
A Unified View of Label Shift Estimation

Saurabh Garg, Yifan Wu, Sivaraman Balakrishnan, Zachary C. Lipton
Machine Learning Department,
Department of Statistics and Data Science,
Carnegie Mellon University
{sgarg2, yw4, sbalakri, zlipton}@andrew.cmu.edu

Abstract

Under label shift, the label distribution $p(y)$ might change but the class-conditional distributions $p(x|y)$ do not. There are two dominant approaches for estimating the label marginal. BBSE, a moment-matching approach based on confusion matrices, is provably consistent and provides interpretable error bounds. However, a maximum likelihood estimation approach, which we call MLLS, dominates empirically. In this paper, we present a unified view of the two methods and the first theoretical characterization of MLLS. Our contributions include (i) consistency conditions for MLLS, which include calibration of the classifier and a confusion matrix invertibility condition that BBSE also requires; (ii) a unified framework, casting BBSE as roughly equivalent to MLLS for a particular choice of calibration method; and (iii) a decomposition of MLLS’s finite-sample error into terms reflecting miscalibration and estimation error. Our analysis attributes BBSE’s statistical inefficiency to a loss of information due to coarse calibration. Experiments on synthetic data, MNIST, and CIFAR10 support our findings.

1 Introduction

Supervised algorithms are typically developed and evaluated assuming independent and identically distributed (iid) data. However, the real world seldom abides, presenting domain adaptation problems in which the *source distribution* P_s , from which we sample labeled training examples, differs from the *target distribution* P_t , from which we only observe unlabeled data. Absent assumptions on the nature of shift, the problem is underspecified. Multiple assumptions may be compatible with the same observations while implying different courses of action. Fortunately, some assumptions can render shift detection, estimation, and on-the-fly updates to our classifiers possible.

This paper focuses on *label shift* [22, 18, 16], which aligns with the *anticausal* setting in which the labels y cause the features x [19]. Label shift arises in diagnostic problems because diseases cause symptoms. In this setting, an intervention on $p(y)$ induces the shift, but the process generating x given y is fixed ($p_s(x|y) = p_t(x|y)$). Under label shift, the optimal predictor may change, e.g., the probability that a patient suffers from a disease given their symptoms can increase under a pandemic. Contrast label shift with the better-known *covariate shift* assumption, which aligns with the assumption that x causes y , yielding the reverse implication that $p_s(y|x) = p_t(y|x)$.

Under label shift, our first task is to estimate the ratios $w(y) = p_t(y)/p_s(y)$ for all labels y . Two dominant approaches leverage off-the-shelf classifiers to estimate w : (i) *Black Box Shift Estimation* (BBSE) [16] and a variant called *Regularized Learning under Label Shift* (RLLS) [2]: moment-matching based estimators that leverage (possibly biased, uncalibrated, or inaccurate) predictions to estimate the shift; and (ii) *Maximum Likelihood Label Shift* (MLLS) [18]: an Expectation Maximization (EM) algorithm that assumes access to a classifier that outputs the true source distribution conditional probabilities $p_s(y|x)$.

Given a predictor \hat{f} with an invertible confusion matrix, BBSE and RLLS have known consistency results and finite-sample guarantees [16, 2]. However, MLLS, in combination with a calibration heuristic called Bias-Corrected Temperature Scaling (BCTS), outperforms them empirically [1].

In this paper, we theoretically characterize MLLS, establishing conditions for consistency and bounding its finite-sample error. To start, we observe that given the true label conditional $p_s(y|x)$, MLLS is simply a concave Maximum Likelihood Estimation (MLE) problem and standard results apply. However, because we never know $p_s(y|x)$ exactly, MLLS is always applied with an estimated model \hat{f} and thus the procedure consists of MLE under model misspecification.

First, we prove that (i) *canonical calibration* (Definition 1) and (ii) an invertible confusion matrix (as required by BBSE) are *sufficient conditions* to ensure MLLS’s consistency (Proposition 1, Theorems 1 and 2). We also show that calibration can sometimes be *necessary* for consistency (Example 1 in Section 4.3). Recall that neural network classifiers tend to be uncalibrated absent post-hoc adjustments [10]. Second, we observe that confusion matrices can be instruments for calibrating a classifier. Applying MLLS with this technique, BBSE and MLLS are distinguished only by their objective functions. Through extensive experiments, we show that they perform similarly, concluding that MLLS’s superior performance (when applied with more granular calibration techniques) is not due to its objective but rather to the information lost by BBSE via confusion matrix calibration. Third, we analyze the finite-sample error of the MLLS estimator by decomposing its error into terms reflecting the miscalibration error and finite-sample error (Theorem 3). Depending on the calibration method, the miscalibration error can further be divided into two terms: finite sample error due to re-calibration on a validation set and the minimum achievable calibration error with that technique.

We validate our results on synthetic data, MNIST, and CIFAR-10. Empirical results show that MLLS can have 2–10 \times lower Mean Squared estimation Error (MSE) depending on the magnitude of the shift. Our experiments relate MLLS’s MSE to the granularity of the calibration.

In summary, we contribute the following: (i) Sufficient conditions for MLLS’s consistency; (ii) Unification of MLLS and BBSE methods under a common framework, with BBSE corresponding to a particular choice of calibration method; (iii) Finite-sample error bounds for MLLS; (iv) Experiments on synthetic and image recognition datasets that support our theoretical arguments.

2 Problem Setup

Let \mathcal{X} be the input space and $\mathcal{Y} = \{1, 2, \dots, k\}$ the output space. Let $P_s, P_t : \mathcal{X} \times \mathcal{Y} \rightarrow [0, 1]$ be the source and target distributions and let p_s and p_t denote the corresponding probability density (or mass) functions. We use \mathbb{E}_s and \mathbb{E}_t to denote expectations over the source and target distributions. In unsupervised domain adaptation, we possess labeled source data $\{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ and unlabeled target data $\{x_{n+1}, x_{n+2}, \dots, x_{n+m}\}$. We also assume access to a black-box predictor $\hat{f} : \mathcal{X} \mapsto \Delta^{k-1}$, e.g., a model trained to approximate the true probability function f^* , where $f^*(x) := p_s(\cdot|x)$. Here and in the rest of the paper, we use Δ^{k-1} to denote the standard k -dimensional probability simplex. For a vector v , we use v_y to access the element at index y .

Absent assumptions relating the source and target distributions, domain adaptation is underspecified [3]. We work with the *label shift* assumption, i.e., $p_s(x|y) = p_t(x|y)$, focusing on multiclass classification. Moreover, we assume non-zero support for all labels in the source distribution: for all $y \in \mathcal{Y}$, $p_s(y) \geq c > 0$ [16, 2]. Under label shift, three common goals are (i) detection—determining whether distribution shift has occurred; (ii) quantification—estimating the target label distribution; and (iii) correction—producing a predictor that minimizes error on the target distribution [16].

This paper focuses on goal (ii), estimating importance weights $w(y) = p_t(y)/p_s(y)$ for all $y \in \mathcal{Y}$. Given w , we can update our classifiers on the fly, either by retraining in an importance-weighted ERM framework [20, 9, 16, 2]—a practice that may be problematic for overparameterized neural networks [5], or by applying an analytic correction [1, 18]. Within the ERM framework, the generalization result from Azizzadenesheli et al. [2] (Theorem 1) depends only on the error of the estimated weights, and hence any method that improves weight estimates tightens this bound.

There are multiple definitions of calibration in the multiclass setting. Guo et al. [10] study the calibration of the arg-max prediction, while Kumar et al. [14] study a notion of per-label calibration. We use canonical calibration [24] and the expected canonical calibration error on the source data defined as follows:

Definition 1 (Canonical calibration). *A prediction model $f : \mathcal{X} \mapsto \Delta^{k-1}$ is canonically calibrated on the source domain if for all $x \in \mathcal{X}$ and $j \in \mathcal{Y}$, $P_s(y = j | f(x)) = f_j(x)$.*

Definition 2 (Expected canonical calibration error). *For a predictor f , the expected squared canonical calibration error on the source domain is $\mathcal{E}^2(f) = \mathbb{E}_s \|f - f_c\|^2$, where $f_c = P_s(y = \cdot | f(x))$.*

Calibration methods typically work either by calibrating the model during training or by calibrating a trained classifier on held-out data, post-hoc. We refer the interested reader to Kumar et al. [14] and Guo et al. [10] for detailed studies on calibration. We focus on the latter category of methods. Our experiments follow Alexandari et al. [1], who leverage BCTS¹ to calibrate their models. BCTS extends temperature scaling [10] by incorporating per-class bias terms.

3 Prior Work

Two families of solutions have been explored that leverage a blackbox predictor: BBSE [16], a moment matching method, uses the predictor \hat{f} to compute a confusion matrix $C_{\hat{f}} := p_s(\hat{y}, y) \in \mathbb{R}^{k \times k}$ on the source data. Depending on how \hat{y} is defined, there are two types of confusion matrix for a predictor \hat{f} : (i) the *hard confusion matrix* $\hat{y} = \arg \max_{\mathcal{Y}} \hat{f}(x)$; and (ii) the *soft confusion matrix*, where \hat{y} is defined as a random prediction that follows the discrete distribution $\hat{f}(x)$ over \mathcal{Y} . Both soft and hard confusion matrix can be estimated from labeled source data samples. The estimate \hat{w} is computed as $\hat{w} := \hat{C}_{\hat{f}}^{-1} \hat{\mu}$, where $\hat{C}_{\hat{f}}$ is the estimate of confusion matrix and $\hat{\mu}$ is an estimate of $p_t(\hat{y})$, computed by applying the predictor \hat{f} to the target data. In a related vein, RLLS [2] incorporates an additional regularization term of the form $\|w - 1\|$ and solves a constrained optimization problem to estimate the shift ratios w .

MLLS estimates w as if performing maximum likelihood estimation, but substitutes the predictor outputs for the true probabilities $p_s(y|x)$. Saerens et al. [18], who introduce this procedure, describe it as an application of EM. However, as observed in [8, 1], the likelihood objective is concave, and thus a variety of optimization algorithms may be applied to recover the MLLS estimate. Alexandari et al. [1] also showed that MLLS underperforms BBSE when applied naively, a phenomenon that we shed more light on in this paper.

4 A Unified View of Label Shift Estimation with Black Box Predictors

We now present a unified view that subsumes MLLS and BBSE and demonstrate how each is instantiated under this framework. We also establish identifiability and consistency conditions for MLLS, deferring a treatment of finite-sample issues to Section 5. For convenience, throughout Sections 3 and 4, we use the term *calibration* exclusively to refer to canonical calibration (Definition 1) on the source data. We relegate all technical proofs to Appendix D.

4.1 A Unified Distribution Matching View

To start, we introduce a *generalized* distribution matching approach for estimating w . Under label shift, for any (possibly randomized) mapping from \mathcal{X} to \mathcal{Z} , we have that $p_s(z|y) = p_t(z|y)$ since, $p_s(z|y) = p_t(z|y) = \int_{\mathcal{X}} p(z|x)p(x|y)dx$. Throughout the paper, we use the notation $p(z|y)$ to represent either $p_s(z|y)$ or $p_t(z|y)$ (which are identical). We now define a family of distributions over \mathcal{Z} parameterized by $w \in \mathcal{W}$ as

$$p_w(z) = \sum_{y=1}^k p(z|y)p_s(y)w_y = \sum_{y=1}^k p_s(z, y)w_y, \quad (1)$$

where $\mathcal{W} = \{w \mid \forall y, w_y \geq 0 \text{ and } \sum_{y=1}^k w_y p_s(y) = 1\}$. When $w = w^*$, we have that $p_w(z) = p_t(z)$. For fixed $p(z|x)$, $p_t(z)$ and $p_s(z, y)$ are known because $p_t(x)$ and $p_s(x, y)$ are known. So one potential strategy to estimate w^* is to find a weight vector w such that

$$\sum_{y=1}^k p_s(z, y)w_y = p_t(z) \quad \forall z \in \mathcal{Z}. \quad (2)$$

At least one such weight vector w must exist as w^* satisfies (2). We now characterize conditions under which the weight vector w satisfying (2) is unique:

¹Motivated by the strong empirical results in Alexandari et al. [1], we use BCTS in our experiments as a surrogate to canonical calibration.

Lemma 1 (Identifiability). *If the set of distributions $\{p(z|y) : y = 1, \dots, k\}$ are linearly independent, then for any w that satisfies (2), we must have $w = w^*$. This condition is also necessary in general: if the linear independence does not hold then there exists a problem instance where we have $w, w^* \in \mathcal{W}$ satisfying (2) while $w \neq w^*$.*

Lemma 1 follows from the fact that (2) is a linear system with at least one solution w^* . This solution is unique when $p_s(z, y)$ is of rank k . The linear independence condition in Lemma 1, in general, is sufficient for identifiability of discrete \mathcal{Z} . However, for continuous \mathcal{Z} , the linear dependence condition has the undesirable property of being sensitive to changes on sets of measure zero. By changing a collection of linearly dependent distributions on a set of measure zero, we can make them linearly independent. As a consequence, we impose a *stronger* notion of identifiability i.e., the set of distributions $\{p(z|y) : y = 1, \dots, k\}$ are such that there does not exist $v \neq 0$ for which $\int_{\mathcal{Z}} |\sum_y p(z|y)v_y| dz = 0$. We refer this condition as *strict linear independence*.

In generalized distribution matching, one can set $p(z|x)$ to be the Dirac delta function at δ_x^2 such that \mathcal{Z} is the same space as \mathcal{X} , which leads to solving (2) with z replaced by x . In practice where \mathcal{X} is high-dimensional and/or continuous, approximating the solution to (2) from finite samples can be hard when choosing $z = x$. Our motivation for generalizing distribution matching from \mathcal{X} to \mathcal{Z} is that the solution to (2) can be better approximated using finite samples when \mathcal{Z} is chosen carefully. Under this framework, the design of a label shift estimation algorithm can be decomposed into two parts: (i) the choice of $p(z|x)$ and (ii) how to approximate the solution to (2). Later on, we consider how these design choices may affect label shift estimation procedures in practice.

4.2 The Confusion Matrix Approach

If \mathcal{Z} is a discrete space, one can first estimate $p_s(z, y) \in \mathbb{R}^{|\mathcal{Z}| \times k}$ and $p_t(z) \in \mathbb{R}$, and then subsequently attempt to solve (2). Confusion matrix approaches use $\mathcal{Z} = \mathcal{Y}$, and construct $p(z|x)$ using a black box predictor \hat{f} . There are two common choices to construct the confusion matrix: (i) The soft confusion matrix approach: We set $p(z|x) := \hat{f}(x) \in \Delta^{k-1}$. We then define a random variable $\hat{y} \sim \hat{f}(x)$ for each x . Then we construct $p_s(z, y) = p_s(\hat{y}, y)$ and $p_t(z) = p_t(\hat{y})$. (ii) The hard confusion matrix approach: Here we set $p(z|x) = \delta_{\arg \max \hat{f}(x)}$. We then define a random variable $\hat{y} = \arg \max \hat{f}(x)$ for each x . Then again we have $p_s(z, y) = p_s(\hat{y}, y)$ and $p_t(z) = p_t(\hat{y})$.

Since $p_s(z, y)$ is a square matrix, the identifiability condition becomes the invertibility of the confusion matrix. Given an estimated confusion matrix, one can find w by inverting the confusion matrix (BBSE) or minimizing some distance between the vectors on the two sides of (2).

4.3 Maximum Likelihood Label Shift Estimation

When \mathcal{Z} is a continuous space, the set of equations in (2) indexed by \mathcal{Z} is intractable. In this case, one possibility is to find a weight vector \tilde{w} by minimizing the KL-divergence $\text{KL}(p_t(z), p_w(z)) = \mathbb{E}_t [\log p_t(z)/p_w(z)]$, for p_w defined in (1). This is equivalent to maximizing the population log-likelihood: $\tilde{w} := \arg \max_{w \in \mathcal{W}} \mathbb{E}_t [\log p_w(z)]$. One can further show that $\mathbb{E}_t [\log p_w(z)] = \mathbb{E}_t [\log \sum_{y=1}^k p_s(z, y)w_y] = \mathbb{E}_t [\log \sum_{y=1}^k p_s(y|z)p_s(z)w_y] = \mathbb{E}_t [\log \sum_{y=1}^k p_s(y|z)w_y] + \mathbb{E}_t [\log p_s(z)]$. Therefore we can equivalently define:

$$\tilde{w} := \arg \max_{w \in \mathcal{W}} \mathbb{E}_t \left[\log \sum_{y=1}^k p_s(y|z)w_y \right]. \quad (3)$$

This yields a straightforward convex optimization problem whose objective is bounded from below [1, 8]. Assuming access to labeled source data and unlabeled target data, one can maximize the empirical counterpart of the objective in (3), using either EM or an alternative iterative optimization scheme. Saerens et al. [18] derived an EM algorithm to maximize the objective (3) when $z = x$, assuming access to $p_s(y|x)$. Absent knowledge of the ground truth $p_s(y|x)$, we can plug in any approximate predictor f and optimize the following objective:

$$w_f := \arg \max_{w \in \mathcal{W}} \mathcal{L}(w, f) := \arg \max_{w \in \mathcal{W}} \mathbb{E}_t [\log f(x)^T w]. \quad (4)$$

In practice, f is fit from a finite sample drawn from $p_s(x, y)$ and standard machine learning methods often produce uncalibrated predictors. While BBSE and RLLS are provably consistent whenever the

²For simplicity we will use $z = x$ to denote that $p(z|x) = \delta_x$.

predictor f yields an invertible confusion matrix, to our knowledge, no prior works have established sufficient conditions to guarantee MLLS' consistency when f differs from $p_s(y|x)$.

It is intuitive that for some values of $f \neq p_s(y|x)$, MLLS will yield inconsistent estimates. Supplying empirical evidence, Alexandari et al. [1] show that MLLS performs poorly when f is a vanilla neural network predictor learned from data. However, Alexandari et al. [1] also show that in combination with a particular post-hoc calibration technique, MLLS achieves low error, significantly outperforming BBSE and RLLS. As the calibration error is not a distance metric between f and $p_s(y|x)$ (zero calibration error does not indicate $f = p_s(y|x)$), a calibrated predictor f may still be substantially different from $p_s(y|x)$. Some natural questions then arise:

1. *Why does calibration improve MLLS so dramatically?*
2. *Is calibration necessary or sufficient to ensure the consistency of MLLS?*
3. *What accounts for the comparative efficiency of MLLS over BBSE?* (Addressed in Section 5)

To address the first two questions, we make the following observations. Suppose we define z (for each x) with distribution $p(z|x) := \delta_{f(x)}$, for some calibrated predictor f . Then, because f is calibrated, it holds that $p_s(y|z) = f(x)$. Note that in general, the MLLS objective (4) can differ from (3). However, when $p(z|x) := \delta_{f(x)}$, the two objectives are identical. We can formalize this as follows:

Lemma 2. *If f is calibrated, then the two objectives (3) and (4) are identical when \mathcal{Z} is chosen as Δ^{k-1} and $p(z|x)$ is defined to be $\delta_{f(x)}$.*

Lemma 2 follows from changing the variable of expectation in (4) from x to $f(x)$ and applying $f(x) = p_s(y|f(x))$ (definition of calibration). It shows that MLLS with a calibrated predictor on the input space \mathcal{X} is in fact equivalent to performing distribution matching in the space \mathcal{Z} . Building on this observation, we now state our population-level consistency theorem for MLLS:

Theorem 1 (Population consistency of MLLS). *If a predictor $f : \mathcal{X} \mapsto \Delta^{k-1}$ is calibrated and the distributions $\{p(f(x)|y) : y = 1, \dots, k\}$ are strictly linearly independent, then w^* is the unique maximizer of the MLLS objective (4).*

We now turn our attention to establishing consistency of the sample-based estimator. Let $x_1, x_2, \dots, x_m \stackrel{iid}{\sim} p_t(x)$. The finite sample objective for MLLS can be written as

$$\hat{w}_f := \arg \max_{w \in \mathcal{W}} \frac{1}{m} \sum_{i=1}^m \log f(x_i)^T w := \arg \max_{w \in \mathcal{W}} \mathcal{L}_m(w, f). \quad (5)$$

Theorem 2 (Consistency of MLLS). *If f satisfies the conditions in Theorem 1, then \hat{w}_f in (5) converges to w^* almost surely.*

The main idea of the proof of Theorem 2 is to derive a metric entropy bound on the class of functions $\mathcal{G} = \{(f^T w)/(f^T w + f^T w^*) | w \in \mathcal{W}\}$ to prove Hellinger consistency (Theorem 4.6 [25]). The consistency of MLLS relies on the linear independence of the collection of distributions $\{p(f(x)|y) : y = 1, \dots, k\}$. The following result develops several alternative equivalent characterizations of this linear independence condition.

Proposition 1. *For a calibrated predictor f , the following statements are equivalent:*

- (1) $\{p(f(x)|y) : y = 1, \dots, k\}$ are strictly linearly independent.
- (2) $\mathbb{E}_s [f(x)f(x)^T]$ is invertible.
- (3) The soft confusion matrix of f is invertible.

Proposition 1 shows that with a calibrated predictor, the invertibility condition as required by BBSE (or RLLS) is exactly the same as the linear independence condition required for MLLS's consistency.

Having provided sufficient conditions, we consider a binary classification example to provide intuition for why we need calibration for consistency. In this example, we relate the estimation error to the miscalibration error, showing that calibration is not only sufficient but also necessary to achieve zero estimation error for a certain class of predictors.

Example 1. Consider a mixture of two Gaussians with $p_s(x|y = 0) := \mathcal{N}(\mu, 1)$ and $p_s(x|y = 1) := \mathcal{N}(-\mu, 1)$. We suppose that the source mixing coefficients are both $\frac{1}{2}$, while the target mixing

coefficients are $\alpha(\neq \frac{1}{2}), 1 - \alpha$. Assume a class of probabilistic threshold classifiers: $f(x) = [1 - c, c]$ for $x \geq 0$, otherwise $f(x) = [c, 1 - c]$ with $c \in [0, 1]$. Then the population error of MLLS is given by

$$4 \left| \frac{(1 - 2\alpha)(p_s(x \geq 0|y = 0) - c)}{1 - 2c} \right|,$$

which is zero only if $c = p_s(x \geq 0|y = 0)$ for a non-degenerate classifier.

The expression for estimation error arising from our example yields two key insights: (i) an uncalibrated thresholded classifier has an estimation error proportional to the true shift in label distribution i.e. $1 - 2\alpha$; (ii) the error is also proportional to the canonical calibration error which is $p_s(x \geq 0|y = 0) - c$. While earlier in this section, we concluded that canonical calibration is sufficient for consistency, the above example provides some intuition for why it might also be necessary. In Appendix C, we show that marginal calibration [14, 10, 24], a less restricted definition is insufficient to achieve consistency.

4.4 MLLS with Confusion Matrix

So far, we have shown that MLLS with any calibrated predictor can be viewed as distribution matching in a latent space. Now we discuss a method to construct a predictor f to perform MLLS given any $p(z|x)$, e.g., those induced by confusion matrix approaches. Recall, we already have the maximum log-likelihood objective. It just remains to construct a calibrated predictor f from the confusion matrix.

This is straightforward when $p(z|x)$ is deterministic, i.e., $p(z|x) = \delta_{g(x)}$ for some function g : setting $f(x) = p_s(y|g(x))$ makes the objectives (3) and (4) to be the same. Recall that for the hard confusion matrix, the induced latent space is $p(z|x) = \delta_{\arg \max \hat{f}(x)}$. So the corresponding predictor in MLLS is $f(x) = p_s(y|\hat{y}_x)$, where $\hat{y}_x = \arg \max \hat{f}(x)$. Then we obtain the MLLS objective for the hard confusion matrix:

$$\max_{w \in \mathcal{W}} \mathbb{E}_t \left[\log \sum_{y=1}^k p_s(y|\hat{y}_x) w_y \right]. \quad (6)$$

The confusion matrix $C_{\hat{f}}$ and predictor \hat{f} directly give us $p_s(y|\hat{y}_x)$. Given an input x , one can first get \hat{y}_x from \hat{f} , then normalize the \hat{y}_x -th row of $C_{\hat{f}}$ as $p_s(y|\hat{y}_x)$. We denote MLLS with hard confusion matrix calibration (6) by MLLS-CM.

When $p_s(z|x)$ is stochastic, we need to extend (4) to allow f to be a random predictor: $f(x) = p_s(y|z)$ for $z \sim p(z|x)$ ³. To incorporate the randomness of f , one only needs to change the expectation in (4) to be over both x and $f(x)$, then (4) becomes a rewrite of (3).

Proposition 2 indicates that constructing the confusion matrix is a calibration procedure. Thus, the predictor constructed with constructed using confusion matrix is calibrated and suitable for application with MLLS.

Proposition 2 (Vaicenavicius et al. [24]). *For any function g , $f(x) = p_s(y|g(x))$ is a calibrated predictor.*

We can now summarize the relationship between BBSE and MLLS: A label shift estimator involves two design choices: (i) designing the latent space $p(z|x)$ (which is equivalent to designing a calibrated predictor); and (ii) performing distribution matching in the new space \mathcal{Z} . In BBSE, we design a calibrated predictor via the confusion matrix and then perform distribution matching by directly solving linear equations. In general, MLLS does not specify how to obtain a calibrated predictor, but specifies KL minimization as the distribution matching procedure. One can apply the confusion matrix approach to obtain a calibrated predictor and then plug it into MLLS, which is the BBSE analog under MLLS, and is a special case of MLLS.

5 Theoretical Analysis of MLLS

We now analyze the performance of MLLS estimator. Even when w^* is the unique optimizer of (4) for some calibrated predictor f , assuming convex optimization can be done perfectly, there are still two sources of error preventing us from exactly computing w^* in practice. First, we are optimizing a sample-based approximation (5) to the objective in expectation (4). We call this source of error

³Here, by a random predictor we mean that the predictor outputs a random vector from Δ^{k-1} , not \mathcal{Y} .

finite-sample error. Second, the predictor f we use may not be perfectly calibrated on the source distribution as we only have access to samples from source data distribution $p_s(x, y)$. We call this source of error *miscalibration error*. We will first analyze how these two sources of errors affect the estimate of w^* separately and then give a general error bound that incorporates both. All proofs are relegated to Appendix E.

Before presenting our analysis, we introduce some notation and regularity assumptions. For any predictor $f : \mathcal{X} \mapsto \Delta^{k-1}$, we define w_f and \widehat{w}_f as in (4) and (5). If f satisfies the conditions in Theorem 2 (calibration and linear independence) then we have that $w_f = w^*$. Our goal is to bound $\|\widehat{w}_f - w^*\|$ for a given (possibly miscalibrated) predictor f . We now introduce a regularity condition:

Condition 1 (Regularity condition for a predictor f). *For any x within the support of $p_t(x)$, i.e. $p_t(x) > 0$, we have both $f(x)^T w_f \geq \tau$, $f(x)^T w^* \geq \tau$ for some universal constant $\tau > 0$.*

Condition 1 is mild if f is calibrated since in this case $w_f = w^*$ is the maximizer of $\mathbb{E}_t [\log f(x)^T w]$, and the condition is satisfied if the expectation is finite. Since $f(x)^T w^*$ and $f(x)^T w_f$ are upper-bounded (they are the inner products of two vectors which sum to 1), they also must be lower-bounded away from 0 with arbitrarily high probability without any assumptions. For miscalibrated f , a similar justification holds for assumption that $f(x)^T w_f$ is lower bounded. Turning our attention to the assumption that $f(x)^T w^*$ is lower bounded, we note that it is sufficient if f is close (pointwise) to some calibrated predictor. This in turn is a reasonable assumption on the actual predictor we use for MLLS in practice as it is post-hoc calibrated on source data samples.

Define $\sigma_{f,w}$ to be the minimum eigenvalue of the Hessian $-\nabla_w^2 \mathcal{L}(w, f)$. To state our results compactly we use standard stochastic order notation (see, for instance, [26]). We first bound the estimation error introduced by only having finite samples from the target distribution in Lemma 3. Next, we bound the estimation error introduced by having a miscalibrated f in Lemma 4.

Lemma 3. *For any predictor f that satisfies Condition 1, we have $\|w_f - \widehat{w}_f\| \leq \sigma_{f,w_f}^{-1} \mathcal{O}_p(m^{-1/2})$.*

Lemma 4. *For any predictor f and any calibrated predictor f_c that satisfies Condition 1, we have $\|w_f - w^*\| \leq \sigma_{f,w^*}^{-1} \cdot C \cdot \mathbb{E}_t [\|f - f_c\|]$, for some constant C .*

If we set $f_c(x) = p_s(y|f(x))$, which is a calibrated predictor (Proposition 2), we can bound the error in terms of the calibration error of f on the source data⁴: $\|w_f - w^*\| \leq \sigma_{f,w^*}^{-1} \cdot C \cdot \mathcal{E}(f)$.

Note that since $p_s(y) > 0$ for all y , we can upper-bound the error in Lemma 4 with calibration error on the source data. We combine the two sources of error to bound the estimation error $\|\widehat{w}_f - w^*\|$:

Theorem 3. *For any predictor f that satisfies Condition 1, we have*

$$\|\widehat{w}_f - w^*\| \leq \sigma_{f,w_f}^{-1} \mathcal{O}_p(m^{-1/2}) + C \cdot \sigma_{f,w^*}^{-1} \mathcal{E}(f). \quad (7)$$

The estimation error of MLLS can be decomposed into (i) finite-sample error, which decays at a rate of $m^{-1/2}$; and (ii) the calibration error of the predictor that we use. The proof is a direct combination of Lemma 3 and Lemma 4 applied to the same f with the following error decomposition:

$$\|\widehat{w}_f - w^*\| \leq \underbrace{\|w_f - \widehat{w}_f\|}_{\text{finite-sample}} + \underbrace{\|w_f - w^*\|}_{\text{miscalibration}}.$$

Theorem 3 shows that the estimation error depends inversely on the minimum eigenvalue of the Hessian at two different points w_f and w^* . One can unify these two eigenvalues as a single quantity σ_f , the minimum eigenvalue $\mathbb{E}_t [f(x)f(x)^T]$. We formalize this observation in Appendix E.

If we use the *post-hoc calibration* procedure (as discussed in Section 2 and A) to calibrate a blackbox predictor \widehat{f} , we can obtain a bound on the calibration error of f . In more detail, suppose that the class \mathcal{G} used for calibration satisfies standard regularity conditions (injectivity, Lipschitz-continuity, twice differentiability, non-singular Hessian). We have the following lemma:

Lemma 5. *Let $f = g \circ \widehat{f}$ be the predictor after post-hoc calibration with squared loss l and g belongs to a function class \mathcal{G} that satisfies the standard regularity conditions, we have*

$$\mathcal{E}(f) \leq \min_{g \in \mathcal{G}} \mathcal{E}(g \circ \widehat{f}) + \mathcal{O}_p(n^{-1/2}). \quad (8)$$

⁴We present two upper bounds because the second is more interpretable while the first is tighter.

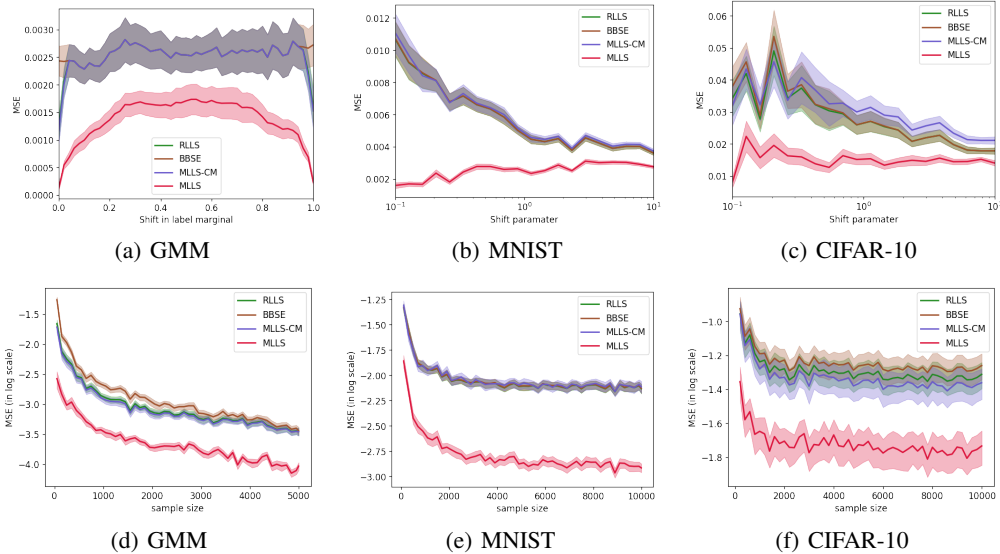


Figure 1: **(top)** MSE vs the degree of shift; For GMM, we control the shift in the label marginal for class 1 with a fixed target sample size of 1000. For multiclass problems—MNIST and CIFAR-10, we control the Dirichlet shift parameter with a fixed sample size of 5000. **(bottom)** MSE (in log scale) vs target sample size; For GMM, we fix the label marginal for class 1 at 0.01 whereas for multiclass problems, MNIST and CIFAR-10, we fix the Dirichlet parameter to 0.1. In all plots, MLLS dominates other methods. All confusion matrix approaches perform similarly, indicating that the advantage of MLLS comes from the choice of calibration but not the way of performing distribution matching.

This result is similar to Theorem 4.1 [14]. For a model class \mathcal{G} that is rich enough to contain a function $g \in \mathcal{G}$ that achieves zero calibration error, i.e., $\min_{g \in \mathcal{G}} \mathcal{E}(g \circ \hat{f}) = 0$, then we obtain an estimation error bound for MLLS of $\sigma_f^{-1} \cdot \mathcal{O}_p(m^{-1/2} + n^{-1/2})$. This bound is similar to rate of RLLS and BBSE, where instead of σ_f they have minimum eigenvalue of the confusion matrix.

The estimation error bound explains the efficiency of MLLS. Informally, the error of MLLS depends inversely on the minimum eigenvalue of the Hessian of the likelihood σ_f . When we apply coarse calibration via the confusion matrix (in MLLS-CM), we only decrease the value of σ_f . Coarse calibration throws away information [13] and thus results in greater estimation error for MLLS. In Section 6, we empirically show that MLLS-CM’s performance is similar to that of BBSE. Moreover, on a synthetic Gaussian mixture model, we show that the minimum eigenvalue of the Hessian obtained using confusion matrix calibration is smaller than the minimum eigenvalue obtained with more granular calibration. Our analysis and observations together suggest MLLS’s superior performance than BBSE (or RLLS) is due to the granular calibration but not due to the difference in the optimization objective.

Finally, we want to highlight one minor point regarding applicability of our result. If f is calibrated, Theorem 3, together with Proposition 3 (in Appendix E), implies that MLLS is consistent if $\mathbb{E}_t[f(x)f(x)^T]$ is invertible. Compared to the consistency condition in Theorem 1 that $\mathbb{E}_s[f(x)f(x)^T]$ is invertible (together with Proposition 1), these two conditions are the same if the likelihood ratio $p_t(f(x))/p_s(f(x))$ is lower-bounded. This is true if all entries in w^* are non-zero. Even if w^* contains non-zero entries, the two conditions are still the same if there exists some $w_y^* > 0$ such that $p(f(x)|y)$ covers the full support of $p_s(f(x))$. In general however, the invertibility of $\mathbb{E}_t[f(x)f(x)^T]$ is a stronger requirement than the invertibility of $\mathbb{E}_s[f(x)f(x)^T]$. We leave further investigation of this gap for future work.

6 Experiments

We experimentally illustrate the performance of MLLS on synthetic data, MNIST [15], and CIFAR10 [12]. Following Lipton et al. [16], we experiment with *Dirichlet shift* simulations. On each run, we sample a target label distribution $p_t(y)$ from a Dirichlet with concentration parameter α . We

then generate each target example by first sampling a label $y \sim p_t(y)$ and then sampling (with replacement) an example conditioned on that label. Note that smaller values of alpha correspond to more severe shift. In our experiments, the source label distribution is uniform.

First, we consider a mixture of two Gaussians (as in Example in Section 4.3) with $\mu = 1$. With CIFAR10 and MNIST, we split the full training set into two subsets: train and valid, and use the provided test set as is. Then according to the label distribution, we randomly sample with replacement train, valid, and test set from each of their respective pool to form the source and target set. To learn the black box predictor on real datasets, we use the same architecture as Lipton et al. [16] for MNIST, and for CIFAR10 we use ResNet-18 [11] as in Azzadenesheli et al. [2]⁵. For simulated data, we use the true $p_s(y|x)$ as our predictor function. For each experiment, we sample 100 datasets for each shift parameter and evaluate the empirical MSE and variance of the estimated weights.

We consider three sets of experiments: (1) MSE vs degree of target shift; (2) MSE vs target sample sizes; and (3) MSE vs calibrated predictors on the source distribution. We refer to MLLS-CM as MLLS with hard confusion matrix calibration as in (6). In our experiments, we compare MLLS estimator with BBSE, RLLS, and MLLS-CM. For RLLS and BBSE, we use the publicly available code⁶. To post-hoc calibration, we use BCTS [1] on the held-out validation set. Using the same validation set, we calculate the confusion matrix for BBSE, RLLS, and MLLS-CM.

We examine the performance of various estimators across all three datasets for various target dataset sizes and shift magnitudes (Figure 1). Across all shifts, MLLS (with BCTS-calibrated classifiers) *uniformly dominates* BBSE, RLLS, and MLLS-CM in terms of MSE (Figure 1). Observe for severe shifts, MLLS is comparatively dominant. As the available target data increased, all methods improve rapidly, with MLLS outperforming all other methods by a significant margin. Moreover, MLLS’s advantages grow more pronounced under extreme shifts. Notice MLLS-CM is roughly equivalent to BBSE across all settings of dataset, target size, and shift magnitude. This concludes MLLS’s superior performance is not because of differences in loss function used for distribution matching but due to differences in the granularity of the predictions, caused by crude confusion matrix aggregation.

Note that given a predictor f_1 , we can partition our input space and produce another predictor f_2 that, for any datapoint gives the expected output of f_1 on points belonging to that partition. If f_1 is calibrated, then f_2 will also be calibrated [24]. On synthetic data, we vary the granularity of calibration (for MLLS) by aggregating $p_s(y|x)$ over a variable number of equal-sized bins. With more bins, less information is lost due to calibration. Consequently, the minimum eigenvalue of the Hessian increases and the MSE decreases, supporting our theoretical bounds (Figure 2). We also verify that the confusion matrix calibration performs poorly (Figure 2). For MLLS-CM, the minimum eigenvalue of the Hessian is 0.195, significantly smaller than for the binned predictor for $\#bin \geq 4$. Thus, the poor performance of MLLS-CM is predicted by its looser upper bound per our analysis. Note that these experiments presume access to the true predictor $p_s(y|x)$ and thus the MSE strictly improves with the number of bins. In practice, with a fixed source dataset size, increasing the number of bins could lead to overfitting, worsening our calibration.

7 Conclusion

This paper provides a unified framework relating techniques that use off-the-shelf predictors for label shift estimation. We argue that these methods all employ calibration, either explicitly or implicitly, differing only in the choice of calibration method and their optimization objective. Moreover, with our analysis we show that the choice of calibration method (and not the optimization objective for distribution matching) accounts for the advantage of MLLS with BCTS calibration over BBSE. In future work, we hope to operationalize these insights to provide guidance for a calibration scheme to improve label shift estimation.

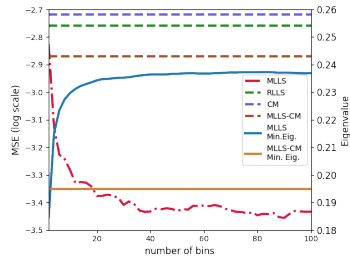


Figure 2: MSE (left-axis) with variation of minimum eigenvalue of the Hessian (right-axis) vs number of bins used for aggregation. With increase in number of bins, MSE decrease and the minimum eigenvalue increases.

⁵We used open source implementation of ResNet-18 <https://github.com/kuangliu/pytorch-cifar>.

⁶BBSE: https://github.com/zackchase/label_shift, RLLS: <https://github.com/Angela0428/labelshift>

Broader Impact

This paper investigates the (statistical) consistency and efficiency of two existing methods for estimating target domain label distributions. While this could potentially guide practitioners to improve detection, estimation, and classification in applications where the label shift assumption holds, we do not believe that it will fundamentally impact how machine learning is used in a way that could conceivably be socially salient. While we take the potential impact of machine learning on society seriously, we believe that this work, which addresses a foundational theoretical problem, does not present a significant societal concern.

Acknowledgments

We thank Zico Kolter and David Childers for their helpful feedback. This material is based on research sponsored by Air Force Research Laboratory (AFRL) under agreement number FA8750-19-1-1000. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation therein. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Air Force Laboratory, DARPA or the U.S. Government. SB acknowledges funding from the NSF grants DMS-1713003 and CIF-1763734. ZL acknowledges Amazon AI, Salesforce Research, Facebook, UPMC, Abridge, and the Center for Machine Learning and Health for their generous support of ACMI Lab’s research on machine learning under distribution shift.

References

- [1] A. Alexandari, A. Kundaje, and A. Shrikumar. Adapting to label shift with bias-corrected calibration. In *arXiv preprint arXiv:1901.06852*, 2019.
- [2] K. Aizzadenesheli, A. Liu, F. Yang, and A. Anandkumar. Regularized learning for domain adaptation under label shifts. In *International Conference on Learning Representations (ICLR)*, 2019.
- [3] S. Ben-David, T. Lu, T. Luu, and D. Pál. Impossibility Theorems for Domain Adaptation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- [4] G. Blanchard, G. Lee, and C. Scott. Semi-supervised novelty detection. *Journal of Machine Learning Research*, 11(Nov):2973–3009, 2010.
- [5] J. Byrd and Z. C. Lipton. What is the effect of importance weighting in deep learning? In *International Conference on Machine Learning (ICML)*, 2019.
- [6] C. Cortes and M. Mohri. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519, 2014.
- [7] C. Cortes, Y. Mansour, and M. Mohri. Learning Bounds for Importance Weighting. In *Advances in Neural Information Processing Systems (NIPS)*, 2010.
- [8] M. C. Du Plessis and M. Sugiyama. Semi-supervised learning of class balance under class-prior change by distribution matching. *Neural Networks*, 50:110–119, 2014.
- [9] A. Gretton, A. J. Smola, J. Huang, M. Schmittfull, K. M. Borgwardt, and B. Schölkopf. Covariate Shift by Kernel Mean Matching. *Journal of Machine Learning Research (JMLR)*, 2009.
- [10] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning (ICML)*, 2017.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [12] A. Krizhevsky and G. Hinton. Learning Multiple Layers of Features from Tiny Images. Technical report, Citeseer, 2009.

- [13] V. Kuleshov and P. S. Liang. Calibrated structured prediction. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.
- [14] A. Kumar, P. S. Liang, and T. Ma. Verified uncertainty calibration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [15] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-Based Learning Applied to Document Recognition. *Proceedings of the IEEE*, 86, 1998.
- [16] Z. C. Lipton, Y.-X. Wang, and A. Smola. Detecting and Correcting for Label Shift with Black Box Predictors. In *International Conference on Machine Learning (ICML)*, 2018.
- [17] H. Ramaswamy, C. Scott, and A. Tewari. Mixture proportion estimation via kernel embeddings of distributions. In *International Conference on Machine Learning*, pages 2052–2060, 2016.
- [18] M. Saerens, P. Latinne, and C. Decaestecker. Adjusting the Outputs of a Classifier to New a Priori Probabilities: A Simple Procedure. *Neural Computation*, 2002.
- [19] B. Schölkopf, D. Janzing, J. Peters, E. Sgouritsa, K. Zhang, and J. Mooij. On Causal and Anticausal Learning. In *International Conference on Machine Learning (ICML)*, 2012.
- [20] H. Shimodaira. Improving Predictive Inference Under Covariate Shift by Weighting the Log-Likelihood Function. *Journal of Statistical Planning and Inference*, 2000.
- [21] C. M. Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981.
- [22] A. Storkey. When Training and Test Sets Are Different: Characterizing Learning Transfer. *Dataset Shift in Machine Learning*, 2009.
- [23] J. A. Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 2015.
- [24] J. Vaicenavicius, D. Widmann, C. Andersson, F. Lindsten, J. Roll, and T. B. Schön. Evaluating model calibration in classification. In *International Conference on Machine Learning (ICML)*, 2019.
- [25] S. van de Geer. *Empirical Processes in M-estimation*, volume 6. Cambridge university press, 2000.
- [26] A. W. van der Vaart and J. A. Wellner. Weak convergence. In *Weak convergence and empirical processes*. Springer, 1996.
- [27] B. Zadrozny. Learning and Evaluating Classifiers Under Sample Selection Bias. In *International Conference on Machine Learning (ICML)*, 2004.
- [28] K. Zhang, B. Schölkopf, K. Muandet, and Z. Wang. Domain Adaptation Under Target and Conditional Shift. In *International Conference on Machine Learning (ICML)*, 2013.

A MLLS Algorithm

Algorithm 1 Maximum Likelihood Label Shift estimation

input : Labeled validation samples from source and unlabeled test samples from target. Trained blackbox model \hat{f} , model class \mathcal{G} and loss function l for calibration (for instance, MSE or negative log-likelihood).

- 1: On validation data minimize the loss l over class \mathcal{G} to obtain $f = g \circ \hat{f}$.
- 2: Solve the optimization problem (5) using f to get \hat{w} .

output : MLLS estimate \hat{w}

Step 1. description. Let the model class used for *post-hoc calibration* be represented by \mathcal{G} . Given a validation dataset $\{(x_{v1}, y_{v1}), \dots, (x_{vn}, y_{vn})\}$ sampled from the source distribution P_s we compute, $\{(\hat{f}(x_{v1}), y_{v1}), (\hat{f}(x_{v2}), y_{v2}), \dots, (\hat{f}(x_{vn}), y_{vn})\}$, applying our classifier \hat{f} to the data. Using this we estimate a function,

$$\hat{g} = \arg \min_{g \in \mathcal{G}} \sum_{i=1}^n \ell(g \circ \hat{f}(x_{vi}), y_{vi}), \quad (9)$$

where the loss function ℓ can be the negative log-likelihood or squared error. Experimentally we observe same performance with both the loss functions. Subsequently, we can apply the calibrated predictor $\hat{g} \circ \hat{f}$.

Our experiments follow Alexandari et al. [1], who leverage BCTS⁷ to calibrate their models. BCTS extends temperature scaling [10] by incorporating per-class bias terms. Formally, a function $g : \Delta^{k-1} \mapsto \Delta^{k-1}$ in the BCTS class \mathcal{G} , is given by

$$g_j(x) = \frac{\exp[\log(x_j)/T + b_j]}{\sum_i \exp[\log(x_i)/T + b_i]} \quad \forall j \in \mathcal{Y}$$

where $\{T, b_1, \dots, b_{|\mathcal{Y}|}\}$ are the $|\mathcal{Y}| + 1$ parameters to be learned.

B Prior Work on Label Shift Estimation

Dataset shifts are predominantly studied under two scenarios: covariate shift and label shift [22]. Schölkopf et al. [19] articulates connections between label shift and covariate shift with anti-causal and causal models respectively. Covariate shift is well explored in past [28, 27, 7, 6, 9].

Approaches for estimating label shift (or prior shift) can be categorized into three classes:

1. Methods that leverage Mixture Proportion Estimation (MPE) [4, 17] techniques to estimate the target label distribution. MPE estimate in general (e.g. Blanchard et al. [4]) needs explicit calculations of $p_s(x|y)$ ($= p_t(x|y)$) which is infeasible for high dimensional data. More recent methods for MPE estimation, i.e. Ramaswamy et al. [17], uses Kernel embeddings, which like many kernel methods, require the inversion of an $n \times n$ Gram matrix. The $\mathcal{O}(n^3)$ complexity makes them infeasible for large datasets, practically used in deep learning these days;
2. Methods that directly operate in RKHS for distribution matching [28, 8]. Zhang et al. [28] extend the kernel mean matching approach due to Gretton et al. [9] to the label shift problem. Instead of minimizing maximum mean discrepancy, Du Plessis and Sugiyama [8] explored minimizing PE divergence between the kernel embeddings to estimate the target label distribution. Again, both the methods involve inversion of an $n \times n$ kernel matrix, rendering them infeasible for large datasets; and
3. Methods that work in low dimensional setting [16, 2, 18] by directly estimating $p_t(y)/p_s(y)$ to avoid the curse of dimensionality. These methods leverage an off-the-shelf predictor to estimate the label shift ratio.

In this paper, we primarily focus on unifying methods that fall into the third category.

⁷Motivated by the strong empirical results in Alexandari et al. [1], we use BCTS in our experiments as a surrogate for canonical calibration.

C Marginal calibration is insufficient to achieve consistency

In this section, we will illustrate insufficiency of *marginal calibration* to achieve consistency. For completeness, we first define margin calibration:

Definition 3 (Marginal calibration). *A prediction model $f : \mathcal{X} \mapsto \Delta^{k-1}$ is marginally calibrated on the source domain if for all $x \in \mathcal{X}$ and $j \in \mathcal{Y}$,*

$$P_s(y = j|f_j(x)) = f_j(x).$$

Intuitively, this definition captures per-label calibration of the classifier which is strictly less restrictive than requiring canonical calibration. In the example, we construct a classifier on discrete \mathcal{X} which is marginally calibrated, but not canonically calibrated. With the constructed example, we show that the population objective (4) yields inconsistent estimates.

Example. Assume $\mathcal{X} = \{x_1, x_2, x_3, x_4, x_5, x_6\}$ and $\mathcal{Y} = \{1, 2, 3\}$. Suppose the predictor $f(x)$ and $P_s(y|f(x))$ are given as,

$f(x)$	y=1	y=2	y=3	$P_s(y f(x))$	y=1	y=2	y=3
x_1	0.1	0.2	0.7	x_1	0.2	0.1	0.7
x_2	0.1	0.7	0.2	x_2	0.0	0.8	0.2
x_3	0.2	0.1	0.7	x_3	0.1	0.2	0.7
x_4	0.2	0.7	0.1	x_4	0.3	0.6	0.1
x_5	0.7	0.1	0.2	x_5	0.8	0.0	0.2
x_6	0.7	0.2	0.1	x_6	0.6	0.3	0.1

Clearly, the prediction $f(x)$ is marginally calibrated. We have one more degree to freedom to choose, which is the source marginal distribution on \mathcal{X} . For simplicity let's assume $p_s(x_i) = 1/6$ for all $i = \{1, \dots, 6\}$. Thus, we have $p_s(y = j) = 1/3$ for all $j = \{1, 2, 3\}$. Note, with our assumption of the source marginal on x , we get $P_t(x_i|y = j) = P_s(x_i|y = j) = P_s(y = j|f(x_i))/2$. This follows as $x \mapsto f(x)$ is an one-to-one mapping.

Now, assume a shift i.e. prior on \mathcal{Y} for the target distribution of the form $[\alpha, \beta, 1 - \alpha - \beta]$. With the label shift assumption, we get

$$\forall i \quad p_t(x_i) = \frac{1}{2} (\alpha P_s(y = 1|f(x_i)) + \beta P_s(y = 2|f(x_i)) + (1 - \beta - \alpha) P_s(y = 3|f(x_i))).$$

Assume the importance weight vector as w . Clearly, we have $w_1 + w_2 + w_3 = 3$. Re-writing the population MLLS objective (4), we get the maximisation problem as

$$\arg \max_w \sum_{i=1}^6 p_t(x_i) \log(f(x_i)^T w). \quad (10)$$

Differentiating (10) with respect to w_1 and w_2 , we get two high order equations, solving which give us the MLLS estimate w_f . To show inconsistency, it is enough to consider one instantiation of α and β such that $|3\alpha - w_1| + |3\beta - w_2| + |w_1 + w_2 - 3\alpha - 3\beta| \neq 0$. Assuming $\alpha = 0.8$ and $\beta = 0.1$ and solving (10) using numerical methods, we get $w_f = [2.505893, 0.240644, 0.253463]$. As $w = [2.4, 0.3, 0.3]$, we have $w_f \neq w$ concluding the proof.

D Proofs from Section 4

Lemma 1 (Identifiability). *If the set of distributions $\{p(z|y) : y = 1, \dots, k\}$ are linearly independent, then for any w that satisfies (2), we must have $w = w^*$. This condition is also necessary in general: if the linear independence does not hold then there exists a problem instance where we have $w, w^* \in \mathcal{W}$ satisfying (2) while $w \neq w^*$.*

Proof. First we prove sufficiency. If there exists $w \neq w^*$ such that (2) holds, then we have $\sum_{y=1}^k p_s(z, y)(w_y - w_y^*) = 0$ for all $z \in \mathcal{Z}$. As $w - w^*$ is not the zero vector, $\{p_s(z, y), y = 1, \dots, k\}$ are linearly dependent. Since $p_s(z, y) = p_s(y)p(z|y)$ and $p_s(y) > 0$ for all y (by assumption), we

also have that $\{p(z|y), y = 1, \dots, k\}$ are linearly dependent. By contradiction, we show that the linear independence is necessary.

To show necessity, assume $w_y^* = \frac{1}{kp_s(y)}$ for $y = 1, \dots, k$. We know that w^* satisfies (2) by definition. If linear independence does not hold, then there exists a vector $v \in \mathbb{R}^k$ such that $v \neq 0$ and $\sum_{y=1}^k p_s(z, y)v_y = 0$ for all $z \in \mathcal{Z}$. Since the w^* we construct is not on the boundary of \mathcal{W} , we can scale v such that $w^* + \alpha v \in \mathcal{W}$ where $\alpha \geq 0$ and $v \neq 0$. Therefore, setting $w = w^* + \alpha v$ gives another solution for (2), which concludes the proof. \square

Lemma 2. *If f is calibrated, then the two objectives (3) and (4) are identical when \mathcal{Z} is chosen as Δ^{k-1} and $p(z|x)$ is defined to be $\delta_{f(x)}$.*

Proof. The proof follows a sequence of straightforward manipulations. In more detail,

$$\begin{aligned} \mathbb{E}_t [\log f(x)^T w] &= \int p_t(x) \log[f(x)^T w] dx \\ &= \int \int p_t(x) p(z|x) \log[f(x)^T w] dx dz \\ &= \int \int p_t(x) p(z|x) \mathbb{1}\{f(x) = z\} \log[f(x)^T w] dx dz \\ &= \int \int p_t(x) p(z|x) \log[z^T w] dx dz \\ &= \int p_t(z) \log[z^T w] dz \\ &= \int p_t(z) \log \left[\sum_{y=1}^k p_s(y|z) w \right] dz, \end{aligned}$$

where the final step uses the fact that f is calibrated. \square

Theorem 1 (Population consistency of MLLS). *If a predictor $f : \mathcal{X} \mapsto \Delta^{k-1}$ is calibrated and the distributions $\{p(f(x)|y) : y = 1, \dots, k\}$ are strictly linearly independent, then w^* is the unique maximizer of the MLLS objective (4).*

Proof. According to Lemma 2 we know that maximizing (4) is the same as maximizing (3) with $p(z|x) = \delta_{f(x)}$, thus also the same as minimizing the KL divergence between $p_t(z)$ and $p_w(z)$. Since $p_t(z) \equiv p_{w^*}(z)$ we know that w^* is a minimizer of the KL divergence such that the KL divergence is 0. We also have that $\text{KL}(p_t(z), p_w(z)) = 0$ if and only if $p_t(z) \equiv p_w(z)$, so all maximizers of (4) should satisfy (2). According to Lemma 1, if the strict linear independence holds, then w^* is the unique solution of (2). Thus w^* is the unique maximizer of (4). \square

Proposition 1. *For a calibrated predictor f , the following statements are equivalent:*

- (1) $\{p(f(x)|y) : y = 1, \dots, k\}$ are strictly linearly independent.
- (2) $\mathbb{E}_s [f(x)f(x)^T]$ is invertible.
- (3) The soft confusion matrix of f is invertible.

Proof. We first show the equivalence of (1) and (2). If f is calibrated, we have $p_s(f(x))f_y(x) = p_s(y)p(f(x)|y)$ for any x, y . Then for any vector $v \in \mathbb{R}^k$ we have

$$\sum_{y=1}^k v_y p(f(x)|y) = \sum_{y=1}^k \frac{v_y}{p_s(y)} p_s(y) p(f(x)|y) = \sum_{y=1}^k \frac{v_y}{p_s(y)} p_s(f(x)) f_y(x) = p_s(f(x)) \sum_{y=1}^k \frac{v_y}{p_s(y)} f_y(x). \quad (11)$$

On the other hand, we can have

$$\mathbb{E}_s [f(x)f(x)^T] = \int f(x)f(x)^T p_s(f(x))d(f(x)). \quad (12)$$

If $\{p(f(x)|y) : y = 1, \dots, k\}$ are linearly dependent, then there exist $v \neq 0$ such that (11) is zero for any x . Consequently, there exists a non-zero vector u with $u_y = v_y/p_s(y)$ such that $u^T f(x) = 0$ for any x satisfying $p_s(f(x)) > 0$, which means $u^T \mathbb{E}_s [f(x)f(x)^T] u = 0$ and thus $\mathbb{E}_s [f(x)f(x)^T]$ is not invertible. On the other hand, if $\mathbb{E}_s [f(x)f(x)^T]$ is non-invertible, then there exist some $u \neq 0$ such that $u^T \mathbb{E}_s [f(x)f(x)^T] u = 0$. Further as $u^T \mathbb{E}_s [f(x)f(x)^T] u = \int u^T f(x)f(x)^T u p_s(x)dx = \int |f(x)^T u| p_s(x)dx$. As a result, the vector v with $v_y = p_s(y)u_y$ satisfies that (11) is zero for any x , which means $\{p(f(x)|y) : y = 1, \dots, k\}$ are not strictly linearly independent.

Let C be the soft confusion matrix of f , then

$$\begin{aligned} C_{ij} &= p_s(\hat{y} = i, y = j) = \int d(f(x)) f_i(x)p(f(x)|y = j)p_s(y = j) \\ &= \int f_i(x)f_j(x)p_s(f(x))d(f(x)). \end{aligned}$$

Therefore, we have $C = \mathbb{E}_s [f(x)f(x)^T]$, which means (2) and (3) are equivalent. \square

We introduce some notation before proving consistency. Let $\mathcal{P} = \{\langle f, w \rangle | w \in \mathcal{W}\}$ be the class of densities⁸ for a given calibrated predictor f . Suppose $\hat{p}_n, p_0 \in \mathcal{P}$ are densities corresponding to MLE estimate and true weights, respectively. We use $h(p_1, p_2)$ to denote the Hellinger distance and $\text{TV}(p_1, p_2)$ to denote the total variation distance between two densities p_1, p_2 . $H_r(\delta, \mathcal{P}, P)$ denotes δ -entropy for class \mathcal{P} with respect to metric $L_r(P)$. Similarly, $H_{r,B}(\delta, \mathcal{P}, P)$ denotes the corresponding bracketing entropy. Moreover, P_n denotes the empirical random distribution that puts uniform mass on observed samples x_1, x_2, \dots, x_n . Before proving consistency we need to re-state two results:

Lemma 6 (Lemma 2.1 [25]). *If P is a probability measure, for all $1 \leq r < \infty$, we have*

$$H_{r,B}(\delta, \mathcal{G}, P) \leq H_\infty(\delta/2, \mathcal{G}) \quad \text{for all } \delta > 0.$$

Lemma 7 (Corollary 2.7.10 [26]). *Let \mathcal{F} be the class of convex functions $f : C \mapsto [0, 1]$ defined on a compact, convex set $C \subset \mathbb{R}^d$ such that $|f(x) - f(y)| \leq L \|x - y\|$ for every x, y . Then*

$$H_\infty(\delta, \mathcal{F}) \leq K \left(\frac{L}{\delta} \right)^{d/2},$$

for a constant K that depends on the dimension d and C .

We can now present our proof of consistency, which is based on Theorem 4.6 from van de Geer [25]:

Lemma 8 (Theorem 4.6 [25]). *Let \mathcal{P} be convex and define class $\mathcal{G} = \left\{ \frac{2p}{p+p_0} | p \in \mathcal{P} \right\}$. If*

$$\frac{1}{n} H_1(\delta, \mathcal{G}, P_n) \rightarrow_P 0, \quad (13)$$

then $h(\hat{p}_n, p_0) \rightarrow 0$ almost surely.

Theorem 2 (Consistency of MLLS). *If f satisfies the conditions in Theorem 1, then \hat{w}_f in (5) converges to w^* almost surely.*

Proof. Assume the maximizer of (5) is \hat{w}_f and $p_0 = \langle f, w^* \rangle$. Define class $\mathcal{G} = \left\{ \frac{2p}{p+p_0} | p \in \mathcal{P} \right\}$. To prove consistency, we first bound the bracketing entropy for class \mathcal{G} using Lemma 6 and Lemma 7.

⁸Note that we use the term *density* loosely here for convenience. The actual density is $\langle f(x), w \rangle \cdot p_s(x)$ but we can ignore $p_s(x)$ because it does not depend on our parameters.

Clearly \mathcal{P} is linear in parameters and hence, convex. Gradient of function $g \in \mathcal{G}$ is given by $\frac{2p_0}{(p+p_0)^2}$ which in turn is bounded by $\frac{2}{p_0}$. Under assumptions of Condition 1, the functions in \mathcal{G} are Lipschitz with constant $2/\tau$. We can bound the bracketing entropy $H_{2,B}(\delta, \mathcal{G}, P)$ using Lemma 7 and Lemma 6 as

$$H_{2,B}(\delta, \mathcal{G}, P) \leq H_\infty(\delta, \mathcal{G}) \leq K_1 \left(\frac{1}{\delta\tau} \right)^{k/2},$$

for some constant K_1 that depends on k .

On the other hand, for cases where p_0 can be arbitrarily close to zero, i.e., Condition 1 doesn't hold true, we define $\tau(\delta)$ and \mathcal{G}_τ as

$$\tau(\delta) = \sup \left\{ \tau \geq 0 \mid \int_{p_0 \leq \tau} p_0 dx \leq \delta^2 \right\}, \quad (14)$$

$$\mathcal{G}_\tau = \left\{ \frac{2p}{p+p_0} \mathbb{1}_{\{p_0 \geq \tau\}} \mid p \in \mathcal{P} \right\}.$$

Using triangle inequality, for any $g_1, g_2 \in \mathcal{G}$, we have

$$\begin{aligned} \int \|g_1 - g_2\|^2 dx &\leq \int \|g_1 - g_2\|^2 \mathbb{1}_{\{p_0 \leq \tau\}} dx + \int \|g_1 - g_2\|^2 \mathbb{1}_{\{p_0 \geq \tau\}} dx \\ &\leq 2 \int \mathbb{1}_{\{p_0 \leq \tau\}} dx + \int \|g_1 - g_2\|^2 \mathbb{1}_{\{p_0 \geq \tau\}} dx. \end{aligned} \quad (15)$$

Assume $\tau(\delta)$ such that (14) is satisfied. Using (15), we have

$$H_{2,B}(\delta, \mathcal{G}, P) \leq H_{2,B}(\sqrt{3}\delta, \mathcal{G}_{\tau(\delta)}, P).$$

Thus, for the cases where p_0 can be arbitrarily close to zero, instead of bounding $H_{2,B}(\delta, \mathcal{G}, P)$, we bound $H_{2,B}(\delta, \mathcal{G}_{\tau(\delta)}, P)$. For any $\delta > 0$, there is a compact subset $K_\delta \in \mathcal{X}$, such that $p_s(X \setminus K_\delta) < \delta$. Using arguments similar to above, function $g \in \mathcal{G}_{\tau(\delta)}$ is Lipschitz with constant $2/\tau(\delta) > 0$. Again using Lemma 7 and Lemma 6, we conclude

$$H_{2,B}(2\delta, \mathcal{G}_{\tau(\delta)}, P) \leq H_\infty(\delta, \mathcal{G}_{\tau(\delta)}) \leq K_2 \left(\frac{1}{\delta\tau(\delta)} \right)^k,$$

for some constant K_2 that depends on k . Finally, we use Lemma 8 to conclude $h(\hat{p}_n, p_0) \rightarrow_{\text{a.s.}} 0$. Further, as $\text{TV}(\hat{p}_n, p_0) \leq h(\hat{p}_n, p_0)$, we have $h(\hat{p}_n, p_0) \rightarrow_{\text{a.s.}} 0$ implies $\text{TV}(\hat{p}_n, p_0) \rightarrow_{\text{a.s.}} 0$. Further

$$\begin{aligned} \|\hat{w}_f - w^*\|^2 &\leq \frac{1}{\lambda_{\min}} \int |f(x)^T (\hat{w}_f - w^*)|^2 p_s(x) dx \\ &\leq \frac{\sup_x \{|f(x)^T (\hat{w}_f - w^*)|\}}{\lambda_{\min}} \underbrace{\int |f(x)^T (\hat{w}_f - w^*)| p_s(x) dx}_{\text{TV}(\hat{p}_n, p_0)}, \end{aligned} \quad (16)$$

where λ_{\min} is the minimum eigenvalue of covariance matrix $[\int f(x)f(x)^T p_s(x) dx]$. Note using Proposition 1, we have $\lambda_{\min} > 0$. Thus, we conclude $\|\hat{w}_f - w^*\| \rightarrow_{\text{a.s.}} 0$. \square

Example 1. Consider a mixture of two Gaussians with $p_s(x|y=0) := \mathcal{N}(\mu, 1)$ and $p_s(x|y=1) := \mathcal{N}(-\mu, 1)$. We suppose that the source mixing coefficients are both $\frac{1}{2}$, while the target mixing coefficients are $\alpha (\neq \frac{1}{2}), 1-\alpha$. Assume a class of probabilistic threshold classifiers: $f(x) = [1-c, c]$ for $x \geq 0$, otherwise $f(x) = [c, 1-c]$ with $c \in [0, 1]$.

Then the population error of MLLS is given by

$$4 \left| \frac{(1-2\alpha)(p_s(x \geq 0|y=0) - c)}{1-2c} \right|,$$

which is zero only if $c = p_s(x \geq 0|y=0)$ for a non-degenerate classifier.

Proof. The intuition behind the construction is, for such an Example, we can get a closed form solution for the population MLLS and hence allows a careful analysis of the estimation error. The classifier $f(x)$ predicts class 0 with probability c and class 1 with probability $1 - c$ for $x \geq 0$, and vice-versa for $x < 0$. Using such a classifier, the weight estimator is given by:

$$\begin{aligned}\hat{w} &= \arg \min_w \mathbb{E} [\log \langle f(x), w \rangle] \\ &\stackrel{(i)}{=} \arg \min_{w_0} \left[\int_{-\infty}^0 \log((1-c)w_0 + c(2-w_0))p_t(x)dx + \int_0^{\infty} \log(cw_0 + (1-c)(2-w_0))p_t(x)dx \right] \\ &\stackrel{(ii)}{=} \arg \min_{w_0} [\log((1-c)w_0 + c(2-w_0))p_t(x \leq 0) + \log(cw_0 + (1-c)(2-w_0))p_t(x \geq 0)] ,\end{aligned}$$

where equality (i) follows from $w_1 = 2 - w_0$ and the predictor function and (ii) follows from the fact that within each integral, the term inside the log is independent of x . Differentiating w.r.t. to w_0 , we have:

$$\begin{aligned}\frac{1-2c}{2c+w_0-2cw_0}p_t(x \leq 0) + \frac{2c-1}{2cw_0+2-2c-w_0}p_t(x \geq 0) &= 0 \\ \frac{1}{2c+w_0-2cw_0}p_t(x \leq 0) + \frac{-1}{2cw_0+2-2c-w_0}(1-p_t(x \leq 0)) &= 0 \\ (2cw_0+2-2c-w_0)p_t(x \leq 0) - (2c+w_0-2cw_0)(1-p_t(x \leq 0)) &= 0 \\ 2p_t(x \leq 0) - 2c - w_0 + 2cw_0 &= 0 ,\end{aligned}$$

which gives $w_0 = \frac{2p_t(x \leq 0) - 2c}{1 - 2c}$. Thus for the population MLLS estimate, the estimation error is given by

$$\|\hat{w} - w^*\| = 2|w_0 - 2\alpha| = 4 \left| \frac{(1-2\alpha)(p_s(x \geq 0|y=0) - c)}{1-2c} \right| .$$

□

E Proofs from Section 5

The gradient of the MLLS objective can be written as

$$\nabla_w \mathcal{L}(w, f) = \mathbb{E}_t \left[\frac{f(x)}{f(x)^T w} \right] , \quad (17)$$

and the Hessian is

$$\nabla_w^2 \mathcal{L}(w, f) = -\mathbb{E}_t \left[\frac{f(x)f(x)^T}{(f(x)^T w)^2} \right] . \quad (18)$$

We use $\lambda_{\min}(X)$ to denote the minimum eigenvalue of the matrix X .

Lemma 9 (Theorem 5.1.1 [23]). *Let X_1, X_2, \dots, X_n be a finite sequence of identically distributed independent, random, symmetric matrices with common dimension k . Assume $0 \preceq X \preceq R \cdot I$ and $\mu_{\min} I \preceq \mathbb{E}[X] \preceq \mu_{\max} I$. With probability at least $1 - \delta$,*

$$\lambda_{\min} \left(\frac{1}{n} \sum_{i=1}^n X_i \right) \geq \mu_{\min} - \sqrt{\frac{2R\mu_{\min} \log(\frac{k}{\delta})}{n}} . \quad (19)$$

Lemma 3. *For any predictor f that satisfies Condition 1, we have $\|w_f - \hat{w}_f\| \leq \sigma_{f, w_f}^{-1} \mathcal{O}_p(m^{-1/2})$.*

Proof. We present our proof in two steps. Step-1 is the non-probabilistic part, i.e., bounding the error $\|\hat{w}_f - w_f\|$ in terms of the gradient difference $\|\nabla_w \mathcal{L}(w_f, f) - \nabla_w \mathcal{L}_m(w_f, f)\|$. This step uses Taylor's expansion upto second order terms for empirical log-likelihood around the true w^* . Step-2 involves deriving a concentration on the gradient difference using the Lipschitz property implied by Condition 1. Combining these two steps along with Lemma 22 concludes the proof. Now we detail each of these steps.

Step-1. We represent the empirical Negative Log-Likelihood (NLL) function with \mathcal{L}_m by absorbing the negative sign to simplify notation. Using a Taylor expansion, we have

$$\mathcal{L}_m(\hat{w}_f, f) = \mathcal{L}_m(w_f, f) + \langle \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle + \frac{1}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(\tilde{w}, f_c)(\hat{w}_f - w_f),$$

where $\tilde{w} \in [\hat{w}_f, w_f]$. With the assumption $f^T w_f \geq \tau$, we have $\nabla_w^2 \mathcal{L}_m(\tilde{w}, f) \geq \frac{\tau^2}{\min p_s(y)^2} \nabla_w^2 \mathcal{L}_m(w_f, f)$. Let $\kappa = \frac{\tau^2}{\min p_s(y)^2}$. Using this we get,

$$\begin{aligned} \mathcal{L}_m(\hat{w}_f, f) &\geq \mathcal{L}_m(w_f, f) + \langle \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle + \frac{\kappa}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f) \\ \underbrace{\mathcal{L}_m(\hat{w}_f, f) - \mathcal{L}_m(w_f, f)}_{\text{I}} - \langle \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle &\geq \frac{\kappa}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f), \end{aligned}$$

where term-I is less than zero as \hat{w}_f is the minimizer of empirical NLL $\mathcal{L}_m(\hat{w}_f, f)$. Ignoring term-I and re-arranging a few terms we get:

$$-\langle \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle \geq \frac{\kappa}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f),$$

With first order optimality on w_f , $\langle \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle \geq 0$. Plugging in this, we have,

$$\langle \nabla_w \mathcal{L}_m(w_f, f) - \nabla_w \mathcal{L}_m(w_f, f), \hat{w}_f - w_f \rangle \geq \frac{\kappa}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f),$$

Using Holder's inequality on the LHS we have,

$$\|\nabla_w \mathcal{L}_m(w_f, f) - \nabla_w \mathcal{L}_m(w_f, f)\| \|\hat{w}_f - w_f\| \geq \frac{\kappa}{2}(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f).$$

Let $\hat{\sigma}_{f, w_f}$ be the minimum eigenvalue of $\nabla_w^2 \mathcal{L}_m(w^*, f_c)$. Using the fact that $(\hat{w}_f - w_f)^T \nabla_w^2 \mathcal{L}_m(w_f, f)(\hat{w}_f - w_f) \geq \hat{\sigma}_{\min} \|\hat{w}_f - w_f\|^2$, we get,

$$\|\nabla_w \mathcal{L}_m(w_f, f) - \nabla_w \mathcal{L}_m(w_f, f)\| \geq \frac{\kappa \hat{\sigma}_{f, w_f}}{2} \|\hat{w}_f - w_f\|. \quad (20)$$

Step-2. The empirical gradient is $\nabla_w \mathcal{L}_m(w_f, f) = \sum_{i=1}^m \frac{\nabla_w \mathcal{L}_1(x_i, w_f, f)}{m}$ where $\nabla \mathcal{L}_1(x_i, w_f, f) = \left[\frac{f_1(x_i)}{\langle f(x_i), w_f \rangle} \cdots \frac{f_l(x_i)}{\langle f(x_i), w_f \rangle} \cdots \frac{f_k(x_i)}{\langle f(x_i), w_f \rangle} \right]_{(k)}$. With the lower bound τ on $f^T w_f$, we can upper bound the gradient terms as

$$\|\nabla_w \mathcal{L}_1(x, w_f, f)\| \leq \frac{\|f\|}{\tau} \leq \frac{\|f\|_1}{\tau} \leq \frac{1}{\tau}.$$

As the gradient terms decompose and are independent, using Hoeffding's inequality we have with probability at least $1 - \frac{\delta}{2}$,

$$\|\nabla_w \mathcal{L}_m(w_f, f) - \nabla_w \mathcal{L}_m(w_f, f)\| \leq \frac{1}{2\tau} \sqrt{\frac{\log(4/\delta)}{m}}. \quad (21)$$

Let σ_{f, w_f} be the minimum eigenvalue of $\nabla_w^2 \mathcal{L}_m(w_f, f)$. Using lemma 9, with probability at least $1 - \frac{\delta}{2}$,

$$\frac{\hat{\sigma}_{f, w_f}}{\sigma_{f, w_f}} \geq 1 - \tau \sqrt{\frac{\log(2k/\delta)}{m}}. \quad (22)$$

Plugging (21) and (22) in (20), and applying a union bound, we conclude that with probability at least $1 - \delta$,

$$\begin{aligned} \|\hat{w}_f - w_f\|_2 &\leq \frac{1}{\kappa\tau} \left(\sigma_{f, w_f} - \sigma_{f, w_f} \tau \sqrt{\frac{\log(2k/\delta)}{m}} \right)^{-1} \left(\sqrt{\frac{\log(4/\delta)}{m}} \right) \\ &\leq \frac{1}{\kappa\tau} \frac{1}{\sigma_{f, w_f}} \left(1 + \tau \sqrt{\frac{\log(2k/\delta)}{m}} \right) \sqrt{\frac{\log(4/\delta)}{m}}. \end{aligned}$$

Neglecting the order m term and letting $c = \frac{1}{\kappa\tau}$, we have

$$\|\hat{w}_f - w_f\| \leq \frac{c}{\sigma_{f, w_f}} \sqrt{\frac{\log(4/\delta)}{m}}.$$

□

Lemma 4. For any predictor f and any calibrated predictor f_c that satisfies Condition 1, we have $\|w_f - w^*\| \leq \sigma_{f,w^*}^{-1} \cdot C \cdot \mathbb{E}_t [\|f - f_c\|]$, for some constant C .

If we set $f_c(x) = p_s(y|f(x))$, which is a calibrated predictor (Proposition 2), we can bound the error in terms of the calibration error of f on the source data⁹: $\|w_f - w^*\| \leq \sigma_{f,w^*}^{-1} \cdot C \cdot \mathcal{E}(f)$.

Proof. We present our proof in two steps. Note, all calculations are non-probabilistic. Step-1 involves bounding the error $\|w_f - w^*\|$ in terms of the gradient difference $\|\nabla_w \mathcal{L}(w^*, f_c) - \nabla_w \mathcal{L}(w^*, f)\|$. This step uses Taylor's expansion on $\mathcal{L}(w_f, f)$ upto the second order term for population log-likelihood around the true w^* . Step-2 involves deriving a bound on the gradient difference in terms of the difference $\|f - f_c\|$ using the Lipschitz property implied by Condition 1. Further, for a crude calibration choice of $f_c(x) = p_s(\cdot|x)$, the gradient difference can be bounded by miscalibration error. We now detail both of these steps.

Step-1. Similar to Lemma 3, we represent with \mathcal{L} by absorbing the negative sign to simplify notation. Using the Taylor expansion, we have

$$\mathcal{L}(w_f, f) \geq \mathcal{L}(w^*, f) + \langle \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle + \frac{1}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(\tilde{w}, f) (w_f - w^*),$$

where $\tilde{w} \in [w_f, w^*]$. With the assumption $f^T w^* \geq \tau$, we have $\nabla_w^2 \mathcal{L}(\tilde{w}, f) \geq \frac{\tau^2}{\min p_s(y)^2} \nabla_w^2 \mathcal{L}(w^*, f)$. Let $\kappa = \frac{\tau^2}{\min p_s(y)^2}$. Using this we get,

$$\begin{aligned} \mathcal{L}(w_f, f) &\geq \mathcal{L}(w^*, f) + \langle \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle + \frac{\kappa}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*) \\ \underbrace{\mathcal{L}(w_f, f) - \mathcal{L}(w^*, f)}_{\text{I}} &\geq \langle \nabla_w \mathcal{L}(w_f, f), w_f - w^* \rangle + \frac{\kappa}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*), \end{aligned}$$

where term-I is less than zero as w_f is the minimizer of NLL $\mathcal{L}(w, f)$. Ignoring that term and re-arranging a few terms we get

$$-\langle \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle \geq \frac{\kappa}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*).$$

With first order optimality on w^* , $\langle \nabla_w \mathcal{L}(w^*, f_c), w_f - w^* \rangle \geq 0$. Using this we have:

$$\begin{aligned} \langle \nabla_w \mathcal{L}(w^*, f_c), w_f - w^* \rangle - \langle \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle &\geq \frac{\kappa}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*), \\ \langle \nabla_w \mathcal{L}(w^*, f_c) - \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle &\geq \frac{\kappa}{2} (w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*). \end{aligned}$$

As before, let $\sigma_{f,w}$ be the minimum eigenvalue of $\nabla_w^2 \mathcal{L}(w^*, f)$. Using the fact that $(w_f - w^*)^T \nabla_w^2 \mathcal{L}(w^*, f) (w_f - w^*) \geq \sigma_{f,w} \|w_f - w^*\|^2$, we get

$$\langle \nabla_w \mathcal{L}(w^*, f_c) - \nabla_w \mathcal{L}(w^*, f), w_f - w^* \rangle \geq \frac{\kappa \sigma_{f,w}}{2} \|w_f - w^*\|^2.$$

Using Holder's inequality on the LHS and re-arranging terms gives

$$\|\nabla_w \mathcal{L}(w^*, f_c) - \nabla_w \mathcal{L}(w^*, f)\| \geq \frac{\kappa \sigma_{f,w}}{2} \|w_f - w^*\|. \quad (23)$$

Step-2. By lower bound assumptions $f_c^T w^* \geq \tau$ and $f^T w^* \geq \tau$, we have

$$\|\nabla_w \mathcal{L}(w^*, f_c) - \nabla_w \mathcal{L}(w^*, f)\| \leq \mathbb{E}_t [\|\nabla \mathcal{L}_1(x, w^*, f_c) - \nabla \mathcal{L}_1(x, w^*, f)\|] \leq \frac{1}{\tau^2} \mathbb{E}_t [\|f_c(x) - f(x)\|], \quad (24)$$

where the first inequality is implied by Jensen's inequality and the second is implied by the Lipschitz property of the gradient. Further, we have

$$\mathbb{E}_t [\|f_c(x) - f(x)\|] = \mathbb{E}_s \left[\frac{p_t(x)}{p_s(x)} \|f_c(x) - f(x)\| \right]$$

⁹We present two upper bounds because the second is more interpretable while the first is tighter.

$$\begin{aligned}
&\leq \mathbb{E}_s \left[\max_y \frac{p_t(y)}{p_s(y)} \|f_c(x) - f(x)\| \right] \\
&\leq \max_y \frac{p_t(y)}{p_s(y)} \mathbb{E}_s [\|f_c(x) - f(x)\|]. \tag{25}
\end{aligned}$$

Combining equations (23), (24), and (25), we have

$$\|w_f - w^*\| \leq \frac{2}{\kappa \sigma_{f,w} \tau^2} \max_y \frac{p_t(y)}{p_s(y)} \mathbb{E}_s [\|f_c(x) - f(x)\|]. \tag{26}$$

Further, if we set $f_c(x) = p_s(\cdot|f(x))$, which is a calibrated predictor according to Proposition 2, we can bound the error on the RHS in terms of the calibration error of f . Moreover, in the label shift estimation problem, we have the assumption that $p_s(y) \geq c > 0$ for all y . Hence, we have $\max_y p_t(y)/p_s(y) \leq 1/c$. Using Jensen's inequality, we get

$$\mathbb{E}_s \|f_c(x) - f(x)\| \leq \left(\mathbb{E}_s \|f_c(x) - f(x)\|^2 \right)^{\frac{1}{2}} = \mathcal{E}(f). \tag{27}$$

Plugging (27) back in (26), we get the required upper bound. \square

Proposition 3. *For any $w \in \mathcal{W}$, we have $\sigma_{f,w} \geq p_{s,\min} \sigma_f$ where σ_f is the minimum eigenvalue of $\mathbb{E}_t [f(x)f(x)^T]$ and $p_{s,\min} = \min_{y \in \mathcal{Y}} p_s(y)$. Furthermore, if f satisfies Condition 1, we have $p_{s,\min}^2 \cdot \sigma_f \leq \sigma_{f,w} \leq \tau^{-2} \cdot \sigma_f$, for $w \in \{w_f, w^*\}$.*

Proof. For any $v \in \mathbb{R}^k$, we have

$$v^T (-\nabla_w^2 \mathcal{L}(w, f)) v = \mathbb{E}_t \left[\frac{(v^T f(x))^2}{(f(x)^T w)^2} \right] \in \left[\frac{1}{a^2}, \frac{1}{b^2} \right] \cdot v^T \mathbb{E}_t [f(x)f(x)^T] v,$$

where

$$a = \max_{x:p_s(x)>0} f(x)^T w \leq \frac{1}{p_{s,\min}}$$

and

$$b = \min_{x:p_s(x)>0} f(x)^T w \geq \tau$$

if f satisfies Condition 1 and $w \in \{w_f, w^*\}$. Therefore, we have

$$p_{s,\min}^2 \cdot \sigma_f \leq \sigma_{f,w} \leq \tau^{-2} \cdot \sigma_f$$

for $w \in \{w_f, w^*\}$. \square

Lemma 5. *Let $f = g \circ \hat{f}$ be the predictor after post-hoc calibration with squared loss l and g belongs to a function class \mathcal{G} that satisfies the standard regularity conditions, we have*

$$\mathcal{E}(f) \leq \min_{g \in \mathcal{G}} \mathcal{E}(g \circ \hat{f}) + \mathcal{O}_p \left(n^{-1/2} \right). \tag{8}$$

Proof. Assume regularity conditions on the model class \mathcal{G}_θ (injectivity, Lipschitz-continuity, twice differentiability, non-singular Hessian, and consistency) as in Theorem 5.23 of Stein [21] hold true. Using the injectivity property of the model class as in Kumar et al. [14], we have for all $g_1, g_2 \in \mathcal{G}$,

$$\text{MSE}(g_1) - \text{MSE}(g_2) = \mathcal{E}(g_1)^2 - \mathcal{E}(g_2)^2. \tag{28}$$

Let $\hat{g}, g^* \in \mathcal{G}$ be models parameterized by $\hat{\theta}$ and θ^* , respectively. Using the strong convexity of the empirical mean squared error we have,

$$\text{MSE}_n(\hat{g}) \geq \text{MSE}_n(g^*) + \langle \nabla_{\theta} \text{MSE}_n(g^*), \hat{\theta} - \theta^* \rangle + \frac{\mu^2}{2} \|\hat{\theta} - \theta^*\|_2^2,$$

where μ is the parameter constant for strong convexity. Re-arranging a few terms, we have

$$\underbrace{\text{MSE}_n(\hat{g}) - \text{MSE}_n(g^*)}_{\text{I}} - \langle \nabla_{\theta} \text{MSE}_n(g^*), \hat{\theta} - \theta^* \rangle \geq \frac{\mu^2}{2} \|\hat{\theta} - \theta^*\|_2^2,$$

where term-I is less than zero because \hat{g} is the empirical minimizer of the mean-squared error. Ignoring term-I, we get:

$$\frac{\mu^2}{2} \|\hat{\theta} - \theta^*\|_2^2 \leq -\langle \nabla_{\theta} \text{MSE}_n(g^*), \hat{\theta} - \theta^* \rangle \leq \|\nabla_{\theta} \text{MSE}_n(g^*)\| \|\hat{\theta} - \theta^*\|.$$

As the assumed model class is Lipschitz w.r.t. θ , the gradient is bounded by Lipschitz constant $L = c_1$. $\mathbb{E}[\nabla_{\theta} \text{MSE}_n(g^*)] = 0$ as g^* is the population minimizer. Using Hoeffding's bound for bounded functions, we have with probability at least $1 - \delta$,

$$\|\hat{\theta} - \theta^*\|_2 \leq \frac{c_1}{\mu^2} \sqrt{\frac{\log(2/\delta)}{n}}. \quad (29)$$

Using the smoothness of the $\text{MSE}(g)$, we have

$$\text{MSE}(\hat{g}) - \text{MSE}(g^*) \leq c_2 \|\hat{\theta} - \theta^*\|_2^2, \quad (30)$$

where c_2 is the operator norm of the $\nabla^2 \text{MSE}(g^*)$. Combining (28), (29), and (30), we have for some universal constant $c = \frac{c_1 c_2}{\mu^2}$ with probability at least $1 - \delta$,

$$\mathcal{E}(\hat{g})^2 - \mathcal{E}(g^*)^2 \leq c \frac{\log(2/\delta)}{n}.$$

□

Moreover, with Lemma 4, depending on the degree of the miscalibration and the method involved to calibrate, we can bound the $\mathcal{E}(f)$. For example, if using vector scaling on a held out training data for calibration, we can use Lemma 5 to bound the calibration error $\mathcal{E}(f)$, i.e., with probability at least $1 - \delta$, we have

$$\mathcal{E}(f) \leq \sqrt{\min_{g \in \mathcal{G}} \mathcal{E}(g \circ f)^2 + c \frac{\log(2/\delta)}{n}} \leq \min_{g \in \mathcal{G}} \mathcal{E}(g \circ f) + \sqrt{c \frac{\log(2/\delta)}{n}}. \quad (31)$$

Plugging (27) and (31) into (26), we have with probability at least $1 - \delta$ that

$$\|w_f - w^*\| \leq \frac{1}{\kappa \sigma_{f,w} \tau^2} \left(\|w^*\|_2 \left(\sqrt{c \frac{\log(2/\delta)}{n}} + \min_{g \in \mathcal{G}} \mathcal{E}(g \circ f) \right) \right).$$