



HAL
open science

Fast and Accurate Importance Weighting for Correcting Sample Bias

Antoine de Mathelin, Francois Deheeger, Mathilde Mougeot, Nicolas Vayatis

► **To cite this version:**

Antoine de Mathelin, Francois Deheeger, Mathilde Mougeot, Nicolas Vayatis. Fast and Accurate Importance Weighting for Correcting Sample Bias. Joint European Conference on Machine Learning and Knowledge Discovery in Databases, 2022, Grenoble, France. pp.659-674. <hal-04455961>

HAL Id: hal-04455961

<https://hal.science/hal-04455961v1>

Submitted on 13 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Fast and Accurate Importance Weighting for Correcting Sample Bias

Antoine de Mathelin^{1,2}✉, Francois Deheeger¹, Mathilde Mougeot², and Nicolas Vayatis²

¹ Manufacture Française des Pneumatiques Michelin, Clermont-Ferrand, France
{antoine.de-mathelin-de-papigny, francois.deheeger}@michelin.com

² Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, Gif-sur-Yvette, France {mathilde.mougeot, nicolas.vayatis}@ens-paris-saclay.fr

Abstract. Bias in datasets can be very detrimental for appropriate statistical estimation. In response to this problem, importance weighting methods have been developed to match any biased distribution to its corresponding target unbiased distribution. The seminal Kernel Mean Matching (KMM) method is, nowadays, still considered as state of the art in this research field. However, one of the main drawbacks of this method is the computational burden for large datasets. Building on previous works by Huang et al. (2007) and de Mathelin et al. (2021), we derive a novel importance weighting algorithm which scales to large datasets by using a neural network to predict the instance weights. We show, on multiple public datasets, under various sample biases, that our proposed approach drastically reduces the computational time on large dataset while maintaining similar sample bias correction performance compared to other importance weighting methods. The proposed approach appears to be the only one able to give relevant reweighting in a reasonable time for large dataset with up to two million data.

1 Introduction

The most common assumption in a traditional learning scenario is that training data are independently and identically distributed (iid) and drawn from the same distribution as the target data. However, in real cases, the training dataset often appears to be biased with respect to the target dataset. This happens in particular in medical applications, when, for example, the age distribution of the patients does not match the distribution of the overall population. In product design, predictive models of product performances may be biased by the large amount of data corresponding to outdated products. For both previous cases, it often happens that the learner has access to the unbiased distribution, either because it is known from an external source (the age distribution in the whole population is known) or because he has access to an unbiased dataset (a sample of data of the recent products).

In this paper we assume that the learner owns a sample drawn from a source biased distribution $p_s(x, y)$ as well as a sample coming from the target marginal

distribution $p_t(x)$ such that $p_s(x) \neq p_t(x)$. Our goal is to estimate $p_t(y|x)$ or $p_t(y)$ where y is the variable of interest (e.g. patient survival expectancy for a clinical model or product performances for product design). Given the bias on the marginals, the estimation of $p_t(y|x)$ on the target domain will be biased as well.

To correct this type of sample bias, importance weighting methods can be used. These methods seek to reweight the source data to debias the marginals by looking for the weights corresponding to $w(x) = p_t(x)/p_s(x)$. A successful non-parametric method in this field is the Kernel Mean matching (KMM) method [16] which reweights the sources in order to minimize the MMD distance between the reweighted sources and the targets [14]. Although KMM is one of the first non-parametric method developed to handle sample bias, it is still used nowadays in modern sample bias correction methods for deep learning [11] or for deriving two-sample hypothesis testing under sample bias [1]. KMM solves a quadratic problem for the minimization of the MMD with as many parameters as the number of source instances. Thus, when the number of source data is large, one faces a computational burden because of the large kernel matrix to compute. Some methods have proposed to reduce the problem in batch and to perform a KMM on each of them [6], [27]. This lightens the memory, but the computational time remains important as the number of KMM sub-problems to compute increases with the number of data. Other importance weighting methods reduce the number of parameters to be optimized by linking the weights of the source instances by a parametric function as done for KLIEP [33] and ULSIF [17]. These two methods propose to write each of the weights as a linear combination of kernels centered on target points. Thus, the number of parameters is fixed (in general by selecting a hundred centers in the target domain), however, the computational cost of the pairwise distance calculations between the centers and all the source data still remains. Moreover, this large matrix of pairwise distances is used in the resolution of the gradient descent algorithm which slows the optimization. A last method, NearestNeighborsWeighting (NNW) consists in computing the weights of the source instances according to their number of target nearest neighbors [21]. This heuristic, not relying on the minimization of a distance between distributions, is quite efficient, and its computation time is generally less than its KMM counterpart. However, the search for the nearest neighbors requires the computation of pairwise distances between source and target data and despite the optimization algorithms of type KDTree [12] and BallTree [29], the method encounters computational burden for datasets with many instances and features.

Finally, all these algorithms rely on hyper-parameters to be tuned. The choice of the kernel and its bandwidth for the KLIEP and KMM methods are very important, as well as the number of nearest neighbors to consider for NNW. To choose these parameters, a cross-validation procedure using an unsupervised metric (which does not require the y data on the target domain) is mainly used such as the J-score for KLIEP [33], the normalized mean squared error (NMSE) between the actual and estimated density ratios for KMM [26], an information

criterion for ULSIF [17], or any divergence metric between distributions such as the linear discrepancy or the domain classifier divergence [2], [8], [22]. This hyper-parameter selection procedure, necessary to use these methods in practice, adds to the computational time.

Considering the drawbacks of the previous mentioned methods, we propose, in this paper, a new importance weighting algorithm that scales to large datasets. Our goal is to obtain the same level of performance than KMM but with less computational time. To do so, we propose to minimize the objective of KMM, i.e. the MMD, by a batch gradient descent to avoid the memory burden of the huge kernel matrix. However, it should be underlined that optimizing the weights of the source instances individually brings no complexity gain since each source weight is only updated in its corresponding batch. Assuming $w(x) = p_t(x)/p_s(x)$ continuous and regular, for two close source points $x_1 \simeq x_2$, the weights will be similar $w(x_1) \simeq w(x_2)$. Consequently, we propose to optimize at each batch the parameters θ of a parametric and continuous function $W_\theta(x)$ in order to minimize the empirical MMD on the batch. Inspired from recent works of weighting adversarial neural network (WANN) [25], this function W_θ is chosen as a neural network. The advantage of the networks is the fast update of the θ parameters by backpropagation of the gradient through the layers which is highly parallelizable [19]. This avoids working with huge matrices of pairwise kernel as done in the KLIEP algorithm. We show on several datasets that this approach allows to obtain importance weighting at least as efficient as KMM in a drastically reduced time. The source code of the experiments is publicly available on GitHub³.

Our contributions can be listed as follows:

- We derive a fast and scalable importance weighting algorithm. This is achieved by using batch gradient descent optimizing the MMD and by parameterizing the weights by a neural network.
- The developed algorithm optimizes the kernel parameters of the MMD in the gradient-descent optimization and thus avoid a time consuming CV process to select it.

2 Problem Setting and Proposed Approach

2.1 Learning scenario

Given an input space $\mathcal{X} \in \mathbb{R}^p$ of dimension $p > 0$ and an output space $\mathcal{Y} \in \mathbb{R}^q$ with $q > 0$, we consider the sample bias scenario in which the learner has access to a source sample $\mathcal{S} = \{(x_1, y_1), \dots, (x_m, y_m)\} \subset \mathcal{X} \times \mathcal{Y}$ drawn iid from a source distribution $p_s(x, y)$ on $\mathcal{X} \times \mathcal{Y}$ and a target sample $\mathcal{T} = \{x_1, \dots, x_n\} \subset \mathcal{X}$ drawn iid according to a target distribution $p_t(x)$ on \mathcal{X} . We suppose that $p_s(x) \neq p_t(x)$ and that $p_t(x)$ is absolutely continuous with respect to $p_s(x)$. Finally, we make the covariate-shift assumption [3] which states that the conditional probabilities of $y|x$ remain unchanged for the two distributions: $p_s(y|x) = p_t(y|x)$.

³ <https://github.com/antoinedemathelin/Importance-Weighting-Network>

2.2 MMD

Let's consider $\phi_\sigma : \mathcal{X} \rightarrow \mathcal{F}_\sigma$ with \mathcal{F}_σ the RKHS of Gaussian kernel k_σ such that $\forall x, x' \in \mathcal{X}$, $k_\sigma(x, x') = \langle \phi_\sigma(x), \phi_\sigma(x') \rangle = \exp(-\sigma \|x - x'\|^2)$ with $\sigma > 0$. The Maximum Mean Discrepancy (MMD) between the source and target distributions is defined as follows:

$$\text{MMD}_\sigma(p_s(x), p_t(x)) = \left\| \mathbb{E}_{x \sim p_s(x)} [\phi_\sigma(x)] - \mathbb{E}_{x \sim p_t(x)} [\phi_\sigma(x)] \right\| \quad (1)$$

The MMD is a distance characterizing how close are the two marginal distributions $p_s(x), p_t(x)$. As we consider a Gaussian kernel, $\text{MMD}_\sigma = 0$ if and only if $p_s(x) = p_t(x)$ [14].

As our goal is to correct the sample bias between the source and target distributions with importance weighting, we aim at finding the weights $w(x) \in \mathbb{R}_+$ that solve the following optimization problem:

$$\begin{aligned} \min_{w: \mathcal{X} \rightarrow \mathbb{R}_+} & \left\| \mathbb{E}_{x \sim p_s(x)} [w(x)\phi_\sigma(x)] - \mathbb{E}_{x \sim p_t(x)} [\phi_\sigma(x)] \right\|^2 \\ \text{subject to } & w(x) \geq 0 \forall x \in \mathcal{X} \text{ and } \mathbb{E}_{x \sim p_s(x)} [w(x)] = 1 \end{aligned} \quad (2)$$

As we consider $p_t(x)$ absolutely continuous with respect to $p_s(x)$, the solution of the optimization problem (2) is the density ratio $w(x) = p_t(x)/p_s(x)$ [16].

In practice, we only have access to samples $\{x_1, \dots, x_m\}$ and $\{x'_1, \dots, x'_n\}$ respectively drawn according to both distributions $p_s(x)$ and $p_t(x)$, we then consider the empirical formulation of the previous optimization problem (2) which is written:

$$\begin{aligned} \min_{w \in \mathbb{R}^m} & \frac{1}{m^2} \sum_{i,j} w_i w_j k_\sigma(x_i, x_j) + \frac{1}{n^2} \sum_{i,j} k_\sigma(x'_i, x'_j) - \frac{2}{nm} \sum_i \sum_j w_i k_\sigma(x_i, x'_j) \\ \text{subject to } & w_i \geq 0 \forall i \in [1, m] \text{ and } \frac{1}{m} \sum_i w_i = 1 \end{aligned} \quad (3)$$

2.3 Importance Weighting Network

The optimization problem (3) is a quadratic optimization problem which can be solved by gradient descent. However computing the MMD requires to compute a kernel matrix of size $\mathcal{O}((n+m)^2)$ which can cause memory burden. We propose, in this paper, to compute the MMD on small batches of size B . At each batch, we impose the constraints on the weights by taking their absolute values and dividing them by their sum. It has been shown that, although self-normalizing the weights creates a biased estimation of the MMD, the estimator is asymptotically unbiased, with the bias decreasing at a rate of $\mathcal{O}(1/B)$ [9], [23].

To obtain a fast update of all weights at each batch, we parameterized the weights through a neural network $W_\theta : \mathcal{X} \rightarrow \mathbb{R}$ such that $w_i = W_\theta(x_i)$ for each $i \in [1, m]$. In this way, at each batch, the parameters θ are updated and then all the parameters w_i are updated with them. Notice that the MMD estimation produced by the batch of size B is approaching the true MMD at a strong rate of $\mathcal{O}(1/\sqrt{B})$ [14] which comforts the idea that the update of θ at each batch will be in favor of finding the optimal weights for all x_i .

It should be stressed that the MMD quantity depends on σ which corresponds to the kernel bandwidth. In the seminal paper of KMM [16] the choice of σ is not clearly motivated, but a method proposed by KLIEP [33] consists in choosing between several predefined σ and compute the optimal weights for each σ value. The value which provides the best matching of the target distribution is finally selected. This type of selection is time consuming and requires fixing a pre-selection of σ values.

We propose, instead, to optimize the σ parameter at the same time as the weights. Inspired by the works on MMD-GAN [20], the kernel parameter σ is modified at each batch in order to maximize the MMD. The idea behind this choice of implementation is to increase the discriminative power of the MMD and thus reduce the risk of estimating, from finite samples, that the source and target distributions are the same when this is not the case. By maximizing over σ , we end with an alternate gradient descent-ascent algorithm, where we aim at finding a saddle point. The final optimization formulation can be written as follows:

$$\begin{aligned} \max_{\sigma} \min_{\theta} & \frac{\sum_{i,j}^B |W_\theta(x_i)W_\theta(x_j)|k_\sigma(x_i, x_j)}{\sum_{i,j}^B |W_\theta(x_i)W_\theta(x_j)|} \\ & + \frac{1}{B^2} \sum_{i,j}^B k_\sigma(x'_i, x'_j) \\ & - \frac{2 \sum_i^B \sum_j^B |W_\theta(x_i)|k_\sigma(x_i, x'_j)}{B \sum_i^B |W_\theta(x_i)|} \end{aligned} \quad (4)$$

We therefore introduce the Importance Weighting Network (IWN) which searches for the saddle points that solve the above optimization problem (cf Algorithm 1).

3 Related work

Instance-based domain adaptation. Our work is in line with instance-based unsupervised transfer learning or domain adaptation [30]. Most of the instance-based methods have already been introduced previously as KMM, KLIEP, ULSIF and RULSIF [16], [33], [17], [36]. All of these methods aim to compute the source weights which minimize a distance between the input distributions like the MMD

Algorithm 1 Importance Weighting Network

Inputs: Source and target datasets $\mathcal{S}_X, \mathcal{T}$, initial bandwidth σ , batch size B , neural network W_θ , learning rate ν

Initialization: Fit W_θ with loss $\mathcal{L} = \sum_i \|W_\theta(x_i) - 1\|^2$

while stopping criterion is not reached **do**

 Take batches $\{x_1, \dots, x_B\} \subset \mathcal{S}_X$ and $\{x'_1, \dots, x'_B\} \subset \mathcal{T}$

Forward propagation

$w_i \leftarrow |W_\theta(x_i)| / \sum_j^B |W_\theta(x_j)| \forall x_i$

$\text{MMD}_{\sigma, \theta} = \sum_{i,j}^B w_i w_j k_\sigma(x_i, x_j) + \frac{1}{B^2} \sum_{i,j}^n k_\sigma(x'_i, x'_j) - \frac{2}{B} \sum_i^B \sum_j^B w_i k_\sigma(x_i, x'_j)$

Backward propagation

$\theta \leftarrow \theta - \nu \nabla_\theta \text{MMD}_{\sigma, \theta}$

$\sigma \leftarrow \sigma + \nu \nabla_\sigma \text{MMD}_{\sigma, \theta}$

or the Kullback-Leibler. Other methods have also proposed to take into account the model used to estimate y using appropriate metrics such as the discrepancy [7], [22]. Most of the unsupervised instance-based approaches make the assumption of covariate-shift [3].

Importance weighting and deep learning. This work is related to existing works in importance weighting using deep learning. Recently, Fang et al. [11] have developed a task-oriented sample-bias correction method where a KMM is performed at each batch in different depths of the neural network. In a different context from ours, Diesendruck et al. [9] have developed a sample bias correction method for deep generative models. In this approach, the MMD is also minimized by batch, however the weights are not parameterized but assumed to be known (e.g. a uniform proportion of classes is desired). Importance weighting methods have also been used along with deep feature transformation in partial domain adaptation [4], i.e. when the number or the proportion of classes differ between targets and sources. In this category of methods, the output of a domain classifier network is often used to reweight the source instances. The domain classifier is either trained in parallel to the feature transformation [5], [37], [38] or after it [31], [34]. Other methods in this field consider the uncertainty of a task classifier to reweight both source and target instances during the feature transformation [15], [35]. These works are interesting from a computational point of view and may be seen as an alternative to MMD minimization. Finally, the weighting adversarial neural network (WANN) [25], explicitly proposes to use a neural network to learn the source weights minimizing a distance between distribution called the \mathcal{V} -discrepancy [28]. Their approach, however, is developed in the supervised context and involves a task network fitted at the same time as the weights. Their approach is then deep learning specific. The present work generalizes this last approach as any estimator can be used once the weights are computed.

4 Experiments

We conduct the experiments on a synthetic dataset and 15 UCI datasets⁴ [10] of various size and number of features. The experiments are conducted on a 3.3Ghz computer with 64G RAM and 24 Cores. The source code of the experiments is available on GitHub⁵.

4.1 IWN Settings

The purpose of IWN is to provide a simple and fast tool to perform importance weighting. We observe that the choice of network has little incidence on the learned weights (see Section 4.5), we then arbitrarily choose a three layers neural network with 100 neurons each and a ReLU activation. This architecture is used in all experiments without fine-tuning. We choose the Adam optimizer [18]. The optimization parameters are also fixed for all experiments to a learning rate of 0.001, a batch size of 256 and a maximal number of iterations set to $5 \cdot 10^4$. Early stopping on the objective function is used, if the objective has not improved after $2 \cdot 10^4$ iterations, the learning is stopped.

We remind that the kernel bandwidth σ used to compute the MMD is learned in the gradient descent (cf Algorithm 1) and does not require a cross-validation process. The initial value of σ is set to 0.1 for all experiments.

4.2 Competitors Settings

We consider the following competitors which have already been introduced previously in this paper:

- KMM [16]. We use a Gaussian kernel and the default optimization parameters $B = 1000$ and $\epsilon = \sqrt{m-1}/\sqrt{m}$. The bandwidth σ of the kernel is selected in the set $\{10^{(i-4)}\}_{i \in \llbracket 0,8 \rrbracket}$ with unsupervised cross-validation using the linear discrepancy [22].
- KLIEP [33]. We use a Gaussian kernel and a learning rate of 0.01 with a maximum number of iterations of 1000 as parameters for the gradient descent. These parameters have been selected to obtain an important decrease of the objective function with a fast convergence for most of the datasets. The bandwidth σ of the kernel is selected in the set $\{10^{(i-4)}\}_{i \in \llbracket 0,8 \rrbracket}$ with the native Likelihood Cross-Validation (LCV) procedure of KLIEP.
- NNW [21]. The nearest neighbors are computed with the optimized NearestNeighbors algorithm of scikit-learn [32] which optimizes the computation approach between brute force and KD-Ball-Tree in function of the number of features and samples in the dataset. The Euclidean distance is used and the number of nearest neighbors for averaging is chosen in the set $\{1, 5, 10, 20, 50, 100\}$ with unsupervised cross-validation using the linear discrepancy.

⁴ <https://archive.ics.uci.edu/ml/datasets.php>

⁵ <https://github.com/antoinedemathelin/Importance-Weighting-Network>

The implementation of the competitors are provided by the ADAPT library⁶ [24]. The library also provides the metric for the cross-validation processes.

To offer the best chance to the competitors, the parameters selection with cross-validation are performed with parallel computing for KMM and NNW. For KLIEP, the parallel computing is not available in ADAPT and thus not used, this explains why its computational time is higher than the others.

4.3 Synthetic dataset

We consider the synthetic experiment, inspired from [7], where $p_s(x)$ is a mixture of $M = 10$ Gaussians, i.e. $p_s(x) = \sum_{k=1}^M \pi_k \mathcal{N}(\mu_k, 0.2)(x)$ where the centers $\mu_k \in \mathbb{R}^N$ are drawn according to the distribution $\mathcal{N}(0, 1)$ in \mathbb{R}^N , the ratios π_k are set such that $\pi_k = 0.8/(M - 1) \forall k \neq M$ and $\pi_M = 0.2$. The output variable is written $y = \beta_k^T x$ for any x drawn according to the k^{th} Gaussian. The coefficients $\beta_k \in \mathbb{R}^N$ are drawn according to the distribution $\mathcal{N}(0, 1)$ in \mathbb{R}^N . The target distribution $p_t(x)$ is drawn according to the same mixture of Gaussians but with ratios $\pi'_k = 0.1/(M - 2) \forall k < M - 1, \pi'_{M-1} = 0.1$ and $\pi'_M = 0.8$. We suppose that the learner has access to a sample of size n of source labeled instance $\{(x_i, y_i)\}_{1 \leq i \leq n}$ drawn according to the source distribution $p_s(x, y)$ on \mathbb{R}^N and an unlabeled set of size n , $\{x'_j\}_{1 \leq j \leq n}$ drawn according to the target distribution $p_t(x)$ on \mathbb{R}^N . An illustration of the problem for dimension $N = 2$ is given in Figure 1.A.

We conduct several experiments on this synthetic dataset. First, we fit IWN for the setting $N = 32, n = 10000$. We make 4000 batch updates with batch size 256, at each of them, we use the weights w_i returned by the weighting network to fit a weighted Ridge regression model of parameters β on the set $\{(x_i, y_i, w_i)\}_{1 \leq i \leq n}$, we then record the mean absolute error (MAE) of this model on the target dataset: $\frac{1}{n} \sum_i |\beta^T x'_i - y'_i|$. We also record, at each batch, the computed MMD on the batch and the "true" MMD computed with the whole samples. We also record the current value of the parameter σ which is also updated during the optimization (see Section 2.3). We report the results of this experiment on Figure 1.B, 1.C, 1.D. We first represent the final importance weights returned by the weighting network at the end of the 4000 iterations in Figure 1.B. As we can observe, the learned weights are very close to the true sampling probability. On Figure 1.C, we report the evolution of the recorded MAE in plain orange, the batch and true MMD in blue and the value of σ in green. We also report, for comparison with "MAE (IWN)", the MAE of a Ridge model fitted with uniform weights: "MAE (Unif)" and the MAE of the model fitted with the weights obtained using KMM: "MAE (KMM)". We observe that the importance weighting produced by IWN helps to learn the task on the target domain as the error decreases of 20% compared to the error of the model fitted with uniform weights. Concerning the recorded MMD, we observe that both the batch and "true" MMD decrease very fast, but an offset remains between the two due to the estimation error made with finite samples. We then see on the zoom of Figure 1.D that the error decreases very

⁶ <https://github.com/adapt-python/adapt>

fast as well as the MMD. After 100 iterations, the MMD is minimized and the target error of IWN is on the same level than the error produced by KMM. We notice that the MAE increases a little after some iterations which may indicates some overfitting effect of the weighting network. This observation argues for the use of early stopping based on the evolution of the MMD.




Fig. 1. Visualization of the synthetic experiments. The corresponding experimental settings are: (A) $n = 10000, N = 2$, (B, C, D) $n = 10000, N = 32$, (E) $n \in [100, 2 \cdot 10^5], N = 128$, (F) $n = 2 \cdot 10^4, N \in [16, 4096]$

We finally remark, in Figure 1.C, that the σ parameter becomes relatively stable around 0.15 after some iterations which comforts the idea that the parameter can be efficiently set during the optimization thanks to the adversarial learning (see Section 2.3). We could however argue, that the increase of the target error after the 100th iteration is correlated to the increase of σ from 0.1 to 0.15. This is a plausible explanation as σ is updated in order to make the MMD more discriminative which may provoke overfitting. This is a limitation of the proposed approach which we propose to study in future work.

We then conduct several experiments to observe the evolution of the computational time in function of the sample size and the number of features. We first fix the number of features to $N = 128$ and vary the number of samples from 100 to $2 \cdot 10^5$ and report the computational time in Figure 1.E. Then, we fix the number of samples to $n = 20000$ and vary N from 16 to 4096 and report the results in Figure 1.F. We observe, on these two Figures, that for the number of samples has a stronger impact on the computational time than the number of features. The two methods KMM and NNW have a quadratic complexity $\mathcal{O}(n^2)$ which is well reflected in Figure 1.E. The computational time of KLIEP evolves linearly due to the fixed number of target centers considered. Figure 1.E clearly demonstrates the computational supremacy of IWN compared to the other methods. It should be pointed out that the quality of the importance weighting is similar between methods, the corresponding scores are reported in appendix.

4.4 UCI datasets

We perform the experiments on several UCI datasets [10] with different sizes and dimensions. We record the computational time used by each importance weighting method for computing the source importance weights (cf Figure 3). For each dataset, the score is computed with a Ridge model fitted with the importance weights and without importance weighting. The ratio between the two scores is reported in Figure 2.

To evaluate the importance weighting methods we consider different kind of sample bias following the setting of [16]:

- Sample bias on the input features: the source training set is biased on the input features X with a gaussian weighting on the first component of the PCA of mean $m + (\mu - m)/3$ and standard deviation $(\mu - m)/8$ with m, μ the respective minimum and mean of the first PCA component.
- Sample bias on the output features: the source training set is biased on the output features y with a weighting on the first component of the PCA defined as $\exp(3(y_1 - 1))/(1 + \exp(3(y_1 - 1)))$ with y_1 the first PCA component of y .

The training set is built by taking n data with replacement using the sampling bias as probabilities of selection. The target set is the original dataset without selection bias. For each experiment, we apply the following preprocessing: standard scaling of the numerical inputs (using the mean and standard deviation of the unbiased inputs) and one-hot-encoding of the categorical inputs. The dimension p of the input space corresponds to the dimension after preprocessing.

To evaluate the weighting scheme of each method, we fit a Ridge model with trade-off parameters α selected by leave-one-out process between values $\{10^{(i-4)}\}_{i \in [0,8]}$. For regression datasets, we compute the mean absolute error (MAE) on the target testing set (the original dataset without bias). We then fit another Ridge model with the uniformly weighted biased source data and compute the MAE on the target set. We can then compute a score ratio for each method as follows:

$$\text{Score Ratio} = \frac{\sum_i |\beta_{IW}^T x'_i - y'_i|}{\sum_i |\beta_{Unif}^T x'_i - y'_i|} \quad (5)$$

where β_{IW}, β_{Unif} are respectively the coefficients of the Ridge models fitted with the importance weights and the uniform weights.




Fig. 2. Score ratio for the experiments with the sample bias applied on input features. In each column, the values correspond to the ratios between the mean absolute error (MAE) of the corresponding methods (IWN, NNW, KLIEP and KMM) and the MAE of the Uniform Weighting approach where no reweighting is performed. The MAEs are computed on the target dataset with a Ridge model fitted on the reweighted source data. The two first columns n and p are respectively the number of sample and number of features of the dataset. The lightest colors correspond to the best ratios. The experiment stopped after 500 seconds are marked with Nan.

We repeat each experiment 10 times and report the results of the experiments with the sample bias on the input features in Figures 2 and 3, the standard deviation over the 10 repetitions are given in appendix. The first Figure presents the score ratios computed with Eq (5). We observe that the quality of the weighting scheme provided by IWN is competitive with other methods. IWN is in the top 2 best ratios for all experiments except one. Even more impressive are the computational time of IWN compared to other methods (Figure 3), IWN provides a fast computation of no more than 25 seconds for samples of size $< 5 \times 10^5$. Whereas the other methods failed to provide importance weights in a reasonable time for samples above 5×10^4 samples and > 100 features. The same observations are made on Table 1 which reports the summary results of the experiments conducted on the same datasets with a sample bias on the outputs.

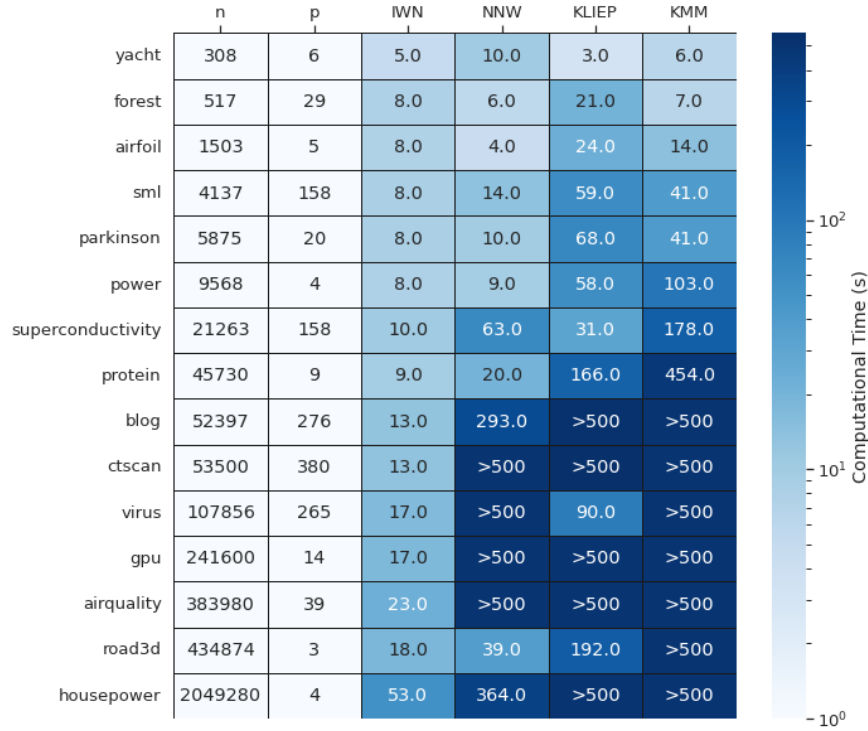


Fig. 3. Computational times (in second) for the experiments with the sample bias applied on input features. The two first columns n and p are respectively the number of sample and number of features of the dataset. The computational time are given in second. The lightest colors correspond to the lowest computational time. The experiment stopped after 500 seconds are marked (>500).

method	Avg Score Ratio	Avg Rank	Avg Comp. Time (in sec.)
IWN	0.9	2.0	8.06
NNW	0.92	2.37	182.81
KLIEP	1.03	2.93	223.16
KMM	0.87	2.7	290.01

Table 1. Summary of the results of the output sample bias experiments (extensive results are reported in appendix).

4.5 Impact of Network Architecture and Batch Size

Finally, we study the impact of the network architecture and the batch size on the solution of IWN. We conduct the experiments on the CTscan dataset⁷ [13] biased through the sample bias on the input features described previously. First, we fix the batch size to 256 and vary the number of hidden layers of the weighting network from 0 to 4 and the number of neurons per layer between [10, 100, 300]. Then, we fix the number of hidden layers to 3 and the number of neurons to 100 and vary the batch size on a geometric scale of ratio 4 from 16 to 4096. We repeat each experiment 10 times and report the means and standard deviations of the scores in Table 2. We observe that the architecture of the network has little impact on the score, however, we observe a slight improvement of the score between the simplest architectures and the more complex ones (0.88 for 0 hidden layer and 0.85 for 4 layers and 300 neurons). The impact of the batch size is more significant on the performance of IWN, enlarging the batches produce better corrections of sample bias. This is due to the more accurate estimation of the MMD made with larger batch. However, increasing the batch size comes with an increase of the computational time as shown in the last column of Table 2.

Neural Network Architecture				Batch Size		
Neurons \ Layers	10	100	300	Size	Score	Time (s)
0	0.88 (0.02)	0.88 (0.02)	0.88 (0.02)	16	0.96 (0.01)	9.7 (8.2)
1	0.86 (0.02)	0.86 (0.02)	0.85 (0.02)	64	0.90 (0.01)	6.6 (0.2)
2	0.86 (0.02)	0.86 (0.02)	0.85 (0.02)	256	0.85 (0.02)	7.4 (0.3)
3	0.86 (0.02)	0.85 (0.02)	0.85 (0.02)	1024	0.82 (0.03)	12.1 (0.1)
4	0.86 (0.02)	0.85 (0.02)	0.85 (0.02)	4096	0.83 (0.04)	117.9 (1.3)

Table 2. Summary of the study on the impact of weighting network architecture and batch size. Standard deviation over the 10 repetitions are given in brackets.

⁷ <https://archive.ics.uci.edu/ml/datasets/Relative+location+of+CT+slices+on+axial+axis>

5 Conclusion

This work introduces a novel algorithm for importance weighting called Importance Weighting Network and shows that sample biases can be efficiently corrected by fitting a weighting neural network with the MMD as loss function. This approach appears to provide very competitive results with state-of-the-art instance-based domain adaptation methods for a minimal cost in term of computational time.

References

1. Bellot, A., van der Schaar, M.: A kernel two-sample test with selection bias. In: *Uncertainty in Artificial Intelligence*. pp. 205–214. PMLR (2021)
2. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F.: Analysis of representations for domain adaptation. In: Schölkopf, B., Platt, J.C., Hoffman, T. (eds.) *Advances in Neural Information Processing Systems 19*, pp. 137–144. MIT Press (2007)
3. Bickel, S., Brückner, M., Scheffer, T.: Discriminative learning under covariate shift. *Journal of Machine Learning Research* **10**(9) (2009)
4. Cao, Z., Long, M., Wang, J., Jordan, M.I.: Partial transfer learning with selective adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2724–2732 (2018)
5. Cao, Z., You, K., Long, M., Wang, J., Yang, Q.: Learning to transfer examples for partial domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2985–2994 (2019)
6. Chandra, S., Haque, A., Khan, L., Aggarwal, C.: Efficient sampling-based kernel mean matching. In: *2016 IEEE 16th International Conference on Data Mining (ICDM)*. pp. 811–816. IEEE (2016)
7. Cortes, C., Mohri, M.: Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science* **519** (2014)
8. Deheeger, F., MOUGEOT, M., Vayatis, N., et al.: Handling distribution shift in tire design. In: *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications* (2021)
9. Diesendruck, M., Elenberg, E.R., Sen, R., Cole, G.W., Shakkottai, S., Williamson, S.A.: Importance weighted generative networks. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 249–265. Springer (2019)
10. Dua, D., Graff, C.: UCI machine learning repository (2017), <http://archive.ics.uci.edu/ml>
11. Fang, T., Lu, N., Niu, G., Sugiyama, M.: Rethinking importance weighting for deep learning under distribution shift. *Advances in Neural Information Processing Systems* **33**, 11996–12007 (2020)
12. Friedman, J.H., Bentley, J.L., Finkel, R.A.: An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software (TOMS)* **3**(3), 209–226 (1977)
13. Graf, F., Kriegel, H.P., Schubert, M., Pölsterl, S., Cavallaro, A.: 2d image registration in ct images using radial image descriptors. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 607–614. Springer (2011)

14. Gretton, A., Borgwardt, K.M., Rasch, M.J., Schölkopf, B., Smola, A.: A kernel two-sample test. *The Journal of Machine Learning Research* **13**(1), 723–773 (2012)
15. Guan, D., Huang, J., Xiao, A., Lu, S., Cao, Y.: Uncertainty-aware unsupervised domain adaptation in object detection. *IEEE Transactions on Multimedia* **24**, 2502–2514 (2021)
16. Huang, J., Gretton, A., Borgwardt, K., Schölkopf, B., Smola, A.J.: Correcting sample selection bias by unlabeled data. In: Schölkopf, B., Platt, J.C., Hoffman, T. (eds.) *Advances in Neural Information Processing Systems* 19, pp. 601–608. MIT Press (2007)
17. Kanamori, T., Hido, S., Sugiyama, M.: A least-squares approach to direct importance estimation. *The Journal of Machine Learning Research* **10**, 1391–1445 (2009)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Bengio, Y., LeCun, Y. (eds.) *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings* (2015)
19. LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *nature* **521**(7553), 436–444 (2015)
20. Li, C.L., Chang, W.C., Cheng, Y., Yang, Y., Póczos, B.: Mmd gan: Towards deeper understanding of moment matching network. *Advances in neural information processing systems* **30** (2017)
21. Loog, M.: Nearest neighbor-based importance weighting. In: *2012 IEEE International Workshop on Machine Learning for Signal Processing*. pp. 1–6. IEEE (2012)
22. Mansour, Y., Mohri, M., Rostamizadeh, A.: Domain adaptation: Learning bounds and algorithms. In: *COLT* (2009)
23. Martino, L., Elvira, V., Louzada, F.: Effective sample size for importance sampling based on discrepancy measures. *Signal Processing* **131**, 386–401 (2017)
24. de Mathelin, A., Deheeger, F., Richard, G., Mougeot, M., Vayatis, N.: Adapt: Awesome domain adaptation python toolbox. *arXiv preprint arXiv:2107.03049* (2021)
25. de Mathelin, A., Richard, G., Deheeger, F., Mougeot, M., Vayatis, N.: Adversarial weighting for domain adaptation in regression. In: *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*. pp. 49–56. IEEE (2021)
26. Miao, Y.Q., Farahat, A.K., Kamel, M.S.: Auto-tuning kernel mean matching. In: *2013 IEEE 13th International Conference on Data Mining Workshops*. pp. 560–567. IEEE (2013)
27. Miao, Y.Q., Farahat, A.K., Kamel, M.S.: Ensemble kernel mean matching. In: *2015 IEEE International Conference on Data Mining*. pp. 330–338. IEEE (2015)
28. Mohri, M., Muñoz Medina, A.: New analysis and algorithm for learning with drifting distributions. In: Bshouty, N.H., Stoltz, G., Vayatis, N., Zeugmann, T. (eds.) *Algorithmic Learning Theory*. pp. 124–138. Springer Berlin Heidelberg, Berlin, Heidelberg (2012)
29. Omohundro, S.M.: *Five balltree construction algorithms*. International Computer Science Institute Berkeley (1989)
30. Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* **22**(10), 1345–1359 (Oct 2010). <https://doi.org/10.1109/TKDE.2009.191>
31. Park, S., Bastani, O., Weimer, J., Lee, I.: Calibrated prediction with covariate shift via unsupervised domain adaptation. In: *International Conference on Artificial Intelligence and Statistics*. pp. 3219–3229. PMLR (2020)

32. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
33. Sugiyama, M., Nakajima, S., Kashima, H., Büna, P.v., Kawanabe, M.: Direct importance estimation with model selection and its application to covariate shift adaptation. In: *Proceedings of the 20th International Conference on Neural Information Processing Systems*. p. 1433–1440. NIPS’07, Curran Associates Inc., Red Hook, NY, USA (2007)
34. Wang, X., Long, M., Wang, J., Jordan, M.: Transferable calibration with lower bias and variance in domain adaptation. *Advances in Neural Information Processing Systems* **33**, 19212–19223 (2020)
35. Wen, J., Zheng, N., Yuan, J., Gong, Z., Chen, C.: Bayesian uncertainty matching for unsupervised domain adaptation. *arXiv preprint arXiv:1906.09693* (2019)
36. Yamada, M., Suzuki, T., Kanamori, T., Hachiya, H., Sugiyama, M.: Relative density-ratio estimation for robust distribution comparison. *Advances in neural information processing systems* **24** (2011)
37. You, K., Long, M., Cao, Z., Wang, J., Jordan, M.I.: Universal domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2720–2729 (2019)
38. Zhang, J., Ding, Z., Li, W., Ogunbona, P.: Importance weighted adversarial nets for partial domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 8156–8164 (2018)