
Correcting a Class of Complete Selection Bias with External Data based on Importance Weight Estimation

Van-Tinh Tran¹, Alex Aussem¹

LIRIS, UMR 5205, University of Lyon 1, 69622 Lyon, France
{van-tinh.tran, aaussem}@univ-lyon1.fr

ABSTRACT. We present a practical bias correction method for classifier and regression models learning under a general class of selection bias. The method hinges on two assumptions: 1) a feature vector, X_s , exists such that S , the variable that controls the inclusion of the samples in the training set, is conditionally independent of (X, Y) given X_s ; 2) one has access to some external samples drawn from the population as a whole in order to approximate the unbiased distribution of X_s . This general framework includes covariate shift and prior probability shift as special cases. We first show how importance weighting can remove this bias. We also discuss the case where our key assumption about X_s is not valid and where X_s is only partially observed in the test set. Experimental results on synthetic and real-world data demonstrate that our method works well in practice.

RÉSUMÉ. Nous présentons une méthode pour corriger une classe générale de biais de sélection dans le cadre de la classification et de la régression. La méthode repose sur deux hypothèses: 1) il existe un vecteurs X_s , tel que S , la variable qui contrôle l'inclusion des individus dans la base d'apprentissage, est indépendante de (X, Y) conditionnellement à X_s ; 2) on dispose d'un échantillon supplémentaire non biaisé permettant d'estimer la distribution de X_s . Le "covariate shift" et le "prior probability shift" sont des cas particuliers de ce cadre. Nous montrons d'abord comment il est possible d'éliminer ce biais avec la technique "importance weighting". Nous discutons ensuite du cas où notre hypothèse d'indépendance sur X_s est invalide et le cas où X_s n'est que partiellement observée. sur la base de test. La pertinence de cette approche est démontrée sur des jeux de données synthétiques et réelles.

KEYWORDS: selection bias, importance weighting, graphical models.

MOTS-CLÉS: biais de sélection, importance weighting, modèles graphiques.

1. Introduction

Selection bias, which occurs when training and test joint distributions are different, i.e. $P_{tr}(x, y) \neq P_{te}(x, y)$, is pervasive in almost all empirical studies, including

Machine Learning, Statistics, Social Sciences, Economics, Bioinformatics, Biostatistics, Epidemiology, Medicine, etc. It is therefore highly desirable to devise algorithms that remain effective under such distribution shifts. In general, the estimation problem with two different distributions $P_{tr}(x, y)$ and $P_{te}(x, y)$ is unsolvable, as the two terms could be arbitrarily far apart. However, when $P_{tr}(x, y)$ and $P_{te}(x, y)$ differ only in $P_{tr}(x)$ and $P_{te}(x)$ (known as *covariate shift*) or only in $P_{tr}(y)$ and $P_{te}(y)$ (known as *prior probability shift*), effective adaptation is possible. In this paper, we present a practical correction method for classifier and regression model learning under a more general class of selection bias. We assume implicitly that there exists a joint probability distribution $P(x, y, s)$ that satisfies: $P_{te}(x, y) = P(x, y) = \sum_s P(x, y, s)$ and $P_{tr}(x, y) = P(x, y | s = 1)$, where the variable S controls the selection of examples in the training set (1 means the example is selected, 0 means the example is not selected). While $P_{te}(x, y)$ and $P_{tr}(x, y)$ are derived from the same distribution $P(x, y, s)$, we assume no independence assumption holds between X , Y , and S . This is termed *complete selection bias* in the literature. In this case, we have to resort to some additional information on the mechanism by which the samples were preferentially selected to the data set to correct the bias.

The recent paper by Bareinboim et al. (Bareinboim *et al.*, 2014) has been very influential in our thinking. Mirroring their work, we show that, if we have a combination of biased data and unbiased data and qualitative probabilistic assumptions that are deemed plausible about our sampling mechanism, our problem becomes solvable. More specifically, we assume we have access to a S -control feature vector, X_s , and some additional sample of the form (x_s) that is drawn from the population as a whole, such that S is conditionally independent of (X, Y) given X_s . Despite being limited to specific or idealized situations, this framework includes covariate shift and prior probability shift as special cases. We also consider the case where X_s is not fully measured in the target population. This situation typically arises in various clinical studies or epidemiological scenarios, where some variables are too difficult or costly to measure in the target population.

We show that one may account for the difference between $P_{tr}(x, y)$ and $P_{te}(x, y)$ by reweighting the training points using the so-called importance weight, denoted as $\beta(x_s)$. If the selection process is explicitly known, then $\beta(x_s)$ is simply given by $\frac{P(s=1)}{P(s=1|x_s)}$, otherwise, we resort to an external (bias-free) data set of X_s in order to estimate $\beta(x_s)$ directly (Huang *et al.*, 2006; Kanamori *et al.*, 2009). Furthermore, as one usually has a partial understanding of the sampling mechanism, we investigate whether covariate shift and prior probability shift corrections may help reduce complete selection bias despite not being valid.

2. Bias correction

In this section, we investigate the interplay between two types of variables, V_B and V_P , where V_B are variables collected under selection bias, $P(V_B | S = 1)$, and V_P are

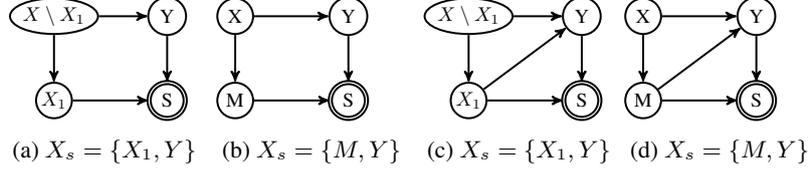


Figure 1: Examples of complete selection bias mechanisms depicted graphically. The S -control vector is shown along each plot. X_1 is some input variable to the model, M denotes a variables that is not observed in the target domain, and thus not used a input variable.

variables collected in the population-level, $P(V_P)$ ¹. We assume that $Y \in V_B$. In (Bareinboim *et al.*, 2014), Bareinbom *et al.* provide a sufficient condition for $P(x, y)$ to be recoverable when no data is gathered over X and Y in the population level. We extend slightly their result by considering also the case where either Y or some input variables in X are not only collected under selection bias, but also in the population-level, (i.e. $\{X, Y\} \cap V_P \cap V_B \neq \emptyset$),

THEOREM 1. — *The bias-free distribution $P(x, y)$ is recoverable from a S -bias training samples if there exists a set of variables $X_s \subseteq V_B \cap V_P$, such that $S \perp\!\!\!\perp (X, Y) | X_s$ and the support of $P(x_s | s = 1)$ contains the support of $P(x_s)$. Let $\beta(x_s) = \frac{P(s=1)}{P(s=1|x_s)}$, $P(x, y)$ is then given by the formula,*

$$P(x, y) = \sum_{x_s \setminus \{x, y\}} P(x, y, x_s | s = 1) \beta(x_s) \quad (1)$$

$\beta(x_s)$ can be reformulated as $\beta(x_s) = \frac{P(x_s)}{P(x_s | s=1)}$. So $\beta(x_s)$ may be estimated from a combination of biased and unbiased data. Theorem 1 relies on qualitative assumptions (X_s controls S over (X, Y)) that may appear difficult to satisfy in practice. However, in certain domains like epidemiology, information about the selection process can sometimes be expressed and modeled in a communicable scientific language (e.g., graphs or structural equations) by the domain experts. The selection bias mechanisms depicted in Figure 1 are common examples². in epidemiology (Hernán *et al.*, 2004). The directed acyclic graphs should be regarded as graphical structures encoding conditional independence between X , Y , and S which may involve other variables as as well variables, like M , that is not observed in the target domain, and thus that is not included as input variable to the model.

THEOREM 2. — *Given that condition of Theorem 1 is satisfied, if \hat{P} is a distribution such that: $\hat{P}(x, y, x_s, s) = P(x, y, x_s, s) \beta(x_s)$ then $\hat{P}(x, y | s = 1) \equiv P(x, y)$.*

1. Upper-case letters in italics denote random variables (e.g., X, Y) and lower-case letters in italics denote their values (e.g., x, y). X denotes the input variables and Y the target.

2. We assume the reader is familiar with the concepts of d -separation (Pearl, 1989)

Theorem 2 states that an unbiased training sample can be obtained by weighting each training example by $\beta(x_s) = \frac{P(x_s)}{P(x_s|s=1)}$. Note however that the support of $P(x_s)$ should be contained in the support of $P(x_s|s=1)$ for the $\beta(x_s)$ to be always defined. A similar technique applied to covariate shift was discussed in (Zadrozny, 2004).

3. Experiments

In this section, we assess the ability of importance weighting to remove complete selection bias based on Theorem 2. In the toy experiment, we investigate whether covariate shift and prior probability shift corrections may help reduce complete selection bias despite our assumptions between the training and test distributions difference being violated (through an invalid choice for X_s).

When the selection process is explicitly known, $\beta(x_s)$ is simply given by $\frac{P(s=1)}{P(s=1|x_s)}$. Otherwise, we resort to an external (bias-free) data set of X_s in order to estimate $\beta(x_s)$ as $\frac{P(x_s)}{P(x_s|s=1)}$. In this study, we use the Kernel Mean Matching (KMM) (Huang *et al.*, 2006; Yu, Szepesvári, 2012) estimator for $\beta(x_s)$ denoted as $\text{KMM}(X_s)$. As one usually has a partial understanding of the sampling mechanism, we investigate whether covariate shift (i.e., $\beta(x) = \frac{P(x)}{P(x|s=1)}$) and prior probability shift (i.e., $\beta(y) = \frac{P(y)}{P(y|s=1)}$), corrections may help reduce complete selection bias despite not being valid. These strategies are denoted as $\text{KMM}(X)$, $\text{KMM}(Y)$. We first apply our method to a simple toy problem and then compare KMM to another estimator called the Unconstrained Least-Square Importance Fitting (uLSIF) (Kanamori *et al.*, 2009) on a variety of regression and classification benchmarks from the UCI Archive.

3.1. Toy problem

Consider the S -bias mechanism displayed in Fig. 1b, where the feature X has a uniform distribution in $[0, 1]$: $P(X) \sim \mathcal{U}(0, 1)$. Note that the influence of M on Y is mediated by $\{X, S\}$. The observations are generated according to $y = 1 - 0.5x$ and are observed in Gaussian noise with standard deviation 0.5 (see Fig. 2c; the black solid line is the noise-free signal). The intermediate variable M , between X and S , is generated according to $M = X + \mathcal{N}(0, 0.3^2)$. As M is only measured in the training set, it is not used as an input variable in our regression model. Therefore, we investigate a case where X_s is partially missing in the test set. The probability of a given example being included in the training set depends on Y and M and is given by

$$P(S = 1|m, y) \sim \begin{cases} y - m, & \text{if } 0.1 \leq (y - m) \leq 1 \\ 0.1, & \text{if } (y - m) \leq 0.1 \\ 1, & \text{otherwise} \end{cases}$$

Note that the minimum value of $P(S = 1|m, y)$ needs to be greater than 0 so that the support of $P(m, y)$ is contained in the support of $P(m, y|s = 1)$, as required by

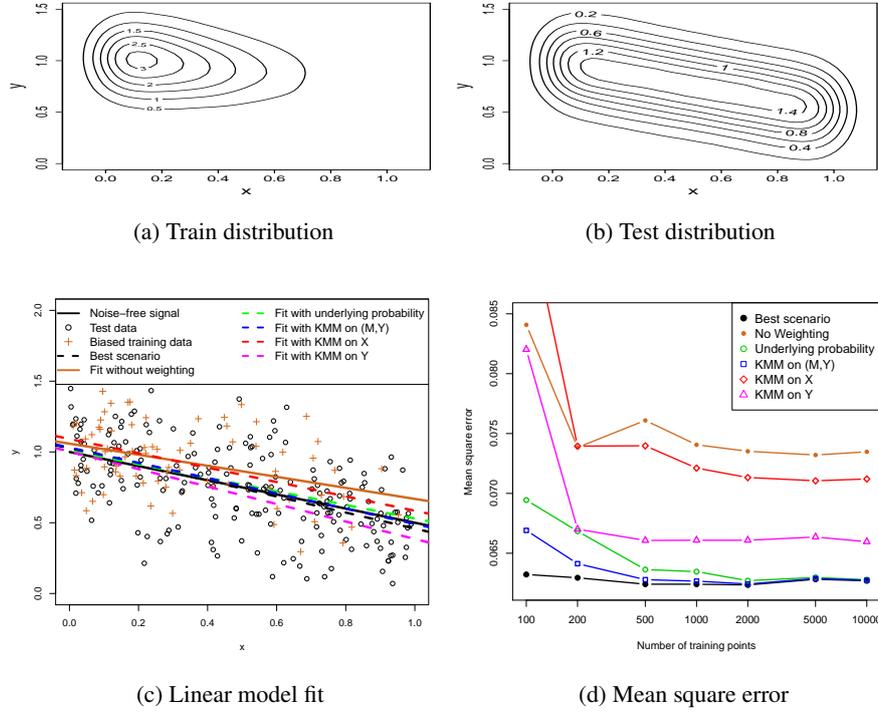


Figure 2: Toy regression problem 1. (a) and (b) Contour plots X - Y on training and test sets; (c) Polynomial models of degree 1 fit with OLS and WOLS; (d) Average performances of four WOLS methods and OLS on the test data as a function of the number of training points.

Theorem 1. The choice of $P(m, y)$ is intended to induce a noticeable discrepancy between $P(y|x, s = 1)$ and $P(y|x)$. We sampled 200 training (red crosses in Fig. 2c) and testing (grey circles) points from P_{tr} and P_{te} respectively. The bias is clearly noticeable from the X - Y contour plots in Fig. 2a and b. The bias-free distribution $P(x, y)$ is recoverable from the S -bias training samples since $\{M, Y\}$ satisfies Theorem 1. Thus we use Theorem 2, to remove selection bias by weighting each example by the importance ratio:

$$\beta(x_s) = \beta(m, y) = \frac{P(m, y)}{P(m, y|s = 1)} = \frac{P(s = 1|m, y)}{P(s = 1)}$$

where $P(s = 1|m, y)$ and $P(s = 1)$ may be obtained from the known selection mechanism shown above or directly estimated by KMM. We attempted to model the observations with a degree 1 polynomial. The black dashed line in Fig. 2c is a best-case scenario given our test points, which is shown for reference purposes: it

represents the model fit using ordinary least squared (OLS) on the test set. The brown line is a second reference result, derived only from the training data via OLS, and predicts the test data very poorly. The green dashed line is a third reference result, fit with weighted ordinary least square (WOLS), using the true $\beta(x_s)$ values calculated from the true data generating mechanism, and predicts the test data quite well. The other three dashed lines are fit with WOLS using the KMM weighting schemes under the three assumptions. Note that the true generating model between X and Y is included in the hypothesis space. We estimated the effect of the number of training points on the estimation of the reweighting factors by examining the average mean square error (MSE) on the test set as a function of the number of training points. As may be observed in Fig. 2d, the error goes down as the sample size increases, until it reaches an asymptotic value. $\text{KMM}(X_s)$ performs well even with relatively moderate amounts of data achieving almost optimal error quite quickly, handily outperforming the reweighting method based on $\text{KMM}(X)$ and $\text{KMM}(Y)$ by a noticeable margin. More interestingly, $\text{KMM}(X_s)$ also outperforms the reweighting method based on the true data generating mechanism, especially when sample size is small. This result may seem counter-intuitive at first sight: the reason is that the exact importance-sampler weights are not always optimal unless we have an infinite sample size. See (Shimodaira, 2000) for a thorough discussion. Remarkably, despite our assumption regarding the difference between the training and test distributions being violated, $\text{KMM}(Y)$ and $\text{KMM}(X)$ improve the test performance.

3.2. Real-world data sets

We now examine whether using importance weighting can reduce selection bias in 10 UCI data sets with 5 classification tasks and 5 regression tasks. We employ three methods to estimate importance weighting: ratio of underlying probability, KMM and uLSIF and compare their performance against the baseline unweighted method. For each data set, X_s is chosen to be the label Y and the most correlated input variable to Y (denoted as X_1 for simplicity). The selection bias mechanism is illustrated in Fig. 1c. The selection variable S for each training example is determined according to two scenarios depending on whether it is regression or classification problem. For regression problem, we use $P(s = 1|x_1, y) = \exp(ax_1 + by + c)/[1 + \exp(ax_1 + by + c)]$, where a, b, c , are parameters that determine the bias. For binary classification problem, we use:

$$P(s = 1|x_1, y) = \begin{cases} 0.5 & \text{if } x_1 > \text{mean}(x_1) \text{ and } y = 1 \\ 1 & \text{if otherwise.} \end{cases}$$

For each data set, we then train 4 predictive models learned under the four weighting schemes discussed above and a model learned from the unbiased data (baseline) using SVM-light (Joachims, 1998) which allows importance weighting to be fed directly to SVM. All classifiers are trained with the common Radial Basis Function

Data set	No weighting	KMM	uLSIF	Underlying P	Unbiased model
India diabetes	0.338 ± 0.049	0.266 ± 0.040	0.332 ± 0.053	0.287 ± 0.055	0.258 ± 0.035
Ionosphere	0.069 ± 0.039	0.066 ± 0.039	0.067 ± 0.040	0.067 ± 0.039	0.065 ± 0.036
BreastCancer	0.044 ± 0.016	0.039 ± 0.015	0.043 ± 0.017	0.040 ± 0.016	0.038 ± 0.015
Haberman	0.264 ± 0.069	0.262 ± 0.071	0.263 ± 0.070	0.262 ± 0.071	0.262 ± 0.071
GermanCredit	0.300 ± 0.044	0.298 ± 0.046	0.298 ± 0.045	0.298 ± 0.046	0.295 ± 0.046
Airfoil self noise*	0.534 ± 0.104	0.470 ± 0.122	0.475 ± 0.082	0.445 ± 0.081	0.403 ± 0.059
Abalone*	0.526 ± 0.048	0.484 ± 0.054	0.521 ± 0.057	0.466 ± 0.041	0.456 ± 0.036
Computer Hardware*	0.326 ± 0.308	0.321 ± 0.304	0.321 ± 0.299	0.319 ± 0.307	0.305 ± 0.201
Auto MGP*	0.268 ± 0.148	0.298 ± 0.192	0.212 ± 0.129	0.203 ± 0.128	0.129 ± 0.063
Boston Housing*	0.323 ± 0.110	0.327 ± 0.127	0.349 ± 0.133	0.332 ± 0.127	0.298 ± 0.112

Table 1: Mean test error averaged over 100 trials of different weighting schemes on UCI data set. Data sets marked with * are for regression problems

(RBF), with a kernel size σ chosen through a 5-fold cross validation. This procedure is repeated 100 times for each data set.

The results are reported in Tab. 1. As may be seen, all importance weighting schemes achieve lower prediction error with respect to the baseline unweighted scheme. The underlying probability weighting scheme performs pretty good. Curiously, on the Boston Housing data set, all three weighting schemes perform worse than the baseline unweighted method. It seems therefore that the effectiveness of bias correction based on importance weighting is data dependent. In order to better assess the overall results obtained for each of the 4 weighting schemes, a non-parametric Friedman test was firstly used to evaluate the rejection of the hypothesis that all the models perform equally well (except the unbiased model of course) at significant level 5%. Statistically significant differences were observed. So we proceeded with the Nemenyi post hoc test. The results along with the average rank diagrams are shown in Fig. 3. The ranks are depicted on the axis, in such a manner that the best ranking algorithms are at the rightmost side of the diagram. The algorithms that do not differ significantly (at $p = 0.05$) are connected with a line. As may be observed in Fig. 3, contrary to uLSIF, KMM is significantly better than no weighting.

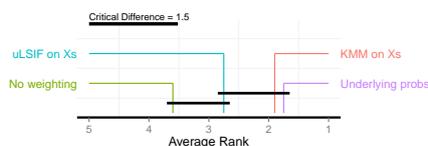


Figure 3: post hoc analysis

4. Discussion & Conclusion

The aim of this paper was to elaborate on the idea of exploiting the assumptions that are deemed plausible about the sampling mechanism to correct or reduce selection bias in machine learning tasks. The method hinges on the existence of a S-control

feature vector, X_s , and an additional (biased-free) sample that allows us to estimate the distribution of X_s . We showed experimentally that direct weighting estimation is able to achieve significant improvements in accuracy over the unweighted method, even in situations where our key assumption is not valid (assuming covariate shift and prior probability shift instead of complete selection bias). However the gain in accuracy is data dependent. In fact, all conclusions are extremely sensitive to which variables we choose for X_s . As the choice of X_s usually reflects the investigator's subjective and qualitative knowledge of statistical influences in the domain, the data analyst must weight the benefit of reducing selection bias against the risk of introducing new bias carried by unmeasured covariates even where none existed before. Nevertheless, we hope this study will convince others about the importance of selection bias correction methods in practical studies and suggest relevant tools which can be used to achieve that goal.

References

- Bareinboim E., Tian J., Pearl J. (2014). Recovering from selection bias in causal and statistical inference.
- Hernán M. A., Hernández-Díaz S., Robins J. M. (2004). A structural approach to selection bias. *Epidemiology*, Vol. 15, No. 5, pp. 615–625.
- Huang J., Smola A. J., Gretton A., Borgwardt K. M., Schölkopf B. (2006). Correcting sample selection bias by unlabeled data. In *Nips*, p. 601-608. MIT Press.
- Joachims T. (1998). *Making large-scale svm learning practical*. LS8-Report No. 24. Universität Dortmund, LS VIII-Report.
- Kanamori T., Hido S., Sugiyama M. (2009, December). A least-squares approach to direct importance estimation. *J. Mach. Learn. Res.*, Vol. 10, pp. 1391–1445. Retrieved from <http://dl.acm.org/citation.cfm?id=1577069.1755831>
- Pearl J. (1989). *Probabilistic reasoning in intelligent systems - networks of plausible inference*. Morgan Kaufmann.
- Shimodaira H. (2000, October 01). Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, Vol. 90, No. 2, pp. 227–244. Retrieved from <http://www.sciencedirect.com/science/article/B6V0M-4136355-5/1/6432c256e0be03b1503bbf79e4e91d1a>
- Yu Y., Szepesvári C. (2012). Analysis of kernel mean matching under covariate shift. *CoRR*, Vol. abs/1206.4650.
- Zadrozny B. (2004). Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on machine learning*, pp. 114–. New York, NY, USA, ACM. Retrieved from <http://doi.acm.org/10.1145/1015330.1015425>