

Sample Selection Bias Correction Theory

Corinna Cortes¹, Mehryar Mohri^{2,1}, Michael Riley¹, and Afshin Rostamizadeh²

¹ Google Research,

76 Ninth Avenue, New York, NY 10011.

² Courant Institute of Mathematical Sciences,

251 Mercer Street, New York, NY 10012.

Abstract. This paper presents a theoretical analysis of sample selection bias correction. The sample bias correction technique commonly used in machine learning consists of reweighting the cost of an error on each training point of a biased sample to more closely reflect the unbiased distribution. This relies on weights derived by various estimation techniques based on finite samples. We analyze the effect of an error in that estimation on the accuracy of the hypothesis returned by the learning algorithm for two estimation techniques: a cluster-based estimation technique and kernel mean matching. We also report the results of sample bias correction experiments with several data sets using these techniques. Our analysis is based on the novel concept of *distributional stability* which generalizes the existing concept of point-based stability. Much of our work and proof techniques can be used to analyze other importance weighting techniques and their effect on accuracy when using a distributionally stable algorithm.

1 Introduction

In the standard formulation of machine learning problems, the learning algorithm receives training and test samples drawn according to the same distribution. However, this assumption often does not hold in practice. The training sample available is *biased* in some way, which may be due to a variety of practical reasons such as the cost of data labeling or acquisition. The problem occurs in many areas such as astronomy, econometrics, and species habitat modeling.

In a common instance of this problem, points are drawn according to the test distribution but not all of them are made available to the learner. This is called the *sample selection bias problem*. Remarkably, it is often possible to correct this bias by using large amounts of unlabeled data.

The problem of sample selection bias correction for linear regression has been extensively studied in econometrics and statistics (Heckman, 1979; Little & Rubin, 1986) with the pioneering work of Heckman (1979). Several recent machine learning publications (Elkan, 2001; Zadrozny, 2004; Zadrozny et al., 2003; Fan et al., 2005; Dudík et al., 2006) have also dealt with this problem. The main correction technique used in all of these publications consists of reweighting the cost of training point errors to more closely reflect that of the test distribution.

This is in fact a technique commonly used in statistics and machine learning for a variety of problems of this type (Little & Rubin, 1986). With the exact weights, this reweighting could optimally correct the bias, but, in practice, the weights are based on an estimate of the sampling probability from finite data sets. Thus, it is important to determine to what extent the error in this estimation can affect the accuracy of the hypothesis returned by the learning algorithm. To our knowledge, this problem has not been analyzed in a general manner.

This paper gives a theoretical analysis of sample selection bias correction. Our analysis is based on the novel concept of *distributional stability* which generalizes the point-based stability introduced and analyzed by previous authors (Devroye & Wagner, 1979; Kearns & Ron, 1997; Bousquet & Elisseeff, 2002). We show that large families of learning algorithms, including all kernel-based regularization algorithms such as Support Vector Regression (SVR) (Vapnik, 1998) or kernel ridge regression (Saunders et al., 1998) are distributionally stable and we give the expression of their stability coefficient for both the l_1 and l_2 distance.

We then analyze two commonly used sample bias correction techniques: a cluster-based estimation technique and kernel mean matching (KMM) (Huang et al., 2006b). For each of these techniques, we derive bounds on the difference of the error rate of the hypothesis returned by a distributionally stable algorithm when using that estimation technique versus using perfect reweighting. We briefly discuss and compare these bounds and also report the results of experiments with both estimation techniques for several publicly available machine learning data sets. Much of our work and proof techniques can be used to analyze other importance weighting techniques and their effect on accuracy when used in combination with a distributionally stable algorithm.

The remaining sections of this paper are organized as follows. Section 2 describes in detail the sample selection bias correction technique. Section 3 introduces the concept of distributional stability and proves the distributional stability of kernel-based regularization algorithms. Section 4 analyzes the effect of estimation error using distributionally stable algorithms for both the cluster-based and the KMM estimation techniques. Section 5 reports the results of experiments with several data sets comparing these estimation techniques.

2 Sample Selection Bias Correction

2.1 Problem

Let X denote the input space and Y the label set, which may be $\{0, 1\}$ in classification or any measurable subset of \mathbb{R} in regression estimation problems, and let \mathcal{D} denote the *true distribution* over $X \times Y$ according to which test points are drawn. In the sample selection bias problem, some pairs $z = (x, y)$ drawn according to \mathcal{D} are not made available to the learning algorithm. The learning algorithm receives a training sample S of m labeled points z_1, \dots, z_m drawn according to a *biased distribution* \mathcal{D}' over $X \times Y$. This sample bias can be represented by a random variable s taking values in $\{0, 1\}$: when $s = 1$ the

point is sampled, otherwise it is not. Thus, by definition of the sample selection bias, the support of the biased distribution \mathcal{D}' is included in that of the true distribution \mathcal{D} .

As in standard learning scenarios, the objective of the learning algorithm is to select a hypothesis h out of a hypothesis set H with a small generalization error $R(h)$ with respect to the true distribution \mathcal{D} , $R(h) = \mathbb{E}_{(x,y) \sim \mathcal{D}}[c(h, z)]$, where $c(h, z)$ is the cost of the error of h on point $z \in X \times Y$.

While the sample S is collected in some biased manner, it is often possible to derive some information about the nature of the bias. This can be done by exploiting large amounts of unlabeled data drawn according to the true distribution \mathcal{D} , which is often available in practice. Thus, in the following let U be a sample drawn according to \mathcal{D} and $S \subseteq U$ a labeled but biased sub-sample.

2.2 Weighted Samples

A *weighted sample* S_w is a training sample S of m labeled points, z_1, \dots, z_m drawn i.i.d. from $X \times Y$, that is augmented with a non-negative weight $w_i \geq 0$ for each point z_i . This weight is used to emphasize or de-emphasize the cost of an error on z_i as in the so-called *importance weighting* or *cost-sensitive learning* (Elkan, 2001; Zadrozny et al., 2003). One could use the weights w_i to derive an equivalent but larger unweighted sample S' where the multiplicity of z_i would reflect its weight w_i , but most learning algorithms, e.g., decision trees, logistic regression, AdaBoost, Support Vector Machines (SVMs), kernel ridge regression, can directly accept a weighted sample S_w . We will refer to algorithms that can directly take S_w as input as *weight-sensitive algorithms*.

The empirical error of a hypothesis h on a weighted sample S_w is defined as

$$\widehat{R}_w(h) = \sum_{i=1}^m w_i c(h, z_i). \quad (1)$$

Proposition 1. *Let \mathcal{D}' be a distribution whose support coincides with that of \mathcal{D} and let S_w be a weighted sample with $w_i = \Pr_{\mathcal{D}}(z_i) / \Pr_{\mathcal{D}'}(z_i)$ for all points z_i in S . Then,*

$$\mathbb{E}_{S \sim \mathcal{D}'}[\widehat{R}_w(h)] = R(h) = \mathbb{E}_{z \sim \mathcal{D}}[c(h, z)]. \quad (2)$$

Proof. Since the sample points are drawn i.i.d.,

$$\mathbb{E}_{S \sim \mathcal{D}'}[\widehat{R}_w(h)] = \frac{1}{m} \sum_z \mathbb{E}_{S \sim \mathcal{D}'}[w_i c(h, z_i)] = \mathbb{E}_{z_1 \sim \mathcal{D}'}[w_1 c(h, z_1)]. \quad (3)$$

By definition of w and the fact that the support of \mathcal{D} and \mathcal{D}' coincide, the right-hand side can be rewritten as follows

$$\sum_{\mathcal{D}'(z_1) \neq 0} \frac{\Pr_{\mathcal{D}}(z_1)}{\Pr_{\mathcal{D}'}(z_1)} \Pr_{\mathcal{D}'}(z_1) c(h, z_1) = \sum_{\mathcal{D}(z_1) \neq 0} \Pr_{\mathcal{D}}(z_1) c(h, z_1) = \mathbb{E}_{z_1 \sim \mathcal{D}}[c(h, z_1)]. \quad (4)$$

This last term is the definition of the generalization error $R(h)$. \square

2.3 Bias Correction

The probability of drawing $z = (x, y)$ according to the true but unobserved distribution \mathcal{D} can be straightforwardly related to the observed distribution \mathcal{D}' . By definition of the random variable s , the observed biased distribution \mathcal{D}' can be expressed by $\Pr_{\mathcal{D}'}[z] = \Pr_{\mathcal{D}}[z|s = 1]$. We will assume that all points z in the support of D can be sampled with a non-zero probability so the support of \mathcal{D} and \mathcal{D}' coincide. Thus for all $z \in X \times Y$, $\Pr[s = 1|z] \neq 0$. Then, by the Bayes formula, for all z in the support of D ,

$$\Pr_{\mathcal{D}}[z] = \frac{\Pr[z|s = 1] \Pr[s = 1]}{\Pr[s = 1|z]} = \frac{\Pr[s = 1]}{\Pr[s = 1|z]} \Pr_{\mathcal{D}'}[z]. \quad (5)$$

Thus, if we were given the probabilities $\Pr[s = 1]$ and $\Pr[s = 1|z]$, we could derive the true probability $\Pr_{\mathcal{D}}$ from the biased one $\Pr_{\mathcal{D}'}$ exactly and correct the sample selection bias.

It is important to note that this correction is only needed for the training sample S , since it is the only source of selection bias. With a weight-sensitive algorithm, it suffices to reweight each sample z_i with the weight $w_i = \frac{\Pr[s=1]}{\Pr[s=1|z_i]}$. Thus, $\Pr[s = 1|z]$ need not be estimated for all points z but only for those falling in the training sample S . By Proposition 1, the expected value of the empirical error after reweighting is the same as if we were given samples from the true distribution and the usual generalization bounds hold for $\widehat{R}(h)$ and $R(h)$.

When the sampling probability is independent of the labels, as it is commonly assumed in many settings (Zadrozny 2004; 2003), $\Pr[s = 1|z] = \Pr[s = 1|x]$, and Equation 5 can be re-written as

$$\Pr_{\mathcal{D}}[z] = \frac{\Pr[s = 1]}{\Pr[s = 1|x]} \Pr_{\mathcal{D}'}[z]. \quad (6)$$

In that case, the probabilities $\Pr[s = 1]$ and $\Pr[s = 1|x]$ needed to reconstitute $\Pr_{\mathcal{D}}$ from $\Pr_{\mathcal{D}'}$ do not depend on the labels and thus can be estimated using the unlabeled points in U . Moreover, as already mentioned, for weight-sensitive algorithms, it suffices to estimate $\Pr[s = 1|x_i]$ for the points x_i of the training data; no generalization is needed.

A simple case is when the points are defined over a discrete set.³ $\Pr[s = 1|x]$ can then be estimated from the frequency m_x/n_x , where m_x denotes the number of times x appeared in $S \subseteq U$ and n_x the number of times x appeared in the full data set U . $\Pr[s = 1]$ can be estimated by the quantity $|S|/|U|$. However, since $\Pr[s = 1]$ is a constant independent of x , its estimation is not even necessary.

If the estimation of the sampling probability $\Pr[s = 1|x]$ from the unlabeled data set U were exact, then the reweighting just discussed could correct the sample bias optimally. Several techniques have been commonly used to estimate the reweighting quantities. But, these estimate weights are not guaranteed to be exact. The next section addresses how the error in that estimation affects the error rate of the hypothesis returned by the learning algorithm.

³ This can be as a result of a quantization or clustering technique as discussed later.

3 Distributional Stability

Here, we will examine the effect on the error of the hypothesis returned by the learning algorithm in response to a change in the way the training points are weighted. Since the weights are non-negative, we can assume that they are normalized and define a distribution over the training sample. This study can be viewed as a generalization of stability analysis where a single sample point is changed (Devroye & Wagner, 1979; Kearns & Ron, 1997; Bousquet & Elisseeff, 2002) to the more general case of *distributional stability* where the sample’s weight distribution is changed.

Thus, in this section the sample weight \mathcal{W} of $S_{\mathcal{W}}$ defines a distribution over S . For a fixed learning algorithm L and a fixed sample S , we will denote by $h_{\mathcal{W}}$ the hypothesis returned by L for the weighted sample $S_{\mathcal{W}}$. We will denote by $d(\mathcal{W}, \mathcal{W}')$ a divergence measure for two distributions \mathcal{W} and \mathcal{W}' . There are many standard measures for the divergences or distances between two distributions, including the relative entropy, the Hellinger distance, and the l_p distance.

Definition 1 (Distributional β -Stability). *A learning algorithm L is said to be distributionally β -stable for the divergence measure d if for any two weighted samples $S_{\mathcal{W}}$ and $S_{\mathcal{W}'}$,*

$$\forall z \in X \times Y, \quad |c(h_{\mathcal{W}}, z) - c(h_{\mathcal{W}'}, z)| \leq \beta d(\mathcal{W}, \mathcal{W}'). \quad (7)$$

Thus, an algorithm is distributionally stable when small changes to a weighted sample’s distribution, as measured by a divergence d , result in a small change in the cost of an error at any point. The following proposition follows directly from the definition of distributional stability.

Proposition 2. *Let L be a distributionally β -stable algorithm and let $h_{\mathcal{W}}$ ($h_{\mathcal{W}'}$) denote the hypothesis returned by L when trained on the weighted sample $S_{\mathcal{W}}$ (resp. $S_{\mathcal{W}'}$). Let \mathcal{W}_T denote the distribution according to which test points are drawn. Then, the following holds*

$$|R(h_{\mathcal{W}}) - R(h_{\mathcal{W}'})| \leq \beta d(\mathcal{W}, \mathcal{W}'). \quad (8)$$

Proof. By the distributional stability of the algorithm,

$$\mathbb{E}_{z \sim \mathcal{W}_T} [|c(z, h_{\mathcal{W}}) - c(z, h_{\mathcal{W}'})|] \leq \beta d(\mathcal{W}, \mathcal{W}'), \quad (9)$$

which implies the statement of the proposition. \square

3.1 Distributional Stability of Kernel-Based Regularization Algorithms

Here, we show that kernel-based regularization algorithms are distributionally β -stable. This family of algorithms includes, among others, Support Vector Regression (SVR) and kernel ridge regression. Other algorithms such as those based on

the relative entropy regularization can be shown to be distributionally β -stable in a similar way as for point-based stability. Our results also apply to classification algorithms such as Support Vector Machine (SVM) (Cortes & Vapnik, 1995) using a margin-based loss function l_γ as in (Bousquet & Elisseeff, 2002).

We will assume that the cost function c is σ -admissible, that is there exists $\sigma \in \mathbb{R}_+$ such that for any two hypotheses $h, h' \in H$ and for all $z = (x, y) \in X \times Y$,

$$|c(h, z) - c(h', z)| \leq \sigma |h(x) - h'(x)|. \quad (10)$$

This assumption holds for the quadratic cost and most other cost functions when the hypothesis set and the set of output labels are bounded by some $M \in \mathbb{R}_+$: $\forall h \in H, \forall x \in X, |h(x)| \leq M$ and $\forall y \in Y, |y| \leq M$. We will also assume that c is differentiable. This assumption is in fact not necessary and all of our results hold without it, but it makes the presentation simpler.

Let $N: H \rightarrow \mathbb{R}_+$ be a function defined over the hypothesis set. Regularization-based algorithms minimize an objective of the form: $F_{\mathcal{W}}(h) = \widehat{R}_{\mathcal{W}}(h) + \lambda N(h)$, where $\lambda \geq 0$ is a trade-off parameter. We denote by B_F the Bregman divergence associated to a convex function F , $B_F(f||g) = F(f) - F(g) - \langle f - g, \nabla F(g) \rangle$, and define Δh as $\Delta h = h' - h$.

Lemma 1. *Let the hypothesis set H be a vector space. Assume that N is a proper closed convex function and that N is differentiable. Assume that $F_{\mathcal{W}}$ admits a minimizer $h \in H$ and $F_{\mathcal{W}'}$ a minimizer $h' \in H$. Then, the following bound holds,*

$$B_N(h'||h) + B_N(h||h') \leq \frac{\sigma l_1(\mathcal{W}, \mathcal{W}')}{\lambda} \sup_{x \in S} |\Delta h(x)|. \quad (11)$$

Proof. Since $B_{F_{\mathcal{W}}} = B_{\widehat{R}_{\mathcal{W}}} + \lambda B_N$ and $B_{F_{\mathcal{W}'}} = B_{\widehat{R}_{\mathcal{W}'}} + \lambda B_N$, and a Bregman divergence is non-negative, $\lambda(B_N(h'||h) + B_N(h||h')) \leq B_{F_{\mathcal{W}}}(h'||h) + B_{F_{\mathcal{W}'}}(h||h')$. By the definition of h and h' as the minimizers of $F_{\mathcal{W}}$ and $F_{\mathcal{W}'}$,

$$B_{F_{\mathcal{W}}}(h'||h) + B_{F_{\mathcal{W}'}}(h||h') = \widehat{R}_{F_{\mathcal{W}}}(h') - \widehat{R}_{F_{\mathcal{W}}}(h) + \widehat{R}_{F_{\mathcal{W}'}}(h) - \widehat{R}_{F_{\mathcal{W}'}}(h'). \quad (12)$$

Thus, by the σ -admissibility of the cost function c , using the notation $\mathcal{W}_i = \mathcal{W}(x_i)$ and $\mathcal{W}'_i = \mathcal{W}'(x_i)$,

$$\begin{aligned} \lambda(B_N(h'||h) + B_N(h||h')) &\leq \widehat{R}_{F_{\mathcal{W}}}(h') - \widehat{R}_{F_{\mathcal{W}}}(h) + \widehat{R}_{F_{\mathcal{W}'}}(h) - \widehat{R}_{F_{\mathcal{W}'}}(h') \\ &= \sum_{i=1}^m \left[c(h', z_i) \mathcal{W}_i - c(h, z_i) \mathcal{W}_i + c(h, z_i) \mathcal{W}'_i - c(h', z_i) \mathcal{W}'_i \right] \\ &= \sum_{i=1}^m \left[(c(h', z_i) - c(h, z_i)) (\mathcal{W}_i - \mathcal{W}'_i) \right] = \sum_{i=1}^m \left[\sigma |\Delta h(x_i)| (\mathcal{W}_i - \mathcal{W}'_i) \right] \\ &\leq \sigma l_1(\mathcal{W}, \mathcal{W}') \sup_{x \in S} |\Delta h(x)|, \end{aligned} \quad (13)$$

which establishes the lemma. \square

Given $x_1, \dots, x_m \in X$ and a positive definite symmetric (PDS) kernel K , we denote by $\mathbf{K} \in \mathbb{R}^{m \times m}$ the kernel matrix defined by $\mathbf{K}_{ij} = K(x_i, x_j)$ and by $\lambda_{\max}(\mathbf{K}) \in \mathbb{R}_+$ the largest eigenvalue of \mathbf{K} .

Lemma 2. *Let H be a reproducing kernel Hilbert space with kernel K and let the regularization function N be defined by $N(\cdot) = \|\cdot\|_K^2$. Then, the following bound holds,*

$$B_N(h' \| h) + B_N(h \| h') \leq \frac{\sigma \lambda_{\max}^{\frac{1}{2}}(\mathbf{K}) l_2(\mathcal{W}, \mathcal{W}')}{\lambda} \|\Delta h\|_2. \quad (14)$$

Proof. As in the proof of Lemma 1,

$$\lambda(B_N(h' \| h) + B_N(h \| h')) \leq \sum_{i=1}^m \left[(c(h', z_i) - c(h, z_i))(\mathcal{W}_i - \mathcal{W}'_i) \right]. \quad (15)$$

By definition of a reproducing kernel Hilbert space H , for any hypothesis $h \in H$, $\forall x \in X, h(x) = \langle h, K(x, \cdot) \rangle$ and thus also for any $\Delta h = h' - h$ with $h, h' \in H$, $\forall x \in X, \Delta h(x) = \langle \Delta h, K(x, \cdot) \rangle$. Let $\Delta \mathcal{W}_i$ denote $\mathcal{W}'_i - \mathcal{W}_i$, $\Delta \mathcal{W}$ the vector whose components are the $\Delta \mathcal{W}_i$'s, and let V denote $B_N(h' \| h) + B_N(h \| h')$. Using σ -admissibility, $V \leq \sigma \sum_{i=1}^m |\Delta h(x_i) \Delta \mathcal{W}_i| = \sigma \sum_{i=1}^m |\langle \Delta h, \Delta \mathcal{W}_i K(x_i, \cdot) \rangle|$. Let $\epsilon_i \in \{-1, +1\}$ denote the sign of $\langle \Delta h, \Delta \mathcal{W}_i K(x_i, \cdot) \rangle$. Then,

$$\begin{aligned} V &\leq \sigma \left\langle \Delta h, \sum_{i=1}^m \epsilon_i \Delta \mathcal{W}_i K(x_i, \cdot) \right\rangle \leq \sigma \|\Delta h\|_K \left\| \sum_{i=1}^m \epsilon_i \Delta \mathcal{W}_i K(x_i, \cdot) \right\|_K \\ &= \sigma \|\Delta h\|_K \left(\sum_{i,j=1}^m \epsilon_i \epsilon_j \Delta \mathcal{W}_i \Delta \mathcal{W}_j K(x_i, x_j) \right)^{1/2} \\ &= \sigma \|\Delta h\|_K \left[\Delta(\mathcal{W}\epsilon)^\top \mathbf{K} \Delta(\mathcal{W}\epsilon) \right]^{\frac{1}{2}} \leq \sigma \|\Delta h\|_K \|\Delta \mathcal{W}\|_2 \lambda_{\max}^{\frac{1}{2}}(\mathbf{K}). \end{aligned} \quad (16)$$

In this derivation, the second inequality follows from the Cauchy-Schwarz inequality and the last inequality from the standard property of the Rayleigh quotient for PDS matrices. Since $\|\Delta \mathcal{W}\|_2 = l_2(\mathcal{W}, \mathcal{W}')$, this proves the lemma. \square

Theorem 1. *Let K be a kernel such that $K(x, x) \leq \kappa < \infty$ for all $x \in X$. Then, the regularization algorithm based on $N(\cdot) = \|\cdot\|_K^2$ is distributionally β -stable for the l_1 distance with $\beta \leq \frac{\sigma^2 \kappa^2}{2\lambda}$, and for the l_2 distance with $\beta \leq \frac{\sigma^2 \kappa \lambda_{\max}^{\frac{1}{2}}(\mathbf{K})}{2\lambda}$.*

Proof. For $N(\cdot) = \|\cdot\|_K^2$, we have $B_N(h' \| h) = \|h' - h\|_K^2$, thus $B_N(h' \| h) + B_N(h \| h') = 2\|\Delta h\|_K^2$ and by Lemma 1,

$$2\|\Delta h\|_K^2 \leq \frac{\sigma l_1(\mathcal{W}, \mathcal{W}')}{\lambda} \sup_{x \in S} |\Delta h(x)| \leq \frac{\sigma l_1(\mathcal{W}, \mathcal{W}')}{\lambda} \kappa \|\Delta h\|_K. \quad (17)$$

Thus $\|\Delta h\|_K \leq \frac{\sigma \kappa l_1(\mathcal{W}, \mathcal{W}')}{2\lambda}$. By σ -admissibility of c ,

$$\forall z \in X \times Y, |c(h', z) - c(h, z)| \leq \sigma |\Delta h(x)| \leq \kappa \sigma \|\Delta h\|_K. \quad (18)$$

Therefore,

$$\forall z \in X \times Y, |c(h', z) - c(h, z)| \leq \frac{\sigma^2 \kappa^2 l_1(\mathcal{W}, \mathcal{W}')}{2\lambda}, \quad (19)$$

which shows the distributional stability of a kernel-based regularization algorithm for the l_1 distance. Using Lemma 2, a similar derivation leads to

$$\forall z \in X \times Y, |c(h', z) - c(h, z)| \leq \frac{\sigma^2 \kappa \lambda_{\max}^{\frac{1}{2}}(\mathbf{K}) l_2(\mathcal{W}, \mathcal{W}')}{2\lambda}, \quad (20)$$

which shows the distributional stability of a kernel-based regularization algorithm for the l_2 distance. \square

Note that the standard setting of a sample with no weight is equivalent to a weighted sample with the uniform distribution \mathcal{W}_U : each point is assigned the weight $1/m$. Removing a single point, say x_1 , is equivalent to assigning weight 0 to x_1 and $1/(m-1)$ to others. Let $\mathcal{W}_{U'}$ be the corresponding distribution, then $l_1(\mathcal{W}_U, \mathcal{W}_{U'}) = \frac{1}{m} + \sum_{i=1}^{m-1} \left| \frac{1}{m} - \frac{1}{m-1} \right| = \frac{2}{m}$. Thus, in the case of kernel-based regularized algorithms and for the l_1 distance, standard uniform β -stability is a special case of distributional β -stability. It can be shown similarly that $l_2(\mathcal{W}_U, \mathcal{W}_{U'}) = \frac{1}{\sqrt{m(m-1)}}$.

4 Effect of Estimation Error for Kernel-Based Regularization Algorithms

This section analyzes the effect of an error in the estimation of the weight of a training example on the generalization error of the hypothesis h returned by a weight-sensitive learning algorithm. We will examine two estimation techniques: a straightforward histogram-based or cluster-based method, and kernel mean matching (KMM) (Huang et al., 2006b).

4.1 Cluster-Based Estimation

A straightforward estimate of the probability of sampling is based on the observed empirical frequencies. The ratio of the number of times a point x appears in S and the number of times it appears in U is an empirical estimate of $\Pr[s = 1|x]$. Note that generalization to unseen points x is not needed since reweighting requires only assigning weights to the seen training points. However, in general, training instances are typically unique or very infrequent since features are real-valued numbers. Instead, features can be discretized based on a partitioning of the input space X . The partitioning may be based on a simple histogram buckets or the result of a clustering technique. The analysis of this section assumes such a prior partitioning of X .

We shall analyze how fast the resulting empirical frequencies converge to the true sampling probability. For $x \in U$, let U_x denote the subsample of U containing exactly all the instances of x and let $n = |U|$ and $n_x = |U_x|$. Furthermore,

let n' denote the number of unique points in the sample U . Similarly, we define S_x , m , m_x and m' for the set S . Additionally, denote by $p_0 = \min_{x \in U} \Pr[x] \neq 0$.

Lemma 3. *Let $\delta > 0$. Then, with probability at least $1 - \delta$, the following inequality holds for all x in S :*

$$\left| \Pr[s = 1|x] - \frac{m_x}{n_x} \right| \leq \sqrt{\frac{\log 2m' + \log \frac{1}{\delta}}{p_0 n}}. \quad (21)$$

Proof. For a fixed $x \in U$, by Hoeffding's inequality,

$$\begin{aligned} \Pr_U \left[\left| \Pr[s = 1|x] - \frac{m_x}{n_x} \right| \geq \epsilon \right] &= \sum_{i=1}^n \Pr_x \left[\left| \Pr[s = 1|x] - \frac{m_x}{i} \right| \geq \epsilon \mid n_x = i \right] \Pr[n_x = i] \\ &\leq \sum_{i=1}^n 2e^{-2i\epsilon^2} \Pr_U[n_x = i]. \end{aligned}$$

Since n_x is a binomial random variable with parameters $\Pr_U[x] = p_x$ and n , this last term can be expressed more explicitly and bounded as follows:

$$\begin{aligned} 2 \sum_{i=1}^n e^{-2i\epsilon^2} \Pr_U[n_x = i] &\leq 2 \sum_{i=0}^n e^{-2i\epsilon^2} \binom{n}{i} p_x^i (1-p_x)^{n-i} = 2(p_x e^{-2\epsilon^2} + (1-p_x))^n \\ &= 2(1-p_x(1-e^{-2\epsilon^2}))^n \leq 2 \exp(-p_x n(1-e^{-2\epsilon^2})). \end{aligned}$$

Since for $x \in [0, 1]$, $1 - e^{-x} \geq x/2$, this shows that for $\epsilon \in [0, 1]$,

$$\Pr_U \left[\left| \Pr[s = 1|x] - \frac{m_x}{n_x} \right| \geq \epsilon \right] \leq 2e^{-p_x n \epsilon^2}.$$

By the union bound and the definition of p_0 ,

$$\Pr_U \left[\exists x \in S : \left| \Pr[s = 1|x] - \frac{m_x}{n_x} \right| \geq \epsilon \right] \leq 2m' e^{-p_0 n \epsilon^2}.$$

Setting δ to match the upper bound yields the statement of the lemma. \square

The following proposition bounds the distance between the distribution \mathcal{W} corresponding to a perfectly reweighted sample ($S_{\mathcal{W}}$) and the one corresponding to a sample that is reweighted according to the observed bias ($S_{\widehat{\mathcal{W}}}$). For a sampled point $x_i = x$, these distributions are defined as follows:

$$\mathcal{W}(x_i) = \frac{1}{m} \frac{1}{p(x_i)} \quad \text{and} \quad \widehat{\mathcal{W}}(x_i) = \frac{1}{m} \frac{1}{\hat{p}(x_i)}, \quad (22)$$

where, for a *distinct* point x equal to the *sampled* point x_i , we define $p(x_i) = \Pr[s = 1|x]$ and $\hat{p}(x_i) = \frac{m_x}{n_x}$.

Proposition 3. *Let $B = \max_{i=1, \dots, m} \max(1/p(x_i), 1/\hat{p}(x_i))$. Then, the l_1 and l_2 distances of the distributions \mathcal{W} and $\widehat{\mathcal{W}}$ can be bounded as follows,*

$$l_1(\mathcal{W}, \widehat{\mathcal{W}}) \leq B^2 \sqrt{\frac{\log 2m' + \log \frac{1}{\delta}}{p_0 n}} \quad \text{and} \quad l_2(\mathcal{W}, \widehat{\mathcal{W}}) \leq B^2 \sqrt{\frac{\log 2m' + \log \frac{1}{\delta}}{p_0 n m}}. \quad (23)$$

Proof. By definition of the l_2 distance,

$$\begin{aligned} l_2^2(\mathcal{W}, \widehat{\mathcal{W}}) &= \frac{1}{m^2} \sum_{i=1}^m \left(\frac{1}{p(x_i)} - \frac{1}{\hat{p}(x_i)} \right)^2 = \frac{1}{m^2} \sum_{i=1}^m \left(\frac{p(x_i) - \hat{p}(x_i)}{p(x_i)\hat{p}(x_i)} \right)^2 \\ &\leq \frac{B^4}{m} \max_i (p(x_i) - \hat{p}(x_i))^2. \end{aligned}$$

It can be shown similarly that $l_1(\mathcal{W}, \widehat{\mathcal{W}}) \leq B^2 \max_i |p(x_i) - \hat{p}(x_i)|$. The application of the uniform convergence bound of Lemma 3 directly yields the statement of the proposition. \square

The following theorem provides a bound on the difference between the generalization error of the hypothesis returned by a kernel-based regularization algorithm when trained on the perfectly unbiased distribution, and the one trained on the sample bias-corrected using frequency estimates.

Theorem 2. *Let K be a PDS kernel such that $K(x, x) \leq \kappa < \infty$ for all $x \in X$. Let $h_{\mathcal{W}}$ be the hypothesis returned by the regularization algorithm based on $N(\cdot) = \|\cdot\|_K^2$ using $S_{\mathcal{W}}$, and $h_{\widehat{\mathcal{W}}}$ the one returned after training the same algorithm on $S_{\widehat{\mathcal{W}}}$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the difference in generalization error of these hypotheses is bounded as follows*

$$\begin{aligned} |R(h_{\mathcal{W}}) - R(h_{\widehat{\mathcal{W}}})| &\leq \frac{\sigma^2 \kappa^2 B^2}{2\lambda} \sqrt{\frac{\log 2m' + \log \frac{1}{\delta}}{p_0 n}} \\ |R(h_{\mathcal{W}}) - R(h_{\widehat{\mathcal{W}}})| &\leq \frac{\sigma^2 \kappa \lambda_{\max}^{\frac{1}{2}}(\mathbf{K}) B^2}{2\lambda} \sqrt{\frac{\log 2m' + \log \frac{1}{\delta}}{p_0 n m}}. \end{aligned} \tag{24}$$

Proof. The result follows from Proposition 2, the distributional stability and the bounds on the stability coefficient β for kernel-based regularization algorithms (Theorem 1), and the bounds on the l_1 and l_2 distances between the correct distribution \mathcal{W} and the estimate $\widehat{\mathcal{W}}$. \square

Let n_0 be the number of occurrences, in U , of the least frequent training example. For large enough n , $p_0 n \approx n_0$, thus the theorem suggests that the difference of error rate between the hypothesis returned after an optimal reweighting versus the one based on frequency estimates goes to zero as $\sqrt{\frac{\log m'}{n_0}}$. In practice, $m' \leq m$, the number of distinct points in S is small, a fortiori, $\log m'$ is very small, thus, the convergence rate depends essentially on the rate at which n_0 increases. Additionally, if $\lambda_{\max}(K) \leq m$ (such as with Gaussian kernels), the l_2 -based bound will provide convergence that is at least as fast.

4.2 Kernel Mean Matching

The following definitions introduced by Steinwart (2002) will be needed for the presentation and discussion of the kernel mean matching (KMM) technique. Let

X be a compact metric space and let $C(X)$ denote the space of all continuous functions over X equipped with the standard infinite norm $\|\cdot\|_\infty$. Let $K: X \times X \rightarrow \mathbb{R}$ be a PDS kernel. There exists a Hilbert space F and a map $\Phi: X \rightarrow F$ such that for all $x, y \in X$, $K(x, y) = \langle \Phi(x), \Phi(y) \rangle$. Note that for a given kernel K , F and Φ are not unique and that, for these definitions, F does not need to be a reproducing kernel Hilbert space (RKHS).

Let \mathcal{P} denote the set of all probability distributions over X and let $\mu: \mathcal{P} \rightarrow F$ be the function defined by $\forall p \in \mathcal{P}$, $\mu(p) = \mathbb{E}_{x \sim p}[\Phi(x)]$. A function $g: X \rightarrow \mathbb{R}$ is said to be *induced* by K if there exists $w \in F$ such that for all $x \in X$, $g(x) = \langle w, \Phi(x) \rangle$. K is said to be *universal* if it is continuous and if the set of functions induced by K are dense in $C(X)$.

Theorem 3 (Huang et al. (2006a)). *Let F be a separable Hilbert space and let K be a universal kernel with feature space F and feature map $\Phi: X \rightarrow F$. Then, μ is injective.*

Proof. We give a full proof of the main theorem supporting this technique in a longer version of this paper. The proof given by Huang et al. (2006a) does not seem to be complete. \square

The KMM technique is applicable when the learning algorithm is based on a universal kernel. The theorem shows that for a universal kernel, the expected value of the feature vectors induced uniquely determines the probability distribution. KMM uses this property to reweight training points so that the average value of the feature vectors for the training data matches that of the feature vectors for a set of unlabeled points drawn from the unbiased distribution.

Let γ_i denote the perfect reweighting of the sample point x_i and $\hat{\gamma}_i$ the estimate derived by KMM. Let B' denote the largest possible reweighting coefficient γ and let ϵ be a positive real number. We will assume that ϵ is chosen so that $\epsilon \leq 1/2$. Then, the following is the KMM constraint optimization

$$\min_{\gamma} G(\gamma) = \left\| \frac{1}{m} \sum_{i=1}^m \gamma_i \Phi(x_i) - \frac{1}{n} \sum_{i=1}^n \Phi(x'_i) \right\| \quad \text{s.t. } \gamma_i \in [0, B'] \wedge \left| \frac{1}{m} \sum_{i=1}^m \gamma_i - 1 \right| \leq \epsilon.$$

Let $\hat{\gamma}$ be the solution of this optimization problem, then $\frac{1}{m} \sum_{i=1}^m \hat{\gamma}_i = 1 + \epsilon'$ with $-\epsilon \leq \epsilon' \leq \epsilon$. For $i \in [1, m]$, let $\hat{\gamma}'_i = \hat{\gamma}_i / (1 + \epsilon')$. The normalized weights used in KMM's reweighting of the sample are thus defined by $\hat{\gamma}'_i / m$ with $\frac{1}{m} \sum_{i=1}^m \hat{\gamma}'_i = 1$.

As in the previous section, given $x_1, \dots, x_m \in X$ and a strictly positive definite universal kernel K , we denote by $\mathbf{K} \in \mathbb{R}^{m \times m}$ the kernel matrix defined by $\mathbf{K}_{ij} = K(x_i, x_j)$ and by $\lambda_{\min}(\mathbf{K}) > 0$ the smallest eigenvalue of \mathbf{K} . We also denote by $\text{cond}(\mathbf{K})$ the condition number of the matrix \mathbf{K} : $\text{cond}(\mathbf{K}) = \lambda_{\max}(\mathbf{K}) / \lambda_{\min}(\mathbf{K})$. When K is universal, it is continuous over the compact $X \times X$ and thus bounded, and there exists $\kappa < \infty$ such that $K(x, x) \leq \kappa$ for all $x \in X$.

Proposition 4. *Let K be a strictly positive definite universal kernel. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the l_2 distance of the distributions $\hat{\gamma}'/m$*

and γ/m is bounded as follows:

$$\frac{1}{m} \|(\hat{\gamma}' - \gamma)\|_2 \leq \frac{2\epsilon B'}{\sqrt{m}} + \frac{2\kappa^{\frac{1}{2}}}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})} \sqrt{\frac{B'^2}{m} + \frac{1}{n}} \left(1 + \sqrt{2 \log \frac{2}{\delta}}\right). \quad (25)$$

Proof. Since the optimal reweighting γ verifies the constraints of the optimization, by definition of $\hat{\gamma}$ as a minimizer, $G(\hat{\gamma}) \leq G(\gamma)$. Thus, by the triangle inequality,

$$\left\| \frac{1}{m} \sum_{i=1}^m \hat{\gamma}_i \Phi(x_i) - \frac{1}{m} \sum_{i=1}^m \gamma_i \Phi(x_i) \right\| \leq G(\hat{\gamma}) + G(\gamma) \leq 2G(\gamma). \quad (26)$$

Let L denote the left-hand side of this inequality: $L = \frac{1}{m} \left\| \sum_{i=1}^m (\hat{\gamma}_i - \gamma_i) \Phi(x_i) \right\|$. By definition of the norm in the Hilbert space, $L = \frac{1}{m} \sqrt{(\hat{\gamma} - \gamma)^\top \mathbf{K} (\hat{\gamma} - \gamma)}$. Then, by the standard bounds for the Rayleigh quotient of PDS matrices, $L \geq \frac{1}{m} \lambda_{\min}^{\frac{1}{2}}(\mathbf{K}) \|(\hat{\gamma} - \gamma)\|_2$. This combined with Inequality 26 yields

$$\frac{1}{m} \|(\hat{\gamma} - \gamma)\|_2 \leq \frac{2G(\gamma)}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})}. \quad (27)$$

Thus, by the triangle inequality,

$$\begin{aligned} \frac{1}{m} \|(\hat{\gamma}' - \gamma)\|_2 &\leq \frac{1}{m} \|(\hat{\gamma}' - \hat{\gamma})\|_2 + \frac{1}{m} \|(\hat{\gamma} - \gamma)\|_2 \leq \frac{|\epsilon'|/m}{1 + \epsilon'} \|\gamma\|_2 + \frac{2G(\gamma)}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})} \\ &\leq \frac{2|\epsilon'|B'\sqrt{m}}{m} + \frac{2G(\gamma)}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})} \leq \frac{2\epsilon B'}{\sqrt{m}} + \frac{2G(\gamma)}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})}. \end{aligned} \quad (28)$$

It is not difficult to show using McDiarmid's inequality that for any $\delta > 0$, with probability at least $1 - \delta$, the following holds (Lemma 4, (Huang et al., 2006a)):

$$G(\gamma) \leq \kappa^{\frac{1}{2}} \sqrt{\frac{B'^2}{m} + \frac{1}{n}} \left(1 + \sqrt{2 \log \frac{2}{\delta}}\right). \quad (29)$$

This combined with Inequality 28 yields the statement of the proposition. \square

The following theorem provides a bound on the difference between the generalization error of the hypothesis returned by a kernel-based regularization algorithm when trained on the true distribution, and the one trained on the sample bias-corrected KMM.

Theorem 4. *Let K be a strictly positive definite symmetric universal kernel. Let h_γ be the hypothesis returned by the regularization algorithm based on $N(\cdot) = \|\cdot\|_K^2$ using $S_{\gamma/m}$ and $h_{\hat{\gamma}'}$ the one returned after training the same algorithm on $S_{\hat{\gamma}'/m}$. Then, for any $\delta > 0$, with probability at least $1 - \delta$, the difference in generalization error of these hypotheses is bounded as follows*

$$|R(h_\gamma) - R(h_{\hat{\gamma}'})| \leq \frac{\sigma^2 \kappa \lambda_{\max}^{\frac{1}{2}}(\mathbf{K})}{\lambda} \left(\frac{\epsilon B'}{\sqrt{m}} + \frac{\kappa^{\frac{1}{2}}}{\lambda_{\min}^{\frac{1}{2}}(\mathbf{K})} \sqrt{\frac{B'^2}{m} + \frac{1}{n}} \left(1 + \sqrt{2 \log \frac{2}{\delta}}\right) \right).$$

For $\epsilon = 0$, the bound becomes

$$|R(h_\gamma) - R(h_{\hat{\gamma}})| \leq \frac{\sigma^2 \kappa^{\frac{3}{2}} \text{cond}^{\frac{1}{2}}(\mathbf{K})}{\lambda} \sqrt{\frac{B'^2}{m} + \frac{1}{n}} \left(1 + \sqrt{2 \log \frac{2}{\delta}}\right). \quad (30)$$

Proof. The result follows from Proposition 2 and the bound of Proposition 4. \square

Comparing this bound for $\epsilon = 0$ with the l_2 bound of Theorem 4, we first note that B and B' are essentially related modulo the constant $\Pr[s = 1]$ which is not included in the cluster-based reweighting. Thus, the cluster-based convergence is of the order $O(\lambda_{\max}^{\frac{1}{2}}(\mathbf{K})B^2\sqrt{\frac{\log m'}{p_0nm}})$ and the KMM convergence of the order $O(\text{cond}^{\frac{1}{2}}(\mathbf{K})\frac{B}{\sqrt{m}})$. Taking the ratio of the former over the latter and noticing $p_0^{-1} \approx O(B)$, we obtain the expression $O\left(\sqrt{\frac{\lambda_{\min}(\mathbf{K})B \log m'}{n}}\right)$. Thus, for $n > \lambda_{\min}(\mathbf{K})B \log(m')$ the convergence of the cluster-based bound is more favorable, while for other values the KMM bound converges faster.

5 Experimental Results

In this section, we will compare the performance of the cluster-based reweighting technique and the KMM technique empirically. We will first discuss and analyze the properties of the clustering method and our particular implementation.

The analysis of Section 4.1 deals with discrete points possibly resulting from the use of a quantization or clustering technique. However, due to the relatively small size of the public training sets available, clustering could leave us with few cluster representatives to train with. Instead, in our experiments, we only used the clusters to estimate sampling probabilities and applied these weights to the full set of training points. As the following proposition shows, the l_1 and l_2 distance bounds of Proposition 5 do not change significantly so long as the cluster size is roughly uniform and the sampling probability is the same for all points within a cluster. We will refer to this as the *clustering assumption*. In what follows, let $\Pr[s = 1 | C_i]$ designate the sampling probability for all $x \in C_i$. Finally, define $q(C_i) = \Pr[s = 1 | C_i]$ and $\hat{q}(C_i) = |C_i \cap S| / |C_i \cap U|$.

Proposition 5. *Let $B = \max_{i=1, \dots, m} \max(1/q(C_i), 1/\hat{q}(C_i))$. Then, the l_1 and l_2 distances of the distributions \mathcal{W} and $\widehat{\mathcal{W}}$ can be bounded as follows,*

$$l_1(\mathcal{W}, \widehat{\mathcal{W}}) \leq B^2 \sqrt{\frac{|C_M|k(\log 2k + \log \frac{1}{\delta})}{q_0nm}} \quad l_2(\mathcal{W}, \widehat{\mathcal{W}}) \leq B^2 \sqrt{\frac{|C_M|k(\log 2k + \log \frac{1}{\delta})}{q_0nm^2}},$$

where $q_0 = \min q(C_i)$ and $|C_M| = \max_i |C_i|$.

Proof. By definition of the l_2 distance,

$$\begin{aligned} l_2^2(\mathcal{W}, \widehat{\mathcal{W}}) &= \frac{1}{m^2} \sum_{i=1}^k \sum_{x \in C_i} \left(\frac{1}{p(x)} - \frac{1}{\hat{p}(x)} \right)^2 = \frac{1}{m^2} \sum_{i=1}^k \sum_{x \in C_i} \left(\frac{1}{q(C_i)} - \frac{1}{\hat{q}(C_i)} \right)^2 \\ &\leq \frac{B^4 |C_M|}{m^2} \sum_{i=1}^k \max_i (q(C_i) - \hat{q}(C_i))^2. \end{aligned}$$

Table 1. Normalized mean-squared error (NMSE) for various regression data sets using unweighted, ideal, clustered and kernel-mean-matched training sample reweightings.

DATA SET	$ U $	$ S $	n_{test}	UNWEIGHTED	IDEAL	CLUSTERED	KMM
ABALONE	2000	724	2177	.654±.019	.551±.032	.623±.034	.709±.122
BANK32NH	4500	2384	3693	.903±.022	.610±.044	.635±.046	.691±.055
BANK8FM	4499	1998	3693	.085±.003	.058±.001	.068±.002	.079±.013
CAL-HOUSING	16512	9511	4128	.395±.010	.360±.009	.375±.010	.595±.054
CPU-ACT	4000	2400	4192	.673±.014	.523±.080	.568±.018	.518±.237
CPU-SMALL	4000	2368	4192	.682±.053	.477±.097	.408±.071	.531±.280
HOUSING	300	116	206	.509±.049	.390±.053	.482±.042	.469±.148
KIN8NM	5000	2510	3192	.594±.008	.523±.045	.574±.018	.704±.068
PUMA8NH	4499	2246	3693	.685±.013	.674±.019	.641±.012	.903±.059

The right-hand side of the first line follows from the clustering assumption and the inequality then follows from exactly the same steps as in Proposition 5 and factoring away the sum over the elements of C_i . Finally, it is easy to see that the $\max_i(q(C_i) - \hat{q}(C_i))$ term can be bounded just as in Lemma 3 using a uniform convergence bound, however now the union bound is taken over the clusters rather than unique points. \square

Note that when the cluster size is uniform, then $|C_M|k = m$, and the bound above leads to an expression similar to that of Proposition 5.

We used the leaves of a decision tree to define the clusters. A decision tree selects binary cuts on the coordinates of $x \in X$ that greedily minimize a node impurity measure, e.g., MSE for regression (Breiman et al., 1984). Points with similar features and labels are clustered together in this way with the assumption that these will also have similar sampling probabilities.

Several methods for bias correction are compared in Table 1. Each method assigns corrective weights to the training samples. The *unweighted* method uses weight 1 for every training instance. The *ideal* method uses weight $\frac{1}{\Pr[s=1|x]}$, which is optimal but requires the sampling distribution to be known. The *clustered* method uses weight $|C_i \cap U|/|C_i \cap S|$, where the clusters C_i are regression tree leaves with a minimum count of 4 (larger cluster sizes showed similar, though declining, performance). The KMM method uses the approach of Huang et al. (2006b) with a Gaussian kernel and parameters $\sigma = \sqrt{d}/2$ for $x \in \mathbb{R}^d$, $B = 1000$, $\epsilon = 0$. Note that we know of no principled way to do cross-validation with KMM since it cannot produce weights for a held-out set (Sugiyama et al., 2008).

The regression datasets are from LIAAD⁴ and are sampled with $P[s = 1|x] = \frac{e^v}{1+e^v}$ where $v = \frac{4w \cdot (x - \bar{x})}{\sigma_{w \cdot (x - \bar{x})}}$, $x \in \mathbb{R}^d$ and $w \in \mathbb{R}^d$ chosen at random from $[-1, 1]^d$. In our experiments, we chose ten random projections w and reported results with the w , for each data set, that maximizes the difference between the unweighted and ideal methods over repeated sampling trials. In this way, we selected bias samplings that are good candidates for bias correction estimation.

⁴ www.liaad.up.pt/~ltorgo/Regression/DataSets.html.

For our experiments, we used a version of SVR available from LibSVM⁵ that can take as input weighted samples, with parameter values $C = 1$, and $\epsilon = 0.1$ combined with a Gaussian kernel with parameter $\sigma = \sqrt{d/2}$. We report results using normalized mean-squared error (NMSE): $\frac{1}{n_{test}} \sum_{i=1}^{n_{test}} \frac{(y_i - \hat{y}_i)^2}{\sigma_y^2}$, and provide mean and standard deviations for ten-fold cross-validation.

Our results show that reweighting with more reliable counts, due to clustering, can be effective in the problem of sample bias correction. These results also confirm the dependence that our theoretical bounds exhibit on the quantity n_0 . The results obtained using KMM seem to be consistent with those reported by the authors of this technique.⁶

6 Conclusion

We presented a general analysis of sample selection bias correction and gave bounds analyzing the effect of an estimation error on the accuracy of the hypotheses returned. The notion of distributional stability and the techniques presented are general and can be of independent interest for the analysis of learning algorithms in other settings. In particular, these techniques apply similarly to other importance weighting algorithms and can be used in other contexts such that of learning in the presence of uncertain labels. The analysis of the discriminative method of (Bickel et al., 2007) for the problem of covariate shift could perhaps also benefit from this study.

7 Acknowledgements

This work was partially funded by the New York State Office of Science Technology and Academic Research (NYSTAR) and was also sponsored in part by the Department of the Army Award Number W23RYX-3275-N605. The U.S. Army Medical Research Acquisition Activity, 820 Chandler Street, Fort Detrick MD 21702-5014 is the awarding and administering acquisition office. The content of this material does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

References

- Bickel, S., Brückner, M., & Scheffer, T. (2007). Discriminative learning for differing training and test distributions. *ICML 2007* (pp. 81–88).
- Bousquet, O., & Elisseeff, A. (2002). Stability and generalization. *JMLR*, 2, 499–526.

⁵ www.csie.ntu.edu.tw/~cjlin/libsvmtools.

⁶ We thank Arthur Gretton for discussion and help in clarifying the choice of the parameters and design of the KMM experiments reported in (Huang et al., 2006b), and for providing the code used by the authors for comparison studies.

- Breiman, L., Friedman, J., Olshen, R., & Stone, C. (1984). *Classification and regression trees*. Boca Raton, FL: CRC Press.
- Cortes, C., & Vapnik, V. N. (1995). Support-Vector Networks. *Machine Learning*, 20, 273–297.
- Devroye, L., & Wagner, T. (1979). Distribution-free performance bounds for potential function rules. *IEEE Trans. on Information Theory* (pp. 601–604).
- Dudík, M., Schapire, R. E., & Phillips, S. J. (2006). Correcting sample selection bias in maximum entropy density estimation. *NIPS 2005*.
- Elkan, C. (2001). The foundations of cost-sensitive learning. *IJCAI* (pp. 973–978).
- Fan, W., Davidson, I., Zadrozny, B., & Yu, P. S. (2005). An improved categorization of classifier’s sensitivity on sample selection bias. *ICDM 2005* (pp. 605–608). IEEE Computer Society.
- Heckman, J. J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47, 151–161.
- Huang, J., Smola, A., Gretton, A., Borgwardt, K., & Schölkopf, B. (2006a). *Correcting Sample Selection Bias by Unlabeled Data*. Technical Report CS-2006-44). University of Waterloo.
- Huang, J., Smola, A. J., Gretton, A., Borgwardt, K. M., & Schölkopf, B. (2006b). Correcting sample selection bias by unlabeled data. *NIPS 2006* (pp. 601–608).
- Kearns, M., & Ron, D. (1997). Algorithmic stability and sanity-check bounds for leave-one-out cross-validation. *COLT 1997* (pp. 152–162).
- Little, R. J. A., & Rubin, D. B. (1986). *Statistical analysis with missing data*. New York, NY, USA: John Wiley & Sons, Inc.
- Saunders, C., Gammernan, A., & Vovk, V. (1998). Ridge Regression Learning Algorithm in Dual Variables. *ICML 1998* (pp. 515–521).
- Steinwart, I. (2002). On the influence of the kernel on the consistency of support vector machines. *JMLR*, 2, 67–93.
- Sugiyama, M., Nakajima, S., Kashima, H., von Bünau, P., & Kawanabe, M. (2008). Direct importance estimation with model selection and its application to covariate shift adaptation. *NIPS 2008*.
- Vapnik, V. N. (1998). *Statistical learning theory*. New York: Wiley-Interscience.
- Zadrozny, B. (2004). Learning and evaluating classifiers under sample selection bias. *ICML 2004*.
- Zadrozny, B., Langford, J., & Abe, N. (2003). Cost-sensitive learning by cost-proportionate example weighting. *ICDM 2003*.