees in particular for the small sample regime, which BBSL does not allow for. RLLS is also equipped with a sample-size-dependent regularization technique and further improves the classifier in both regimes.

We consider this work a necessary step in the direction of solving shifts of this type, although the label shift assumption itself might be too simplified in the real world. In future work, we plan to also study the setting when it is slightly violated. For instance, x in practice cannot be solely explained by the wanted label y, but may also depend on attributes z which might not be observable. In the disease prediction task for example, the symptoms might not only depend on the disease but also on the country and living conditions of its population. In such a case, the label shift assumption only holds in a slightly modified sense, i.e. $\mathbb{P}(X|Y = y, Z = z) = \mathbb{Q}(X|Y = y, Z = z)$. If the attributes Z are observed, then our framework can readily be used to perform importance weighting.

Furthermore, it is not clear whether the final predictor is in fact "better" or more robust to shifts just because it achieves a better target accuracy than a vanilla unweighted estimator. In fact, there is reason to believe that under certain shift scenarios, the predictor might learn to use spurious correlations to boost accuracy. Finding a procedure which can both learn a robust model and achieve high accuracies

on new target sets remains to be an ongoing challenge. Moreover, the current choice of regularization depends on the number of samples rather than data driven regularization which is more desirable.

An interesting direction towards active learning for the same disease-symptoms scenario is when we also have an expert for diagnosing a limited number of patients in country B. Now the question is which patients would be most "useful" to diagnose to obtain a high accuracy on the entire target set? Furthermore, in the case of high risk, we might be able to send some of the patients to country A for further medical diagnosis or treatment, up to some varying cost. We plan to extend the current framework to the active learning setting where we actively query the label of certain x's (Beygelzimer et al., 2009) as well as the cost-sensitive setting where we also consider the cost of querying labels (Krishnamurthy et al., 2017).

REFERENCES

Animashree Anandkumar, Daniel Hsu, and Sham M Kakade. A method of moments for mixture models and hidden markov models. In *Conference on Learning Theory*, pp. 33–1, 2012.

Kamyar Azizzadenesheli, Alessandro Lazaric, and Animashree Anandkumar. Reinforcement learning of pomdps using spectral methods. *arXiv preprint arXiv:1602.07764*, 2016.

Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002.

Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.

Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance weighted active learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 49–56. ACM, 2009.

Gilles Blanchard, Gyemin Lee, and Clayton Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11(Nov):2973–3009, 2010.

AA Buck, JJ Gart, et al. Comparison of a screening test and a reference test in epidemiologic studies. ii. a probabilistic model for the comparison of diagnostic tests. *American Journal of Epidemiology*, 83(3):593–602, 1966.

Yee Seng Chan and Hwee Tou Ng. Word sense disambiguation with distribution estimation. In *IJCAI*, volume 5, pp. 1010–5, 2005.

Corinna Cortes and Mehryar Mohri. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519:103–126, 2014.

Corinna Cortes, Yishay Mansour, and Mehryar Mohri. Learning bounds for importance weighting. In *Advances in neural information processing systems*, pp. 442–450, 2010.

Koby Crammer, Michael Kearns, and Jennifer Wortman. Learning from multiple sources. *Journal of Machine Learning Research*, 9(Aug):1757–1774, 2008.

Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, pp. 1–52.

George Forman. Quantifying counts and costs via classification. *Data Mining and Knowledge Discovery*, 17(2):164–206, 2008.

Arthur Gretton, Alexander J Smola, Jiayuan Huang, Marcel Schmittfull, Karsten M Borgwardt, and Bernhard Schölkopf. Covariate shift by kernel mean matching. 2009.

Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(Mar):723–773, 2012.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Daniel Hsu, Sham M Kakade, and Tong Zhang. A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences*, 78(5):1460–1480, 2012.

Jiayuan Huang, Arthur Gretton, Karsten M Borgwardt, Bernhard Schölkopf, and Alex J Smola. Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pp. 601–608, 2007.

Sham M Kakade, Karthik Sridharan, and Ambuj Tewari. On the complexity of linear prediction: Risk bounds, margin bounds, and regularization. In *Advances in neural information processing systems*, pp. 793–800, 2009.

Akshay Krishnamurthy, Alekh Agarwal, Tzu-Kuo Huang, Hal Daume III, and John Langford. Active learning for cost-sensitive classification. *arXiv preprint arXiv:1703.01014*, 2017.

Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

Yann LeCun and Corinna Cortes. MNIST handwritten digit database. 2010. URL `http://yann.lecun.com/exdb/mnist/`.

Zachary C Lipton, Yu-Xiang Wang, and Alex Smola. Detecting and correcting for label shift with black box predictors. *arXiv preprint arXiv:1802.03916*, 2018.

Song Liu, Akiko Takeda, Taiji Suzuki, and Kenji Fukumizu. Trimmed density ratio estimation. In *Advances in Neural Information Processing Systems*, pp. 4518–4528, 2017.

David Lopez-Paz and Maxime Oquab. Revisiting classifier two-sample tests. *arXiv preprint arXiv:1610.06545*, 2016.

Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In *Advances in Neural Information Processing Systems*, pp. 2208–2216, 2016.

Bernardo Avila Pires and Csaba Szepesvári. Statistical linear estimation with penalized estimators: an application to reinforcement learning. *arXiv preprint arXiv:1206.6444*, 2012.

Harish Ramaswamy, Clayton Scott, and Ambuj Tewari. Mixture proportion estimation via kernel embeddings of distributions. In *International Conference on Machine Learning*, pp. 2052–2060, 2016.

Marco Saerens, Patrice Latinne, and Christine Decaestecker. Adjusting the outputs of a classifier to new a priori probabilities: a simple procedure. *Neural computation*, 14(1):21–41, 2002.

Tyler Sanderson and Clayton Scott. Class proportion estimation with application to multiclass anomaly rejection. In *Artificial Intelligence and Statistics*, pp. 850–858, 2014.

Clayton Scott, Gilles Blanchard, and Gregory Handy. Classification with asymmetric label noise: Consistency and maximal denoising. In *Conference On Learning Theory*, pp. 489–511, 2013.

Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.

Amos Storkey. When training and test sets are different: characterizing learning transfer. *Dataset shift in machine learning*, pp. 3–28, 2009.

Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.

Vladimir Naumovich Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999.

M. J. Wainwright. *High-dimensional statistics: A non-asymptotic viewpoint*. Cambridge University Press, 2019.

Bianca Zadrozny. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, pp. 114. ACM, 2004.

Wojciech Zaremba, Arthur Gretton, and Matthew Blaschko. B-test: A non-parametric, low variance kernel two-sample test. In *Advances in neural information processing systems*, pp. 755–763, 2013.

Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, pp. 819–827, 2013.

A PROOFS

A.1 PROOF OF LEMMA 1

From Thm. 3.7 in (Pires & Szepesvári, 2012) we know that for $\widehat{\theta}$ as defined in equation (3), if with probability at least $1-\delta$, $\|\widehat{C}-C\|_2 \leq \Delta_C$ and $\|\widehat{b}-b\|_2 \leq \Delta_b$ hold simultaneously, then

$$\Upsilon(\widehat{\theta}) \leq \inf_{\theta' \in \mathcal{R}^k} \{\Upsilon(\theta') + 3\Delta_C \|\theta'\|_\alpha\} + 3\Delta_b. \tag{8}$$

where we use the shorthand $\Upsilon(\theta') = \|C\theta' - b\|_2$.

We can get an upper bound on the right hand side of (8) is the infimum by simply choosing a feasible $\theta' = \theta$. We then have $\|C\theta - b\|_2 = 0$ and hence

$$\inf_{\theta'} \{\Upsilon(\theta') + 3\Delta_C \|\theta'\|_2\} \leq 3\Delta_C \|\theta\|_2$$

as a consequence,

$$\Upsilon(\widehat{\theta}) = \|C\widehat{\theta} - b\|_2 = \|C\left(\widehat{\theta} - \theta\right)\|_2 \leq 3\Delta_C \|\theta\|_2 + 3\Delta_b$$

Since $\|C\left(\widehat{\theta} - \theta\right)\|_2 \geq \sigma_{\min}(C)\|\widehat{\theta} - \theta\|_2$ by definition of the minimum singular value, we thus have

$$\|\widehat{\theta} - \theta\|_2 \leq \frac{1}{\sigma_{min}} \left(3\Delta_C \|\theta\|_2 + 3\Delta_b\right)$$

Let us first notice that

$$b_h = q_h - C\mathbf{1} = q_h - p_h$$

The mathematical definition of the finite sample estimates $\widehat{C}_h, \widehat{b}_h$ (in matrix and vector representation) with respect to some hypothesis h are as follows

$$[\widehat{C}_h]_{ij} = \frac{1}{(1-\beta)n_p} \sum_{(x,y)\in D_p} \mathbb{I}_{h(x)=i, y=j}$$

$$[\widehat{b}_h](i) = \frac{1}{m} \sum_{(x,y)\in D_q} \mathbb{I}_{h(x)=i} - \frac{1}{(1-\beta)n_p} \sum_{(x,y)\in D_p^{\text{weight}}} \mathbb{I}_{h(x)=i}$$

where $m = |D_q|$ and $\mathbb{I}$ is the indicator function. C_h, b_h can equivalently be expressed with the population over $\mathbb{P}$ for C_h and over $\mathbb{Q}$ for b_h respectively. We now use the following concentration Lemmas to bound the estimation errors of $\widehat{C}, \widehat{b}$ where we drop the subscript h for ease of notation.

Lemma 2 (Concentration of measurement matrix $\widehat{C}$) *For finite sample estimate $\widehat{C}$ we have*

$$\|\widehat{C} - C\|_2 \leq \frac{2\log(2k/\delta)}{3(1-\beta)n_p} + \sqrt{\frac{2\log(2k/\delta)}{(1-\beta)n_p}}$$

with probability at least $(1-\delta)$.

Lemma 3 (Concentration of label measurements) *For the finite sample estimate $\widehat{b}$ with respect to any hypothesis h it holds that*

$$\|\widehat{b}_h - b_h\|_2 \leq \frac{2}{\sqrt{\log(2)}} \left(\frac{\sqrt{\log(1/\delta)}}{\sqrt{(1-\beta)n_p}} + \frac{\sqrt{\log(1/\delta)}}{\sqrt{n_q}} \right)$$

with probability at least $1-2\delta$.

By Lemma. 2 for concentration of C and Lemma. 3 for concentration of b we now have with probability at least $1-\delta$

$$\|\widehat{\theta} - \theta\|_2 \leq \frac{1}{\sigma_{min}} \Bigg(\frac{2\|\theta\|_2 \log(2k/\delta)}{(1-\beta)n_p} + \|\theta\|_2 \sqrt{\frac{18\log(4k/\delta)}{(1-\beta)n_p}} + \sqrt{\frac{36\log(2/\delta)}{n_q}} + \sqrt{\frac{36\log(2/\delta)}{(1-\beta)n_p}} \Bigg).$$

which, considering that $O(\frac{1}{\sqrt{n}})$ dominates $O(\frac{1}{n})$, yields the statement of the Lemma 1.

A.2 PROOF OF LEMMA 2

We prove this lemma using the theorem 1.4[Matrix Bernstein] and Dilations technique from Tropp (2012). We can rewrite $C_h = \mathbb{E}_{(x,y)\sim\mathbb{P}}\left[e_{h(x)}e_y^\top\right]$ where $e(i)$ is the one-hot-encoding of index i. Consider a finite sequence $\{\Psi(i)\}$ of independent random matrices with dimension k. By dilations, lets construct another sequence of self-adjoint random matrices of $\{\widetilde{\Psi}(i)\}$ of dimension $2k$, such that for all i

$$\widetilde{\Psi}(i) = \begin{bmatrix} \mathbf{0} & \Psi(i) \\ \Psi(i)^\top & \mathbf{0} \end{bmatrix}$$

therefore,

$$\widetilde{\Psi}(i)^2 = \begin{bmatrix} \Psi(i)\Psi(i)^\top & \mathbf{0} \\ \mathbf{0} & \Psi(i)^\top\Psi(i) \end{bmatrix} \tag{9}$$

which results in $\|\widetilde{\Psi}(i)\|_2 = \|\Psi(i)\|_2$. The dilation technique translates the initial sequence of random matrices to the sequence of random self-adjoint matrices where we can apply the Matrix Bernstein theorem which states that, for a finite sequence of i.i.d. self-adjoint matrices $\widetilde{\Psi}(i)$, such that, almost surely $\forall i,\ \mathbb{E}\left[\widetilde{\Psi}(i)\right] = 0$ and $\|\widetilde{\Psi}(i)\| \leq R$, then for all $t \geq 0$,

$$\|\frac{1}{t}\sum_{i=1}^{t}\widetilde{\Psi}(i)\| \leq \frac{R\log(2k/\delta)}{3t} + \sqrt{\frac{2\varrho^2\log(2k/\delta)}{t}}$$

with probability at least $1-\delta$ where $\varrho^2 := \|\mathbb{E}\left[\widetilde{\Psi}^2(i)\right]\|_2,\ \forall i$ which is also $\varrho^2 = \|\mathbb{E}\left[\Psi^2(i)\right]\|_2,\ \forall i$ due to Eq. 9. Therefore, thanks to the dilation trick and theorem 1.4[Matrix Bernstein] in Tropp (2012),

$$\|\frac{1}{t}\sum_{i=1}^{t}\Psi(i)\| \leq \frac{R\log(2k/\delta)}{3t} + \sqrt{\frac{2\varrho^2\log(2k/\delta)}{t}}$$

with probability at least $1-\delta$.

Now, by plugging in $\Psi(i) = e_{h(x(i))}e_{y(i)}^\top - C$, we have $\mathbb{E}\left[\widetilde{\Psi}(i)\right] = 0$. Together with $\|\widetilde{\Psi}(i)\| \leq 2$ as well as $\varrho^2 = \|\mathbb{E}\left[\Psi^2(i)\right]\|_2 = 1$ and setting $t = n$, we have

$$\|\widehat{C} - C\|_2 \leq \frac{2\log(2k/\delta)}{3(1-\beta)n} + \sqrt{\frac{2\log(2k/\delta)}{(1-\beta)n}}$$

A.3 PROOF OF LEMMA 3

The proof of this lemma is mainly based on a special case of and appreared at proposition 6 in Azizzadenesheli et al. (2016), Lemma F.1 in Anandkumar et al. (2012) and Proposition 19 of Hsu et al. (2012).

Analogous to the previous section we can rewrite $b_h = \mathbb{E}_{(x,y)\in\mathbb{Q}}[e_{h(x)}] - \mathbb{E}_{(x,y)\in\mathbb{P}}[e_{h(x)}]$ where $e(i)$ is the one-hot-encoding of index i. Note that (dropping the subscript h) we have

$$\|\widehat{b}_h - b_h\|_2 \leq \|\widehat{q}_h - q_h\|_2 + \|\widehat{p}_h - p_h\|_2$$

We now bound both estimates of probability vectors separately.

Consider a fixed multinomial distribution characterized with probability vector of $\overline{\varsigma} \in \Delta_{k-1}$ where Δ_{k-1} is a $k-1$ dimensional simplex. Further, consider t realization of this multinomial distribution $\{\varsigma(i)\}_{i=1}^t$ where $\varsigma(i)$ is the one-hot-encoding of the i'th sample. Consider the empirical estimate mean of this distribution through empirical average of the samples; $\widehat{\varsigma} = \frac{1}{t}\sum(i)^t\varsigma(i)$, then

$$\|\widehat{\varsigma} - \overline{\varsigma}\| \leq \frac{1}{\sqrt{t}} + \sqrt{\frac{\log(1/\delta)}{t}}$$

with probability at least $1-\delta$.

By plugging in $\overline{\varsigma} = q_h, \widehat{\varsigma} = \widehat{q}_h$ with $t = n_q$ and finally $\{\varsigma(i)\}_{i=1}^{n_q} = \{e_{h(x(i))}\}(i)^{n_q}$ and equivalently for p_h we obtain;

$$\|\widehat{b}_h - b_h\|_2 \leq \left(\frac{\sqrt{\log(1/\delta)}}{\sqrt{(1-\beta)n_p}} + \frac{\sqrt{\log(1/\delta)}}{\sqrt{n_q}}\right) + \left(\frac{1}{\sqrt{(1-\beta)n_p}} + \frac{1}{\sqrt{n_q}}\right)$$

with probability at least $1-2\delta$, therefore;

$$\|\widehat{b}_h - b_h\|_2 \leq \frac{2}{\sqrt{\log(2)}}\left(\frac{\sqrt{\log(1/\delta)}}{\sqrt{(1-\beta)n_p}} + \frac{\sqrt{\log(1/\delta)}}{\sqrt{n_q}}\right)$$

resulting in the statement in the Lemma 3.

A.4 Proof of Theorem 1

We want to ultimately bound $|\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}(h^\star)|$. By addition and subtraction we have

$$\begin{aligned}\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}(h^\star) = &\underbrace{\mathcal{L}(\widehat{h}_{\widehat{w}}) - \mathcal{L}_n(\widehat{h}_{\widehat{w}})}_{(b)} + \underbrace{\mathcal{L}_n(\widehat{h}_{\widehat{w}}) - \mathcal{L}_n(\widehat{h}_{\widehat{w}}; \widehat{w})}_{(a)} \\ &+ \underbrace{\mathcal{L}_n(\widehat{h}_{\widehat{w}}; \widehat{w}) - \mathcal{L}_n(h^\star; \widehat{w})}_{\leq 0} + \underbrace{\mathcal{L}_n(h^\star; \widehat{w}) - \mathcal{L}_n(h^\star)}_{(a)} + \underbrace{\mathcal{L}_n(h^\star) - \mathcal{L}(h^\star)}_{(b)} \qquad (10)\end{aligned}$$

where $n = \beta n_p$ and we used optimality of $\widehat{h}_{\widehat{w}}$. Here (a) is the weight estimation error and (b) is the finite sample error.

Uniform law for bounding (b) We bound (b) using standard results for uniform laws for uniformly bounded functions which holds since $\|w\|_\infty \leq d_\infty(q||p)$ and $\|\ell\|_\infty \leq 1$. Since $|w(y)\ell(h(x), y)| \leq d_\infty(q||p),\ \forall x, y \in \mathcal{X} \times \mathcal{Y}$, by deploying the McDiarmid's inequality we then obtain that

$$\sup_{h\in\mathcal{H}} |\mathcal{L}_n(h) - \mathcal{L}(h)| \leq 2\mathcal{R}_n(\mathcal{G}(\ell, \mathcal{H})) + d_\infty(q||p)\sqrt{\frac{\log\frac{2}{\delta}}{n}}$$

where $\mathcal{G}(\ell, \mathcal{H}) = \{g_h(x, y) = w(y)\ell(h(x), y) : h \in \mathcal{H}\}$ and the Rademacher complexity is defined as $\mathcal{R}_n(\mathcal{G}) := \mathbb{E}_{X(i),Y(i)\sim\mathbb{P}:i\in[n]}\left[\mathbb{E}_{\sigma_i:i\in[n]}\frac{1}{n}\left[\sup_{h\in\mathcal{H}}\sum_{i=1}^n \sigma_i w(y_i)\ell(x_i, h(y_i))\right]\right]$.

of the hypothesis class $\mathcal{H}$ (see for example Percy Liang notes on Statistical Learning Theory and chapter 4 in Wainwright (2019))

Bounding term (a) Remember that $k = |\mathcal{Y}|$ is the cardinality of the finite domain of Y, or the number of classes. Let us define $\widetilde{\ell} \in \mathcal{R}^k$ with $\widetilde{\ell}_j = \sum_{i=1}^n \mathbb{I}_{y(i)=j}\ell(y(i), h(x(i)))$. Notice that by definition $\|\widetilde{\ell}\|_1 \leq n$ and $\|\widetilde{\ell}\|_\infty \leq n$ from which it follows by Hoelder's inequality that $\|\widetilde{\ell}\|_2 \leq n$. Furthermore, we slightly abuse notation and use w to denote the k-dimensional vector with $w_i = w(i)$. Therefore, for all h we have via the Cauchy Schwarz inequality that

$$\begin{aligned}|\mathcal{L}_n(h; w) - \mathcal{L}_n(h; \widehat{w})| &= |\frac{1}{n}\sum_{i=1}^n (w(y(i)) - \widehat{w}(y(i)))\ell(h(x(i)), y(i))| \\ &\leq |\frac{1}{n}\sum_{j=1}^k (w(j) - \widehat{w}(j))\widetilde{\ell}(j)| \\ &\leq \frac{1}{n}\|\widehat{w} - w\|_2\|\widetilde{\ell}\|_2 \leq \|\widehat{w} - w\|_2 \\ &\leq \|\lambda\widehat{\theta} - \theta\|_2 \leq (1-\lambda)\|\theta\|_2 + \lambda\|\widehat{\theta} - \theta\|_2 \qquad (11)\end{aligned}$$

It then follows by Lemma 1 that

$$\sup_{h\in\mathcal{H}} |\mathcal{L}_n(h;w) - \mathcal{L}_n(h;\widehat{w})| \leq (1-\lambda)\|\theta\|_2 \\ O\left(\frac{\lambda}{\sigma_{min}}((\|\theta\|_2)\sqrt{\frac{\log(k/\delta)}{(1-\beta)n_p}} + \sqrt{\frac{\log(1/\delta)}{(1-\beta)n_p}}\sqrt{\frac{\log(1/\delta)}{n_q}})\right)$$

A.5 Generalization for finite hypothesis classes

For finite hypothesis classes, one may bound (b) in (10) using Bernstein's inequality.

We bound (b) by first noting that $w(Y)\ell(Y,h(X))$ satisfies the Bernstein conditions so that Bernstein's inequality holds

$$\mathbb{E}_{\mathbb{P}}[w(Y)] = 1, \quad \mathbb{E}_{\mathbb{P}}[w(Y)^2] = d(q||p), \quad \sigma^2_{\mathbb{P}}(w(Y)) = d(q||p) - 1 \tag{12}$$

by definition. Because we assume $\ell \leq 1$, we directly have

$$\mathbb{E}_{\mathbb{P}}\left[w(Y)^2 l^2(Y,h(X))\right] \leq \mathbb{E}_{\mathbb{P}}\left[w(Y)^2\right] = d(q||p) \tag{13}$$

Since we have a bound on the second moment of weighted loss while its first moment is $L(h)$ we can apply Bernstein's inequality to obtain for any fixed h that

$$|\mathcal{L}_n(h) - \mathcal{L}(h)| \leq \frac{2d_\infty(q||p)\log(\frac{2}{\delta})}{3n} + \sqrt{\frac{2\left(d(q||p) - \mathcal{L}(h)^2\right)\log(\frac{2}{\delta})}{n}}$$

For the uniform law for finite hypothesis classes make the union bound on all the hypotheses;

$$\sup_{h\in\mathcal{H}} |\mathcal{L}_n(h) - \mathcal{L}(h)| \leq \frac{2d_\infty(q||p)\log(\frac{2|\mathcal{H}|}{\delta})}{3n} + \sqrt{\frac{2\left(d(q||p)\right)\log(\frac{2|\mathcal{H}|}{\delta})}{n}}$$

The second moment of the importance weighted loss $\mathbb{E}_{\mathbb{P}}\left[\omega(Y)^2\ell^2(h(X),Y)\right]$, given a $h \in \mathcal{H}'$ can be bounded for general $\alpha \geq 0$, potentially leading to smaller values than $d(q||p)$:

$$\begin{aligned}
&\mathbb{E}_{\mathbb{P}}\left[\omega_Y^2\ell^2(Y,h(X))\right] \\
&\quad= \sum_i^k p(i)\frac{q^2(i)}{p^2(i)}\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),X)\right] \\
&\quad= \sum_i^k q(i)^{\frac{1}{\alpha}}\frac{q(i)}{p(i)}q(i)^{\frac{\alpha-1}{\alpha}}\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),i)\right] \\
&\quad\leq \left(\sum_i^k q(i)\frac{q(i)^\alpha}{p(i)^\alpha}\right)^{\frac{1}{\alpha}}\left(\sum_i^k q(i)\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),i)\right]^{\frac{2\alpha}{\alpha-1}}\right)^{\frac{\alpha-1}{\alpha}} \\
&\quad= \left(\sum_i^k q(i)\frac{q(i)^\alpha}{p(i)^\alpha}\right)^{\frac{1}{\alpha}}\left(\sum_i^k q(i)\mathbb{E}_{X\sim p(X|Y=i)}\left[l^2(h(X),i)\right]\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),i)\right]^{\frac{\alpha+1}{\alpha-1}}\right)^{\frac{\alpha-1}{\alpha}} \\
&\quad\leq \left(\sum_i^k q(i)\frac{q(i)^\alpha}{p(i)^\alpha}\right)^{\frac{1}{\alpha}}\sum_i^k q(i)\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),i)\right]^{1-\frac{1}{\alpha}}\mathbb{E}_{X\sim p(X|Y=i)}\left[\ell^2(h(X),i)\right]^{1+\frac{1}{\alpha}}
\end{aligned} \tag{14}$$

where the first inequality follows from Hölder's inequality, the second one follows from Jensen's inequality and the fact that the loss is in $[0,1]$ as well as the fact that the exponentiation function is convex in this region. Moreover, since $1+\frac{1}{\alpha} \geq 1$ and upper bound for the loss square, $l(\cdot,\cdot)^2 \leq 1$, then;

$$\mathbb{E}_{\mathbb{P}}\left[w(Y)^2\ell^2(h(X),Y)\right] \leq \left(\sum_i^k q(i)\frac{q(i)^\alpha}{p(i)^\alpha}\right)^{\frac{1}{\alpha}}\sum_i^k q(i)\mathbb{E}_{X\sim\mathbb{P}|Y=i}\left[\ell^2(h(X),i)\right]^{1-\frac{1}{\alpha}}$$

which gives bound on the second moment of weighted loss.

A.6 MORE EXPERIMENTAL RESULTS

This section contains more experiments that provide more insights about in which settings the advantage of using RLLS vs. BBSL are more or less pronounced.

A.6.1 CIFAR10 EXPERIMENTS UNDER TWEAK-ONE SHIFT AND DIRICHLET SHIFT

Here we compare weight estimation performance between RLLS and BBSL for different types of shifts including the *Tweak-one Shift*, for which we randomly choose one class, e.g. i and set $p(i) = \rho$ while all other classes are distributed evenly. Figure 5 depicts the the weight estimation performance of RLLS compared to BBSL for a variety of values of ρ and α. Note that larger shifts correspond to smaller α and larger ρ. In general, one observes that our RLLS estimator has smaller MSE and that as the shift increases, the error of both methods increases. For tweak-one shift we can additionally see that as the shift increases, RLLS outperforms BBSL more and more as both in terms of bias and variance.

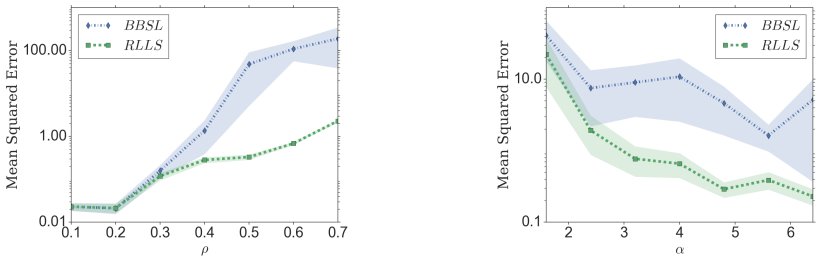

(a) (b)

Figure 5: Comparing MSE of estimated weights using BBSL and RLLS on CIFAR10 with (a) tweak-one shift on source and uniform target, and (b) Dirichlet shift on source and uniform target

A.6.2 MNIST EXPERIMENTS UNDER MINORITY-CLASS SOURCE SHIFTS FOR DIFFERENT VALUES OF p

In order to show weight estimation and classification performance under different level of label shifts, we include several additional sets of experiments here in the appendix. Figure 6 shows the weight estimation error and accuracy comparison under a minority-class shift with p = 0.001. The training and testing sample size is 10000 examples in this case. We can see that whenever the weight estimation of RLLS is better, the accuracy is also better, except in the four classes case when both methods are bad in weight estimation. Figure 7 demonstrates another case in minority-class shift

(a) (b)

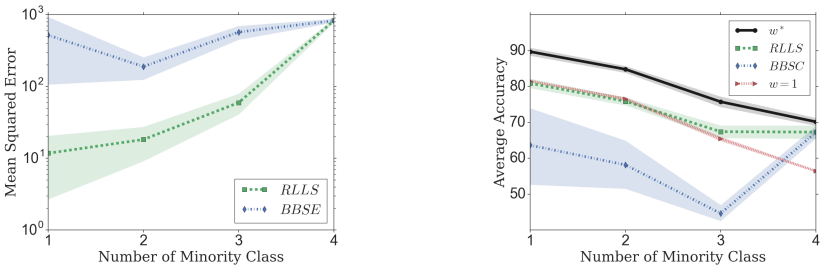

Figure 6: (a) Mean squared error in estimated weights and (b) accuracy on MNIST for minority-class shifted source and uniform target with p = 0.001.

when $p = 0.01$. The black-box classifier is the same two-layers neural network trained on a biased source data set with tweak-one $\rho = 0.5$. We observe that when the number of minority class is small like 1 or 2, the weight estimation is similar between two methods, as well as in the classification accuracy. But when the shift get larger, the weights are worse and the performance in accuracy decreases, getting even worse than the unweighted classifier.

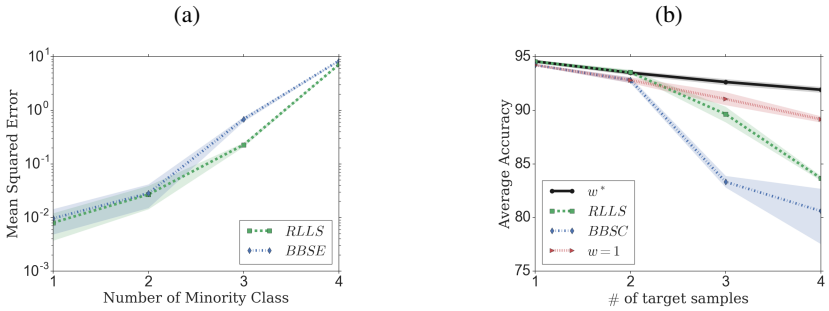

Figure 7: (a) Mean squared error in estimated weights and (b) accuracy on MNIST for minority-class shifted source and uniform target with p = 0.01.

A.6.3 MNIST Experiment under Dirichlet Shift with low target samples

We show the performance of classifier with different regularization λ under a Dirichlet shift with $\alpha = 0.5$ in Figure 8. The training has 5000 examples in this case. We can see that in this low target sample case, $\lambda = 1$ only take over after several hundreds example, while some λ value between 0 and 1 outperforms it at the beginning. Similar as in the paper, we use different black-box classifier that is corrupted in different levels to show the relation between the quality of black-box predictor and the necessary target sample size. We use biased source data with tweak-one $\rho = 0, 0.2, 0.6$ to train the black-box classifier. We see that we need more target samples for the fully weighted version $\lambda = 1$ to take over for a more corrupted black-box classifier.

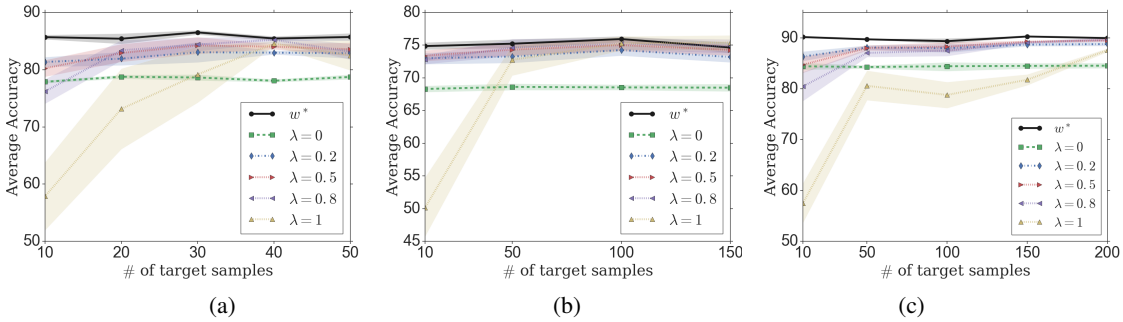

Figure 8: Performance on MNIST for Dirichlet shifted source and uniform target with various target sample size and λ using (a) better predictor, (b) neutral predictor and (c) corrupted predictor.