



**HAL**  
open science

# Nearest Neighbor Sampling for Covariate Shift Adaptation

François Portier, Lionel Truquet, Ikko Yamane

► **To cite this version:**

François Portier, Lionel Truquet, Ikko Yamane. Nearest Neighbor Sampling for Covariate Shift Adaptation. 2024. hal-04645530

**HAL Id: hal-04645530**

**<https://hal.science/hal-04645530>**

Preprint submitted on 11 Jul 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.



Distributed under a Creative Commons Attribution 4.0 International License

---

# NEAREST NEIGHBOR SAMPLING FOR COVARIATE SHIFT ADAPTATION

---

A PREPRINT

**François Portier**

francois.portier@ensai.fr

Department of Statistics,

Univ Rennes, Ensai, CNRS, CREST—UMR 9194, F-35000 Rennes, France

**Lionel Truquet**

lionel.truquet@ensai.fr

Department of Statistics,

Univ Rennes, Ensai, CNRS, CREST—UMR 9194, F-35000 Rennes, France

**Ikko Yamane**

ikko.yamane@ensai.fr

Department of Computer Science,

Univ Rennes, Ensai, CNRS, CREST—UMR 9194, F-35000 Rennes, France

July 1, 2024

## ABSTRACT

Many existing covariate shift adaptation methods estimate sample weights given to loss values to mitigate the gap between the source and the target distribution. However, estimating the optimal weights typically involves computationally expensive matrix inversion and hyper-parameter tuning. In this paper, we propose a new covariate shift adaptation method which avoids estimating the weights. The basic idea is to directly work on unlabeled target data, labeled according to the  $k$ -nearest neighbors in the source dataset. Our analysis reveals that setting  $k = 1$  is an optimal choice. This property removes the necessity of tuning the only hyper-parameter  $k$  and leads to a running time quasi-linear in the sample size. Our results include sharp rates of convergence for our estimator, with a tight control of the mean square error and explicit constants. In particular, the variance of our estimators has the same rate of convergence as for standard parametric estimation despite their non-parametric nature. The proposed estimator shares similarities with some matching-based treatment effect estimators used, e.g., in biostatistics, econometrics, and epidemiology. Our experiments show that it achieves drastic reduction in the running time with remarkable accuracy.

## 1 Introduction

Traditional machine learning methods assume that the source data distribution  $P$  and the target data distribution  $Q$  are identical. However, this assumption can be violated in practice when there is a *distribution shift* (Chen et al., 2022) between them. Various types of shift have been studied in the literature, and one of the most common scenarios is *covariate shift* (Shimodaira, 2000) in which there is a shift in the input distribution:  $P_X \neq Q_X$  while the conditional distribution of the output variable given the input variable is the same:  $P_{Y|X} = Q_{Y|X}$ , where  $X$  is the input and  $Y$  is the output variable. The goal of *covariate shift adaptation* is to adapt a supervised learning algorithm to the target distribution using labeled source data and unlabeled target data.

A standard approach to covariate shift is weighting source examples (Shimodaira, 2000), and many studies focused on improving the weights (Huang et al., 2006; Gretton et al., 2008; Yamada et al., 2013; Kanamori et al., 2009;

Sugiyama et al., 2007, 2008; Aminian et al., 2022) in the same line of research. We refer the reader to Section 6 for more details of related work. Since we rarely know the model for how the input distributions can be shifted a priori, non-parametric methods are particularly useful for covariate shift adaptation. Some of the existing methods allow one to use non-parametric models through kernels. However, such kernel-based methods take at least quadratic times in computing kernel matrices. Some methods further need to solve linear systems and take cubic times in the sample size unless one resorts to approximations (Williams and Seeger, 2000; Le et al., 2013). Moreover, their performance is often sensitive to the choice of hyper-parameters of the kernel. Typically, one performs a grid search  $K$ -fold cross-validation for selecting the hyper-parameters, which amplifies the running time by about  $K|\Gamma|$ , where  $\Gamma$  is the set of candidates for the hyper-parameters. Moreover, the criterion for the hyper-parameter selection is not obvious either because we do not have access to the labels for the target data. One can use weighted validation scores using the labeled source data with importance sampling, but it is not straightforward to choose what weights to be used for the cross-validation when we are choosing weights.

In this paper, we propose a non-parametric covariate shift adaptation method that is scalable and has no hyper-parameter. Our idea is to generate synthetic labels for unlabeled target data using a non-parametric conditional sampler constructed from source data. Under the assumption of covariate shift, the target data attached with the generated labels behave like labeled target data. This sampling technique allows any supervised learning method to be simply applied to the generated data to produce a model already adapted to the target distribution.

While the proposed approach is quite general and can be employed with various sampling methods for the synthetic labeling part, our main result is that a  $k$ -nearest neighbor ( $k$ -NN) based sampling method achieves an error of order  $(k/n)^{1/d} + 1/\sqrt{n} + 1/\sqrt{m}$  for estimating an expectation on the target domain, where  $d$  is the data dimensionality, and  $n$  and  $m$  are the source and the target sample size, respectively. Importantly, our error bounds suggest that  $k = 1$  is the most favorable. This property, which is revealed by a precise scaling of the variance term in  $1/\sqrt{n}$ , is a non-trivial and remarkable result, given the  $1/\sqrt{k}$ -rate of the variance associated to the  $k$ -NN estimator of the conditional distribution (Portier, 2021, Corollary 1), and it contrasts with the well-known application of  $k$ -NN to standard density estimation (Dasgupta and Kpotufe, 2014), classification (Gadat et al., 2016; Cannings et al., 2020), or regression problems (Devroye et al., 1994; Jiang, 2019), in which we typically need to let  $k$  grow in a polynomial rate in the sample size in order to achieve a good balance in the bias-variance trade-off. This important difference in the rate of convergence, leading to a  $k = 1$  number of neighbor, has also been noticed in other estimation problems such as the  $k$ -NN entropy estimator (Berrett et al., 2019) or the integral approximation problem (Leluc et al., 2023; Blanchet et al., 2024). Textbooks dealing with the  $k$ -NN algorithm include (Györfi et al., 2006; Devroye et al., 2013; Biau and Devroye, 2015).

In addition of being optimal with respect to the estimation error, setting  $k = 1$  circumvent the cumbersome hyper-parameter tuning while providing computational efficiency at the same time. Our 1-NN-based algorithm takes only a quasi-linear time  $\mathcal{O}((n + m) \log n)$  on average using the optimized  $k$ -d tree (Bentley, 1975; Friedman et al., 1977). Indeed, our experiments show that the proposed method terminates faster than previous methods, by large margins. Note that the problem of getting a computationally efficient method for covariate shift adaptation, in particular for scalability to large data sets, is a recurrent problem in the existing literature. In fact, many existing methods resorted to implementation heuristics such as using a fixed number of kernel centers for reducing the computational burden at the cost of statistical guarantee (Kanamori et al., 2009; Sugiyama et al., 2007, 2008; Yamada et al., 2013).

Our method simulates the missing labels of the target sample, which in turn can be used for a variety of downstream supervised learning tasks. Even though the main focus of this paper is the estimation of expectations in the target domain, for illustrating the usefulness of our method in a typical machine learning downstream task, we also demonstrate consistency properties of parametric M-estimators in the target domain. This is particularly useful when the parametric model is misspecified, since in this case, even the population minimizer changes when the covariate distribution is shifted.

The problem of interest here is closely related to a well-known matching problem studied in the context of treatment effect estimation. In particular,  $k$ -NN estimators have been used to estimate the so-called average treatment effect to tackle missingness of potential outcomes. See, e.g., Rosenbaum (1995); Abadie and Imbens (2006), for an error bound obtained in this specific instance of the problem. In Section 6, we discuss the main differences between the two problems and why their result is not generally applicable to ours.

In summary, the key contributions of this paper are the following. (i) Our method is non-parametric. It does not introduce a model in covariate shift adaptation so that it will have a minimum impact on the model trained for the downstream task. (ii) Our method is fast. Adaptation only takes a quasi-linear time. (iii) There is no hyper-parameter to be tuned. (iv) The proposed method only incurs an error of order  $(k/n)^{1/d} + 1/\sqrt{n} + 1/\sqrt{m}$  for estimating an expectation on the target domain.

The outline is as follows. In Section 2, the problem of covariate shift adaptation is formally introduced along with the mathematical notation. Section 3 contains the description of the method. Section 4 is dedicated to the main theoretical results while Section 5 investigates the empirical risk minimization problem in presence of covariate shift adaptation. Section 6 provides a description of several alternative approaches to a similar type of problem as well as some points of comparison with our proposal. In Section 7, several avenues for further research are discussed and finally, the numerical experiments are provided in Section 8.

## 2 Problem setup

Let  $\mathcal{X}$  and  $\mathcal{Y}$  be measurable spaces. Let  $P \equiv P_{X,Y}$  and  $Q \equiv Q_{X,Y}$  be probability distributions defined on  $\mathcal{X} \times \mathcal{Y}$ . Throughout the paper, we assume that  $P$  and  $Q$  admit the decomposition

$$P = P_{Y|X}P_X \quad \text{and} \quad Q = Q_{Y|X}Q_X,$$

where  $P_{Y|X=x}$  and  $Q_{Y|X=x}$  are probability distributions defined on  $\mathcal{Y}$  for each  $x \in \mathcal{X}$ .<sup>1</sup> Here,  $P_X$  and  $Q_X$  are the marginal distributions of  $X$  when  $(X, Y)$  is distributed with  $P$  and  $Q$ , respectively. We shall simply call  $P_{Y|X}$  (or  $Q_{Y|X}$ ) the *conditional distribution* of  $Y$  given  $X$  in the source domain (or the target domain).

**Definition 1** (Source sample, source distribution). *For each integer  $n \geq 1$ , let  $(X_i, Y_i)_{i=1}^n$  be a collection of independent and identically distributed random variables with  $P$ . We refer to  $(X_i, Y_i)_{i=1}^n$  as the (labeled) source sample and  $P$  as the source distribution.*

**Definition 2** (Target sample, target distribution). *For each integer  $m \geq 1$ , let  $(X_i^*)_{i=1}^m$  be a collection of independent and identically distributed random variables with  $Q$ . We refer to  $(X_i^*)_{i=1}^m$  as the (unlabeled) target sample and  $Q$  as the target distribution.*

**Definition 3** (Covariate shift). *Covariate shift is a situation in which the source and the target distribution have different marginal distributions for  $X$  while sharing a common conditional distribution:*

$$(C1) \quad P_{Y|X} = Q_{Y|X}, \quad P_X\text{- and } Q_X\text{-a.s.}, \quad \text{but } P_X \neq Q_X.$$

This paper focuses on the following simple but versatile estimation problem under covariate shift.

**Definition 4** (Mean estimation under covariate shift). *For each pair of integers  $n \geq 1$ ,  $m \geq 1$ , and a known integrable function  $h: \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ , the goal of mean estimation under covariate shift is to estimate the mean of  $h$  under the target distribution,*

$$Q(h) \equiv \int h(x, y)Q(dx, dy),$$

*given access to the source sample  $(X_i, Y_i)_{i=1}^n \sim P$  and the target sample  $(X_i^*)_{i=1}^m \sim Q_X$  under Assumption (C1).*

For instance, when  $h(x, y) = \ell(f(x), y)$  for a loss function  $\ell: \mathcal{Y}^2 \rightarrow \mathbb{R}$  and a hypothesis function  $f: \mathcal{X} \rightarrow \mathcal{Y}$ , estimation of  $Q(h)$  becomes risk estimation, which is the central subtask in *empirical risk minimization*.

## 3 Proposed method

The basic idea of our proposed method is to use the source sample for learning to *label the target data*. Specifically, using the source sample  $(X_i, Y_i)_{i=1}^n$ , we will construct a stochastic labeling function  $\hat{S}$  that inputs any target data point  $X_i^*$  and outputs a random label  $Y_{n,i}^* \in \mathcal{Y}$ . (The subscript  $n$  of  $Y_{n,i}^*$  is for explicitly denoting the dependence on the source sample.) Once we succeed in generating labels for target data that behave like true target labels, we will be able to perform any supervised learning method *directly on the target sample* for the downstream task. For our mean estimation problem, we can simply average the output  $h$  evaluated at the target data with the generated labels.

When do the generated labels behave like the true target labels? Let  $\hat{P}_{Y|X_i^*}$  denote the probability distribution of an output  $Y_{n,i}^*$  of  $\hat{S}$  for input  $X_i^*$ . We wish to obtain  $\hat{S}$  such that the probability distribution  $\hat{Q} \equiv \hat{P}_{Y|X}Q_X$  of  $(X_i^*, Y_{n,i}^*)$  will be a good estimate of  $Q = Q_{Y|X}Q_X$ . For this, we want  $\hat{P}_{Y|X_i^*}$  to be a good estimate of  $P_{Y|X_i^*}$ . In fact, if  $\hat{P}_{Y|X_i^*} = P_{Y|X_i^*}$ , the generated sample  $(X_i^*, Y_{n,i}^*)$  will follow the target distribution  $Q \equiv Q_{Y|X}Q_X = P_{Y|X}Q_X$  under Assumption (C1). In this sense, our task boils down to designing a good conditional sampler  $\hat{S}$  mimicking sampling from  $P_{Y|X}$ . Algorithm 1 describes an outline of this general framework.

<sup>1</sup>More formally, we denote by  $P_{Y|X=(\cdot)}(dy)$  a *regular conditional measure* (Bogachev and Ruas, 2007, Definition 10.4.1) such that the marginal distribution of  $Y$  can be expressed as  $P_Y(dy) = \int P_{Y|X=x}(dy)P_X(dx)$ . We also use  $P_{Y|X}(dy|\cdot)$  for  $P_{Y|X=(\cdot)}(dy)$ . The same goes for  $Q$ .

**Algorithm 1** Conditional Sampling Adaptation

---

**Input:** Conditional sampler  $\hat{S}$  and target sample  $(X_j^*)_{j=1}^m$ .  
 $Y_{n,j}^* \leftarrow \hat{S}(X_j^*)$  for each  $j \in \{1, \dots, m\}$ . // Generate a label conditioned on  $X_j^*$ .  
**return**  $m^{-1} \sum_{j=1}^m h(X_j^*, Y_{n,j}^*)$ .

---

In this paper, we propose a method using a non-parametric conditional sampler  $\hat{S}$  based on the  $k$ -Nearest Neighbor ( $k$ -NN) method, which randomly picks one of the  $k$ -nearest neighbors of the input  $X_j^*$  among the source instances  $(X_i)_{i=1}^n$  and output the corresponding label (Algorithm 2). We refer to this method as  $k$ -NN-based Conditional Sampling Adaptation ( $k$ -NN-CSA).

**Algorithm 2**  $k$ -Nearest Neighbor Conditional Sampler

---

**Input:** Source sample  $(X_i, Y_i)_{i=1}^n$  and target input  $X_j^*$ .  
 $(i_1, \dots, i_k) \leftarrow$  the indices of the  $k$ -nearest neighbors of  $X_j^*$  among the source instances  $(X_i)_{i=1}^n$ .  
Pick  $i^* \in \{i_1, \dots, i_k\}$  uniformly at random.  
**return**  $Y_{n,j}^* := Y_{i^*}^*$ .

---

**Computing time** Recent advances for nearest neighbor search rely on tree-search to reduce the computing time. The seminal paper by Bentley (1975) introduced the  $k$ -d tree method. Building such a tree requires  $\mathcal{O}(n \log n)$  and once the tree is available, search for the nearest neighbor of a given point can be done in  $\mathcal{O}(\log n)$  time (Friedman et al., 1977). As a consequence, the time complexity of  $k$ -NN-CSA is  $\mathcal{O}(n \log n + km \log n)$ .

## 4 Theoretical analysis

We now present the theory behind our approach in a didactic way by introducing a key decomposition first and then studying separately each of the terms involved: the sampling error and the  $k$ -NN conditional sampling error. We will see that the  $k$ -NN-CSA with  $k = 1$  (1-NN-CSA for short) achieves the best theoretical performance among those with other  $k$ 's.

### 4.1 The key decomposition

For the analysis of  $k$ -NN-CSA, recall that  $\hat{Q} = \hat{P}_{Y|X} Q_X$  is an estimate of the target distribution  $Q = P_{Y|X} Q_X$  that depends on the source sample  $(X_i, Y_i)_{i=1}^n$ , whose probability distribution is  $P$ . We introduce the bootstrap sample as a collection of random variable generated according to  $\hat{Q}$ .

**Definition 5** (Bootstrap sample). For each  $m \geq 1$  and  $n \geq 1$ , let  $(X_i^*, Y_{n,i}^*)_{1 \leq i \leq m}$  be a collection of random variables identically distributed with  $\hat{Q}$  and conditionally independent given  $(X_i, Y_i)_{i=1}^n$ .

Let  $h : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$  be a measurable function. The quantity of interest is

$$\hat{Q}^*(h) = m^{-1} \sum_{i=1}^m h(X_i^*, Y_{n,i}^*),$$

which is the CSA estimate of  $Q(h) = \int h(x, y) Q(dx, dy)$  as introduced in Algorithm 1. The following decomposition is crucial in our analysis:

$$\begin{aligned} (\hat{Q}^* - Q)(h) &= \underbrace{(\hat{Q}^* - \hat{Q})(h)}_{\text{Marginal sampling error}} + \underbrace{(\hat{Q} - Q)(h)}_{\text{Conditional sampling error}} \\ &= \left( (\hat{Q}_X^* - Q_X) \hat{P}_{Y|X}(h) + Q_X (\hat{P}_{Y|X} - P_{Y|X})(h) \right), \end{aligned} \quad (1)$$

where  $\hat{Q}_X^*(\cdot) \equiv \frac{1}{n} \sum_{i=1}^m \mathbb{1}_{X_i^* = (\cdot)}$  is the empirical measure defined with  $(X_i^*)_{i=1}^m$ . The first term is the error due to the use of  $\hat{Q}_X^*$  in place of  $Q_X$ , which tends to zero as  $m$  grows. The second term represents the error due to the use of  $\hat{P}_{Y|X}$  in place of  $P_{Y|X}$ . When using the  $k$ -nearest neighbor algorithm to obtain  $\hat{P}_{Y|X}$ , we show that this term is of order  $(k/n)^{1/d} + 1/\sqrt{n}$ , which differs from the standard non-parametric convergence rate in  $(k/n)^{1/d} + 1/\sqrt{k}$  found in regression problems.

## 4.2 Marginal sampling error

First, we will show that the marginal sampling error,  $(\hat{Q}^* - \hat{Q})(h)$  (the first term in our decomposition (1)), is of order  $1/\sqrt{m}$ . The analysis relies on martingale tools. Define  $\mathcal{F}_n = \sigma((X_1, Y_1), \dots, (X_n, Y_n))$ . For each  $1 \leq i \leq m$ , we have

$$\mathbb{E}[h(X_i^*, Y_{n,i}^*) | \mathcal{F}_n] = \hat{Q}(h).$$

This property implies that  $\sum_{i=1}^m \{h(X_i^*, Y_{n,i}^*) - \int h(x, y) \hat{Q}(dx, dy)\}$  is a martingale and therefore can be analyzed using the Lindeberg-CLT conditionally on the initial sample hence fixing the distribution  $\hat{Q}$ . The next property is reminiscent of certain results about the bootstrap method where sampling is done with the basic empirical measure, see e.g., Van der Vaart (2000). We need this type of results without specifying the measure  $\hat{Q}$  so that we can incorporate a variety of sampling schemes such as  $\hat{Q} = \hat{P}_{Y|X} Q_X$ . The proof is given in Appendix B.1.

**Proposition 1.** *Suppose that  $\hat{Q}$  satisfies the following strong law of large number: for each  $h$  such that  $Q(h) < \infty$ , we have  $\lim_{n \rightarrow \infty} \hat{Q}(h) = Q(h)$  almost surely. Then, if  $m := m_n \rightarrow \infty$  as  $n \rightarrow \infty$ , we have the following central limit theorem: for each function such that  $Q(h^2) < \infty$ , we have, conditionally to  $\mathcal{F}_n$ , almost surely,*

$$\sqrt{m}\{\hat{Q}^*(h) - \hat{Q}(h)\} \rightsquigarrow \mathcal{N}(0, V) \quad \text{as } n \rightarrow \infty,$$

where  $V = \lim_{n \rightarrow \infty} \{\hat{Q}(h^2) - \hat{Q}(h)^2\}$ .

As a corollary of the previous results, we can already deduce that if  $m$  goes to  $\infty$  and  $\hat{Q}$  satisfies a strong law of large numbers, then  $\hat{Q}^*(h)$  converges to  $Q(h)$  provided that  $Q(h^2)$  exists. This is a general consistency result that justifies the use of any resampling distribution  $\hat{Q}$  that converges to  $Q$ . In practical situations, it is useful to know a finite-sample bound on the error. This is the purpose of the next proposition, in which we give a non-asymptotic control of the sampling error. A proof is given in Appendix B.2.

**Proposition 2.** *Suppose that  $h$  is bounded by a constant  $U_h > 0$ . Let  $\delta \in (0, 1)$ . Then with probability greater than  $1 - \delta$ ,*

$$\left| \hat{Q}^*(h) - \hat{Q}(h) \right| \leq \frac{U_h}{m} \log(2/\delta) + \sqrt{2 \frac{\hat{v}_n}{m} \log(2/\delta)},$$

where  $\hat{v}_n = \hat{Q}(h^2) - (\hat{Q}(h))^2$ .

**Notes.** A natural ‘‘averaging’’ alternative to the above ‘‘sampling’’ estimator  $\hat{Q}^*$  can also be investigated using the same tools. Instead of sampling  $Y_{n,i}^*$  according to  $\hat{P}_n(dy|X_i^*)$ , one might consider taking the expectation, leading to

$$\bar{Q}(h) = \frac{1}{m} \sum_{i=1}^m \int h(X_i^*, y) \hat{P}_n(dy|X_i^*).$$

This estimate can be studied in a similar way as before and the two above results are still valid with small changes. In particular Proposition 2 holds true with smaller variance term as, by Jensen’s inequality,  $Q_X(\int h(X_i^*, y) \hat{P}_n(dy|X_i^*))^2 \leq \hat{Q}(h^2)$ . This alternative  $\bar{Q}(h)$  requires more computing time (when measured in terms of evaluation of  $h$ ) and is less appealing for stochastic gradient descent algorithm or in semiparametric estimation problems, as discussed in Section 7. Estimators similar to  $\bar{Q}(h)$  have been studied in average treatment effects literature (Rosenbaum, 1995; Abadie and Imbens, 2006); see Section 6 for precise discussion.

## 4.3 Conditional sampling error of the nearest neighbor estimate

Our aim in this section is to obtain a bound on  $\hat{Q}(h) - Q(h)$  (the second term in our decomposition (1)) when  $\hat{P}_{Y|X}$  is the  $k$ -nearest neighbor measure.

Let  $x \in \mathbb{R}^d$  and  $\|\cdot\|$  be the Euclidean norm on  $\mathbb{R}^d$ . Denote the closed ball of radius  $\tau \geq 0$  around  $x$  by  $B(x, \tau) := \{z \in \mathbb{R}^d \mid \|x - z\| \leq \tau\}$ . For  $n \geq 1$  and  $k \in \{1, \dots, n\}$ , the  $k$ -nearest neighbor ( $k$ -NN for short) radius at  $x$  is denoted by  $\hat{\tau}_{n,k,x}$  and defined as the smallest radius  $\tau \geq 0$  such that the ball  $B(x, \tau)$  contains at least  $k$  points from the collection  $\{X_1, \dots, X_n\}$ . That is,

$$\hat{\tau}_{n,k,x} := \inf \left\{ \tau \geq 0 : \sum_{i=1}^n 1_{B(x,\tau)}(X_i) \geq k \right\},$$

where  $1_A(x)$  is 1 if  $x \in A$  and 0 elsewhere. The  $k$ -NN estimate of  $P_{Y|X}(dy|x)$  is given by

$$\hat{P}_{Y|X}(dy|x) = k^{-1} \sum_{i=1}^n 1_{\|X_i - x\| \leq \hat{r}_{n,k,x}} \delta_{Y_i}(dy),$$

where  $\delta_y(\cdot)$  is the Dirac measure at  $y \in \mathcal{Y}$  defined by  $\delta_y(A) = 1_A(y)$  for any measurable set  $A \subseteq \mathcal{Y}$ . Consequently, the  $k$ -NN estimate of the integral  $\int h(y,x) P_{Y|X}(dy|x)$  is then defined as

$$\int h(y,x) \hat{P}_{Y|X}(dy|x) = k^{-1} \sum_{i=1}^n 1_{\|X_i - x\| \leq \hat{r}_{n,k,x}} h(Y_i, x).$$

To obtain some guarantee on the behavior of the nearest neighbors estimate, we consider the case in which covariates  $X$  admit a density with respect to the Lebesgue measure. We will need in addition that the support  $S_X$  is well shaped and that the density is lower bounded. These are standard regularity conditions to obtain some upper bound on the  $k$ -NN radius.

(X1) The random variable  $X$  admits a density  $p_X$  with compact support  $S_X \subset \mathbb{R}^d$ .

(X2) There is  $c > 0$  and  $T > 0$  such that

$$\lambda(S_X \cap B(x, \tau)) \geq c\lambda(B(x, \tau)), \quad \forall \tau \in (0, T], \forall x \in S_X,$$

where  $\lambda$  is the Lebesgue measure.

(X3) There is  $0 < b_X \leq U_X < +\infty$  such that  $b_X \leq p_X(x) \leq U_X$ , for all  $x \in S_X$ .

To obtain our main result, on the estimation property of the  $k$ -NN measure, we need some assumptions on the target measure  $Q_X$ .

(X4) The probability measure  $Q_X$  admits a bounded density  $q_X$  with support  $S_X$ . We will take  $U_X$  large enough such that it will also be an upper bound of  $q_X$ .

Two additional assumptions, different from the one before about  $X$ , will be needed to deal with the function  $h$  and the probability distribution of  $(Y, X)$ .

(H1) For any  $x$  in  $S_X$ ,

$$|\mathbb{E}[h(Y, x) | X = x] - \mathbb{E}[h(Y, x) | X = x + u]| \leq g_h(x) \|u\|$$

with  $\int g_h^2(x) Q_X(dx) < \infty$ .

(H2) There exists  $\sigma_+^2 > 0$  such that  $\sup_{x \in S_X} \text{Var}(h(Y, x) | X) \leq \sigma_+^2$  a.s., where  $\text{Var}(h(Y, x) | X)$  is the conditional variance of  $h(Y, x)$  given  $X$ .

In what follows, we give a control of the RMSE of  $\hat{Q}h$ . Let  $\|X\|_2 = \sqrt{\mathbb{E}(X^2)}$ ,  $[x]$  the integer part of a real number  $x$  and let  $\Gamma(x) := \int_0^\infty u^{x-1} \exp(-u) du$  for  $x > 0$ . Finally, we denote by  $V_d := \lambda(B(0, 1)) = \frac{\pi^{d/2}}{\Gamma(d/2+1)}$  the volume of the unit Euclidean ball in dimension  $d$  for the Lebesgue measure.

We give an upper-bound for the RMSE with explicit constants with respect to the dimension  $d$ . Additionally, we give a lower bound for the variance which has a standard parametric rate of convergence. The proof is given in Appendix C.1.

**Proposition 3.** *Suppose that Assumptions (X1), (X2), (X3), (X4), (H1), and (H2) are fulfilled. We have*

$$\hat{Q}h - Qh = S_h + B_h,$$

where  $B_h$  is a bias term (defined in the proof) that satisfies, for any  $n \geq 1$ ,

$$\mathbb{E}|B_h|^2 \leq \frac{2\Gamma(1 + [2/d])}{M_{1,d}^{2/d}} \int g_h^2(x) Q_X(dx) \cdot \frac{k^{2/d}}{n^{2/d}},$$

and  $S_h$  is a variance term (defined in the proof) that satisfies, for any  $n \geq 2$ ,

$$\frac{\sigma_-^2 M_{1,d}^2}{4M_{2,d}^2} n^{-1} \leq \min_{1 \leq k \leq n} \mathbb{E}[S_n^2] \leq \max_{1 \leq k \leq n} \mathbb{E}[S_n^2] \leq \frac{2^{d+3} \sigma_+^2 M_{2,d}^2}{M_{1,d}^2} n^{-1}.$$

For the lower bound to be true, it is assumed that the mapping  $h$  does not depend on  $x$ , i.e.  $h(y, x) = h(y)$  and  $\sigma_-^2 = \inf_{x \in S_X} \text{Var}(h(Y) | X = x)$ .

**Notes.** (i) The two terms  $B_h$  and  $S_h$  correspond respectively to the bias term and the variance term. The upper bound obtained for the bias term is usual in  $k$ -NN regression analysis. However, the upper and lower bound on the variance are particular to our framework as they show that the variance behaves as in usual parametric estimation. Consequently, our rates of convergence are sharper than the optimal rate of convergence  $n^{-\frac{1}{2+d}}$  for nonparametric estimation of Lipschitz functions. This can be explained by the fact that several  $k$ -NN estimators are averaged to estimate  $Qh$ , which is a standard expectation and not a conditional expectation.

(ii) Since the rate of convergence of the variance term  $S_h$  does not depend on  $k$ ,  $k$  might be chosen according to the upper bound on the bias term, which gives  $k = 1$ . One can deduce the following convergence rates, depending on the dimension. For  $d = 1$ , we get the rate  $n^{-1/2}$ . For  $d = 2$ , the contributions of both terms,  $B_h$  and  $S_h$ , coincide and we get the rate  $n^{-1/2}$ . For  $d \geq 3$ , the rate is  $n^{-1/d}$ .

For the global mean square error which incorporate the marginal sampling error as well as the  $k$ -NN conditional sampling error, we give the following result in the optimal case  $k = 1$ . The proof can be found in Appendix D.

**Theorem 1.** *Suppose that Assumptions (X1), (X2), (X3), (X4) and (H1) hold true with  $\sup_{x \in S_X} \mathbb{E} [h^2(Y, x)|X]$  bounded. If  $k = 1$ , there then exists  $C > 0$  only depending on the distribution of  $(X, Y)$ ,  $X^*$ , and on  $h$  such that*

$$\mathbb{E} \left| \hat{Q}^*(h) - Q(h) \right|^2 \leq C \left\{ \frac{1}{m} + \frac{1}{n^{2/d}} + \frac{1}{n} \right\}.$$

We next give a non-asymptotic control of  $\hat{Q}h - Qh$  when  $h$  is a bounded function using Bernstein's concentration inequality. This bound affords a complement with respect to the bound for the MSE. However, for technical reasons, this high-probability bound requires that  $k$  grows at least logarithmically with respect to  $n$ , in contrast to Proposition 3. In our numerical experiments, we will also include the case  $k = \log n$  for comparison. The proof of the next result is given in Appendix C.2.

**Proposition 4.** *Suppose that Assumptions (X1), (X2), (X3), (X4), (H1), and (H2) are fulfilled. Suppose that there exists a constant  $C > 0$  such that  $C \log n \leq k \leq n/2$  and that  $h$  is bounded by  $U_h$ . Let  $\delta \in (0, 1/3)$ . With probability greater than  $1 - 3\delta$ , we have*

$$\left| \hat{Q}h - Qh \right| \leq L_0 \left( \frac{k}{n} \right)^{1/d} + \frac{2L_2}{n} \log(2/\delta) + \sqrt{\frac{2L_1\sigma^2}{n} \log(2/\delta)},$$

where

$$L_0 = \left( \frac{2}{cb_x V_d} \right)^{1/d} \int g_h(x) Q_X(dx), \quad L_1 = \frac{4\sigma_+^2 U_X^2}{b_X^2 c^2}, \quad L_2 = \frac{4U_h U_X}{3b_X c}.$$

**Notes.** (i) The proof needs a bound on  $\sup_{x \in S_X} \hat{\tau}_{n,k,x}$  which is given in (Portier, 2021, Lemma 4). For this we need that  $k$  grows logarithmically w.r.t.  $n$  as stated in the assumptions.

(ii) The sum of the two last terms in the upper-bound, which corresponds to the variance of our estimator, is of order  $1/\sqrt{n}$  and the conditional variance of  $h$  appears as a multiplicative factor. Combining Proposition 2 and Proposition 4, we finally obtain that  $\hat{Q}^*(h) - Q(h)$  is of order  $1/\sqrt{m} + 1/\sqrt{n} + (k/n)^{1/d}$  (up to log factors).

(iii) In Corollary 1 in Appendix E, the upper bound given in Proposition 2 is refined using Proposition 4 that allows to (roughly speaking) replace the empirical variance by the true variance.

## 5 Applications to empirical risk minimization

In this section, we illustrate our results with some applications to empirical risk minimization. This is of particular interest in our context as the optimal linear model for the source distribution might be different from the ideal linear model for the target. In such a case, using covariate adaptation is always better as the source minimizer will be away from the target minimizer.

### 5.1 Mathematical background

Suppose that  $\mathcal{R}_{m,n}^*(\theta) = \frac{1}{m} \sum_{i=1}^m m_\theta(Y_{n,i}^*, X_i^*)$ , where for each  $\theta \in \Theta \subset \mathbb{R}^d$ ,  $m_\theta$  is a measurable function from  $\mathbb{R} \times \mathbb{R}^p$  to  $\mathbb{R}$ . Set

$$\hat{\theta}^* \in \arg \min_{\theta \in \Theta} \mathcal{R}_{m,n}^*(\theta).$$



Similarly, we define

$$\theta^* := \arg \min_{\theta \in \Theta} \mathcal{R}^*(\theta)$$

with  $\mathcal{R}^*(\theta) = \mathbb{E} m_\theta(Y^*, X^*)$  and  $(Y^*, X^*)$  is a copy of  $(Y_i^*, X_i^*)$ . Note that the expected value is taken for the unobserved label  $Y_i^*$  and not the generated label  $Y_{n,i}^*$ . We assume here that for a reference measure  $\mu$  on  $\mathcal{Y}$ , there exists for each  $x \in \mathcal{X}$  a conditional density  $p(\cdot|x)$  such that  $(x, y) \mapsto p(y|x)$  is jointly measurable and for any Borel set  $B$ ,

$$\mathbb{P}(Y \in B|X = x) = \int_B p(y|x)\mu(dy).$$

One can then include the case of classification ( $\mu = \delta_0 + \delta_1$ ), counts ( $\mu$  is the counting measure on the set of nonnegative integers) or regression ( $\mu$  is the Lebesgue measure on  $\mathbb{R}$ ).

## 5.2 Consistency of general empirical risk minimizers

We will use the following assumptions.

(A1) There exist a measurable function  $h : \mathbb{R} \rightarrow \mathbb{R}_+$  and  $\eta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  satisfying

$$\begin{aligned} \sup_{x \in \mathcal{X}} \sup_{\theta \in \Theta} |m_\theta(y, x)| &\leq h(y), \\ \sup_{x \in \mathcal{X}} \sup_{|\theta - \theta'| \leq \delta} |m_\theta(y, x) - m_{\theta'}(y, x)| &\leq h(y)\eta(\delta), \end{aligned}$$

and such that  $\mathbb{E}[h(Y)^2|X]$  is a bounded random variable and  $\lim_{\delta \rightarrow 0} \eta(\delta) = 0$ .

(A2) There exists a measurable function  $g_h : \mathcal{X} \rightarrow \mathbb{R}_+$  such that  $\int g_h(x)Q_X(dx) < \infty$  and

$$\int h(y) |p(y|x+u) - p(y|x)| \mu(dy) \leq g_h(x)|u|, \quad (x, x+u) \in S_X^2.$$

The above assumptions are satisfied, for instance, in the logistic regression framework with compact covariates. In this case,  $h$  is a constant function and  $\eta(\delta) = \delta$ . Note also that  $p(1|x)$  could be different from  $(1 + \exp(-x^T\theta))^{-1}$  as soon as  $x \mapsto p(1|x) \in (0, 1)$  is Lipschitz on  $S_X$ .

In what follows, an assertion of the form  $X_{m,n} = o_{\mathbb{P}}(1)$  as  $m, n \rightarrow \infty$  means that for any  $\epsilon, \zeta > 0$ , there exists  $A > 0$  such that

$$\min(m, n) \geq A \Rightarrow \mathbb{P}(|X_{m,n}| > \epsilon) \leq \zeta.$$

Additionally, the assertion  $X_{m,n} = O_{\mathbb{P}}(1)$  means that for any  $\epsilon > 0$ , there exist  $A, M > 0$  such that

$$\sup_{m, n \geq A} \mathbb{P}(|X_{m,n}| > M) \leq \epsilon.$$

The proof of the following result is in Appendix F.1.

**Theorem 2.** *Suppose that Assumptions (X1), (X2), (X3), (X4), and (A1), (A2) hold true with a compact subset  $\Theta$  of  $\mathbb{R}^d$  and the unique minimizer  $\theta^*$  of  $M$ . Then  $\hat{\theta}^* - \theta^* = o_{\mathbb{P}}(1)$  as  $m, n \rightarrow \infty$ . Moreover, the excess risk satisfies  $\mathcal{R}^*(\hat{\theta}^*) - \mathcal{R}^*(\theta^*) = o_{\mathbb{P}}(1)$ .*

## 5.3 Convergence rate for linear least-squares estimators

We now illustrate our results with an upper-bound on the excess risk for linear least-squares estimators in the misspecified case. Here, the targeted risk is given by

$$\mathcal{R}^*(\theta) = \mathbb{E} \left[ \left( Y^* - X^{*T} \theta \right)^2 \right]$$

and any optimal linear rule should simply be satisfied:

$$\theta^* \in \arg \min_{\theta \in \mathbb{R}^d} \mathcal{R}^*(\theta).$$

Note that  $\theta^*$  is unique the matrix  $\mathbb{E}[X^* X^{*T}]$  is of full rank. The empirical risk is defined by

$$\mathcal{R}_m^*(\theta) = \frac{1}{m} \sum_{i=1}^m \left( Y_{n,i}^* - X_i^{*T} \theta \right)^2$$

and  $\hat{\theta}^*$ , the empirical risk minimizer, is given by

$$\hat{\theta}^* \in \arg \min_{\theta \in \mathbb{R}^d} \mathcal{R}_m^*(\theta).$$

The excess risk satisfies the following upper bound whose proof is given in Appendix F.1.

**Theorem 3.** *Suppose that Assumptions (X1), (X2), (X3), (X4) hold true. Suppose that the mapping  $x \mapsto \mathbb{E}[Y|X=x]$  is Lipschitz and that the conditional expectation  $\mathbb{E}[Y^2|X]$  is bounded. Suppose also that  $\Gamma = \mathbb{E}[X^*X^{*T}]$  is positive definite. Then, we have*

$$\mathcal{R}^*(\hat{\theta}^*) - \mathcal{R}^*(\theta^*) = O_{\mathbb{P}}\left(m^{-1} + n^{-1} + n^{-2/d}\right).$$

**Notes.** The assumptions do not require the linear model for the  $(X_i, Y_i)$ 's to be valid, i.e., one can consider cases where  $\mathbb{E}[Y|X]$  is not linear. Also, when the source data follows a non-linear model of the form  $Y = r(X) + \varepsilon$  where  $\varepsilon$  and  $X$  are independent, our regularity assumptions means that  $r$  is Lipschitz on the compact set  $S_X$ .

## 6 Related work

A standard approach to covariate shift problems is to use some re-weighting in order to “transfer” the source distribution with density  $p_X$  to the target distribution with density  $q_X$ . This approach relies on the following type of estimates:

$$\hat{Q}_w(h) = n^{-1} \sum_{i=1}^n w(X_i) h(X_i, Y_i),$$

where ideally the function  $w$  would take the form  $q_X/p_X$ . Such a choice has the nice property that the expected value  $\mathbb{E}[w(X_i)h(X_i, Y_i)]$  is equal to the targeted quantity  $Q(h)$ . This however cannot be directly computed as  $p_X$  and  $q_X$  are unknown in practice. There are actually different ways to estimate  $w$ , and our goal here is to distinguish between two leading approaches.

**Plug-in approach** The plug-in approach is when the weights are computed using two estimates  $\hat{p}_X$  and  $\hat{q}_X$  in place of  $p_X$  and  $q_X$ , respectively; i.e., simply use  $\hat{w} = \hat{q}_X/\hat{p}_X$  instead  $w$  in the above formula, see for instance (Shimodaira, 2000; Sugiyama et al., 2007, 2008). Note that the selection of hyper-parameters for  $\hat{q}_X$  and  $\hat{p}_X$  is needed and the  $n$  evaluation ( $\hat{q}_X(X_i), \hat{p}_X(X_i)$ ) might be heavy in terms of computing time.

For the sake of clarity, we focus on a specific instance of covariate shift problem in which the target probability density  $q_X$  is known and  $\hat{p}$  is the kernel density estimate (KDE), i.e.,  $\hat{p}_X^{KDE}(x) = (1/n) \sum_{i=1}^n K_b(x - X_i)$ , where  $K_b$  typically is a Gaussian density with mean 0 and variance  $b^2$  (a hyper-parameter to be tuned). Note that such a situation does not involve any changes for our sampling procedure whereas it is clearly advantageous for the weighted approach for which one unknown,  $q_X$ , is now given. In this case, the analysis of  $\hat{Q}_w(h)$  can be carried out using the decomposition  $\hat{Q}_{\hat{w}}(h) - Q(h) = \hat{Q}_{\hat{w}}(h - Th) + \hat{Q}_{\hat{w}}(Th) - Q(h)$ , with  $Th(x) = \mathbb{E}[h(X, Y)|X=x]$ . The first term above is a sum of centered random variable which (provided some conditions) satisfies the so-called Lindeberg condition so that the central limit theorem implies that  $\sqrt{n}\hat{Q}_{\hat{w}}(h - Th)$  is asymptotically Gaussian. The second term above is more complicated and the analysis can be derived using results in Delyon and Portier (2016); Cléménçon and Portier (2018). Those results assert (under some conditions) that  $\hat{Q}_{\hat{w}}(Th) - Q(h) = O_p(nb^d + b)$  (in case  $p_X$  is Lipschitz). As a consequence, we obtain, optimizing over  $b$ , that  $\hat{Q}_{\hat{w}}(h) - Q(h) = O_p(n^{-1/2} + n^{-1/(1+d)})$ . This is easily compared to our bound, when  $q_X$  is known,  $n^{-1/2} + n^{-1/d}$ , which is smaller than the one given before.

**Direct weight estimation** Huang et al. (2006) proposed *Kernel Mean Matching (KMM)* for estimating the ratios of the probability density functions of the source and the target distribution. They used the estimated ratios for weighting the source sample. Gretton et al. (2008) further studied this method theoretically and empirically. Sugiyama et al. (2007, 2008) proposed a method that estimates the ratios as a function by minimizing the Kullback-Leibler divergence between the source density function multiplied by the ratio function and the target density function. The estimated function can predict ratios even outside of the source sample, which enables cross-validation for hyper-parameter tuning. Kanamori et al. (2009) proposed constrained and unconstrained least squares methods for estimating the ratio function called *Least-Squares Importance Fitting (LSIF)* and *unconstrained LSIF (uLSIF)*. Yamada et al. (2013) developed its variant called *Relative uLSIF (RuLSIF)*, which replaces the denominator of the ratio with a convex mixture of the source and the target density functions to circumvent issues caused by near-zero denominators. Zhang et al. (2021) proposed a covariate shift adaptation method that directly minimizes an upper bound of the target risk in order to avoid estimation of weights. The method shows great empirical performance while it does not exactly minimize the target risk and hence the minimizer converges to a biased solution.

**Connection to treatment effect estimation** One of the quantities of great interest in treatment effect estimation is the average treatment effect on the treated (ATT),  $\mathbb{E}[Y^{(1)} - Y^{(0)} | W = 1]$ , where  $W \in \{0, 1\}$  is a treatment assignment variable,  $Y^{(1)}$  and  $Y^{(0)}$  are potential outcomes corresponding to the treatment 1 and 0.<sup>2</sup> Suppose that we wish to estimate the ATT using i.i.d. observations of  $W$  and its outcome  $Y := Y^{(W)}$  together with covariates  $X$ ,  $\{(Y_i, W_i, X_i)\}_{i=1}^N$ . Under the standard assumptions (see e.g., Hernan and Robins (2023)) including the *conditional exchangeability*  $Y^{(w)} \perp\!\!\!\perp W | X$ , the *positivity*  $P(\{W = w | X\}) > 0$ , and the *consistency*  $W = w \implies Y^{(w)} = Y^{(W)} = Y$ , for each  $w \in \{0, 1\}$ , the ATT equals the difference between

$$\begin{aligned} \mathbb{E}[Y^{(1)} | W = 1] &= \int y P_{Y^{(1)} | X, W=1}(dy) P_{X | W=1}(dx) \\ &= \int y P_{Y | X, W=1}(dy) P_{X | W=1}(dx) \\ &= \mathbb{E}[Y | W = 1] \end{aligned} \quad (2)$$

and

$$\begin{aligned} \mathbb{E}[Y^{(0)} | W = 1] &= \int y P_{Y^{(0)} | X, W=0}(dy) P_{X | W=1}(dx) \\ &= \int y P_{Y | X, W=0}(dy) r(x) P_{X | W=0}(dx) \\ &= \mathbb{E}[r(X)Y | W = 0], \end{aligned} \quad (3)$$

where  $r(x)$  is the density ratio defined such that  $r(x)dP_{X | W=0}(x) = dP_{X | W=1}(x)$ . We can easily estimate the first term  $\mathbb{E}[Y^{(1)} | W = 1]$  (Eq. (2)) by the conditional sample average  $\frac{1}{N_1} \sum_{i=1}^N Y_i \cdot \mathbb{1}_{W_i=1}$ , where  $N_1 := \sum_{i=1}^N \mathbb{1}_{W_i=1}$ . Estimating the second term  $\mathbb{E}[Y^{(0)} | W = 1]$  (Eq. (3)) is more involved. The sample average with the condition  $W = 0$ ,  $\frac{1}{N_0} \sum_{i=1}^M Y_i \cdot \mathbb{1}_{W_i=0}$ , where  $N_0 := \sum_{i=1}^N \mathbb{1}_{W_i=0}$ , would be biased to  $\mathbb{E}[Y | W = 0] \neq \mathbb{E}[Y^{(0)} | W = 1]$ , but the bias is only due to the change in the conditional distributions of  $X$  given  $W = 1$  and  $X$  given  $W = 0$  quantified by  $r(X)$ , similarly to the covariate shift (see Eq. (3)). One way to correct the bias is to use an estimate  $\hat{r}$  of the ratio  $r$  for the weighted average  $\frac{1}{N_0} \sum_{i=1}^N \hat{r}(X_i) \cdot Y_i \cdot \mathbb{1}_{W_i=0}$ , similarly to the reweighting approach to covariate shift adaptation, leading to the following estimate:

$$\widehat{\text{ATT}} = \frac{1}{N_1} \sum_{i: W_i=1} Y_i - \frac{1}{N_0} \sum_{i: W_i=0} \hat{r}(X_i) \cdot Y_i.$$

Another popular approach is the nearest neighbor matching Abadie and Imbens (2006). See also Rosenbaum (1995) for a broad introduction to matching problems for evaluating treatment effects. In Abadie and Imbens (2006), the ATT is estimated by

$$\overline{\text{ATT}} = \frac{1}{N_1} \sum_{i: W_i=1} [Y_i - \hat{Y}_i(0)],$$

where  $\hat{Y}_i(0) = \frac{1}{k} \sum_{j=1}^N Y_j (1 - W_j) \mathbb{1}_{\{\|X_i - X_j\| \leq \tau_{n,k,X_i}\}}$  is the average of  $Y_j$ 's over the  $k$  first NNs of  $X_i$  in the untreated group. The estimator takes the form

$$\overline{\text{ATT}} = \frac{1}{N_1} \sum_{i: W_i=1} Y_i - \frac{1}{N_1} \sum_{i: W_i=0} \frac{K_k(i)}{k} Y_i,$$

where  $K_k(i) = \sum_{j=1}^N W_j \mathbb{1}_{\|X_i - X_j\| \leq \tau_{n,k,X_i}}$  is the number of times observation  $i$  is used as a match, i.e., the number of times observation  $i$  is among the  $k$  NNs of variables in the treated group. Note that  $\overline{\text{ATT}}$  coincides with  $\widehat{\text{ATT}}$  if  $\hat{r}(X_i)$  is defined as  $\frac{K_k(i)N_0}{kN_1}$ . Recently, Lin et al. (2023) showed that the latter quantity can be indeed interpreted as an estimate of the density ratio but its consistency requires  $k \rightarrow \infty$  while Abadie and Imbens (2006) considered a fixed value of  $k$ , as in our problem. To see an analogy with our method, one can consider the case in which  $h$  does not depend on  $x$ , i.e.  $h(y, x) = g(y)$  for some function  $g$ . Using the notation from the present paper ( $W = 0$  and  $W = 1$  indicate the target and the source domain, respectively, with  $N_0 = n$  and  $N_1 = m$ ) the second term of  $\overline{\text{ATT}}$  above generalizes to the form

$$\frac{1}{m} \sum_{i=1}^n \frac{K_k(i)}{k} g(Y_i) = \frac{1}{m} \sum_{i=1}^m \int g(y) \hat{P}_n(dy | X_i^*) \quad (4)$$

<sup>2</sup>A common scenario is that we have a treated group (represented by treatment 1) and a non-treated, or controlled group (represented by treatment 0).

with  $K_k(i) = \sum_{j=1}^m 1_{\{\|X_i - X_j^*\| \leq \tau_{n,k,X_i^*}\}}$ . The previous estimate corresponds to the one introduced in the notes following Proposition 2. On the other hand, our estimator applied to this case is given by

$$\frac{1}{m} \sum_{i=1}^m g(Y_{n,i}^*) \quad (5)$$

Both estimators are different when  $k > 1$  but they coincide as soon as  $k = 1$ . In fact,  $\hat{P}_n(dy|X_i^*)$  has one single atom when  $k = 1$ , so that sampling from it and evaluating the average are the same. Here are a few remarks.

- When  $k > 1$ , Eq. (5) requires fewer evaluations of  $g$  than Eq. (4). This is relevant when evaluation of the function is time-consuming or costly such as observation from physical experiments.
- Our theoretical analysis is rather different from that of Abadie and Imbens (2006). Since they rely on the expression in the left side of Eq. (4), it is unclear whether they can or not handle the case when  $h$  depends on  $x$  (required for prediction purpose). In contrast, our approach is based on the decomposition given in Section 4.2, with sampling error and estimation error, leveraging  $\int \int g(y) \hat{P}_n(dy|x) Q(dx)$  as a centering term. Our results are more general because they include the case when  $h$  depends on  $x$  and also we can deal with both estimates (4) and (5) in the meantime, as mentioned in the notes following Proposition 2). Moreover, Proposition 3 implies a lower bound for Eq. (4) and we believe this result to be new in treatment effect literature.

**Other references** The idea of nonparametric sampling is a standard one in the field of texture synthesis. In particular, the choice of 1-NN resampling was often used as a fast method to generate new textures from a small sample. See Truquet (2011) for a literature review in this context. Our conditional sampling framework bears resemblance with traditional bootstrap sampling as there is random generation according to some estimated distribution. In contrast, the original bootstrap method is usually made up using draws from the standard empirical measure  $(1/n) \sum_{i=1}^n \delta_{X_i, Y_i}$ . Here another distribution,  $\hat{P}$ , has been used to generate new samples. Moreover, our goal is totally different here. While the bootstrap technique was initially introduced for making inference, here the goal is to estimate an unknown quantity  $Q(h)$  which appears in many machine learning tasks. Kpotufe and Martinet (2021) theoretically study covariate shift adaptation under the assumption that we have access to a labeled sample both from the source and the target distribution. Although they consider a  $k$ -nearest-neighbor-based method, it is essentially different from ours since they perform the  $k$ -NN method on the union of the source and the target sample. Lee (2013) proposed pseudo-labeling unlabeled data in the context of semi-supervised learning. Wang (2023) proposed a hyper-parameter selection method for kernel ridge regression under covariate shift using pseudo-labeling. The author focuses on model selection in regression problems while we study the mean estimation that can be applied to a wider range of supervised learning problems.

## 7 Extensions

Several ways to extend our method beyond the mean estimation problem are considered in this section.

**Heterogeneity in target distributions** The case where the target covariates distribution  $Q_X$  changes across the data might be of interest if one wishes to aggregate several pieces of target data whose covariates distributions are not necessarily the same. This might occur when the target data is obtained by gathering individuals from different countries, and consequently, the distributions are not the same anymore or when the time between the measurements has caused some changes in the distribution.

While such an heterogeneity in target data might be seen as more complicated at first glance, it actually can be examined using a similar decomposition and the same tools as the one used to obtain the non-asymptotic bound in Theorem 1. More formally, the target distribution is here  $Q = (1/m) \sum_{i=1}^m Q_i$  with  $Q_i = P_{Y|X} Q_{X,i}$ . For each  $m \geq 1$  and  $n \geq 1$ , let  $(X_i^*, Y_{n,i}^*)_{1 \leq i \leq m}$  be a collection of random variables conditionally independent given  $(X_i, Y_i)_{i=1}^n$  and such that for each  $i = 1, \dots, m$ ,  $(X_i^*, Y_{n,i}^*) \sim \hat{Q}_{i,n}$  with  $\hat{Q}_{i,n} = \hat{P}_{Y|X} Q_{X,i}$ . The quantity of interest and the proposed estimator are therefore slightly different from before, given by, respectively,

$$Q(h) = m^{-1} \sum_{i=1}^m Q_i(h), \quad \hat{Q}^*(h) = m^{-1} \sum_{i=1}^m h(X_i^*, Y_{n,i}^*).$$

The decomposition is

$$\hat{Q}^*(h) - Q(h) = m^{-1} \sum_{i=1}^m \left\{ h(X_i^*, Y_{n,i}^*) - \hat{Q}_{i,n}(h) \right\} + m^{-1} \sum_{i=1}^m \left\{ \hat{Q}_{i,n}(h) - Q_i(h) \right\}.$$

The non-asymptotic analysis of the sampling error is similar to before as the Bernstein inequality is tailored to non-identically distributed variables. We obtain that the rate  $O(m^{-1/2})$  as before by simply requiring a bound on the variance of each random variables. The other term concerning the conditional distribution can be analyzed by writing

$$\left\| m^{-1} \sum_{i=1}^m \left\{ \hat{Q}_{i,n}(h) - Q_i(h) \right\} \right\|_{L_2} \leq m^{-1} \sum_{i=1}^m \|\hat{Q}_{i,n}(h) - Q_i(h)\|_{L_2}$$

and therefore we can directly apply Proposition 3 (given the assumptions are satisfied for each  $i$  uniformly). We finally obtain the rate  $m^{-1/2} + n^{-1/d} + n^{-1/2}$ , similar to the one obtained before.

**Stochastic gradient descent** Our sampling approach can be easily combined with the well-known stochastic gradient descent algorithm (and more generally with stochastic approximation) where only a small part of the data is used at each step to update the estimator. This particular property allows to require a small number of operations at each iteration (in contrast with gradient based optimization).

To illustrate this idea, consider the empirical risk minimization problem described in Section 5 where one is interested in solving  $\min_{\theta} \{R^*(\theta) := \mathbb{E}[m_{\theta}(Y^*, X^*)]\}$  where  $\theta \mapsto m_{\theta}(y, x)$  is differentiable. Suppose that  $n$  source samples have been obtained making the conditional distribution  $\hat{P}_{Y|X}$  available for sampling new points. Then the algorithm at step  $i \geq 1$ , might proceed by first generating  $X_i^*$  and then  $Y_{n,i}^* \sim \hat{P}_{Y|X_i^*}$ . This means finding the nearest neighbor to  $X_i^*$  among the source data and represents only  $\log(n)$  operations using the  $kd$ -tree. Having this been done, the update is simply

$$\theta_i = \theta_{i-1} - \gamma_i \nabla_{\theta} m_{\theta}(Y_{n,i}^*, X_i^*).$$

It results that each iteration in the above is similar to standard stochastic gradient descent, the only difference being the additional 1-nearest neighbor search. We stress that this is contrasting with the re-weighting approach for which a new sample, say  $X_i$ , would require evaluating  $\hat{p}_X(X_i)$  and therefore would need to compute all  $n$  distances between  $X_j$ ,  $j = 1, \dots, n$  and the new  $X_i$ .

**Semiparametric estimation** Simulating the labels to obtain a new sample is also convenient in semiparametric problems where quantities of interest often involve additional estimated parameters. Typical semiparametric problems involve expectations of functions that are indexed by an unknown parameter,  $\mathbb{E}[h_{\theta}(X, Y)]$ , and  $\theta$  is estimated from the data using some transformation  $\hat{\theta}$  of the sample. In such a situation, while estimating  $\theta$  using reweighting is unclear without more information on  $\hat{\theta}$ , one can directly use our sampling approach by introducing  $m^{-1} \sum_{j=1}^m h_{\hat{\theta}^*}(Y_j^*, X_j^*)$  where  $\hat{\theta}^* = \hat{\theta}(X_1^*, Y_1^*, \dots, X_m^*, Y_m^*)$ . This allows to obtain a semiparametric estimate with covariate shift adaptation. See Van der Vaart (2000), Chapters 19.4 and 25.8 for more details and examples in parametric or semiparametric estimation.

## 8 Experiments

The main purpose of the experiments is to compare our  $k$ -NN-CSA approach with several state-of-the-art competitors when facing multiple situations from mean estimation to empirical risk minimization with synthetic and real-world data.

We consider the following instances of our proposed method.

**1-NN-CSA:** the *Conditional Sampling Covariate-shift Adaptation (CSA)* (Algorithm 1) with  $k$ -Nearest Neighbor ( $k$ -NN) conditional sampler (Algorithm 2) with  $k = 1$ .

**log  $n$ -NN-CSA:** the same as above but with  $k = \log n$ .

We use the Python module `cKDTree` (Archibald, 2008) from SciPy (Virtanen et al., 2020) for nearest neighbor search in our methods. We compare them with the following existing covariate-shift adaptation methods.

**KDE-R-W (KDE-Ratio-Weighting):** the weighting method using the ratio of the Kernel Density Estimates (KDEs) of  $p_X$  and  $q_X$  (see Section 6).

**KMM-W (KMM-Weighting):** the weighting method estimating  $q_X/p_X$  using the *Kernel Mean Matching (KMM)* (Huang et al., 2006; Gretton et al., 2008). We use the Gaussian kernel.

**KLIEP-W (KLIEP-Weighting):** the weighting method estimating  $q_X/p_X$  by the *Kullback-Leibler Importance Estimation Procedure (KLIEP)* (Sugiyama et al., 2007, 2008). The linear combination of the Gaussian basis functions centered at the sample points are used for modeling the weight function.

**KLIEP100-W** : the same as KLIEP-W but with 100 randomly subsampled basis functions (Sugiyama et al., 2007, 2008) for reducing the time- and space-complexities.

**RuLSIF-W (RuLSIF-Weighting)**: the weighting method using  $q_X / (\alpha p_X + (1 - \alpha)q_X)$  estimated by *Relative unconstrained Least-Squares Importance Fitting (RuLSIF)* (Yamada et al., 2013), where  $\alpha \in [0, 1]$  is a hyperparameter. We use the default value  $\alpha = 0.1$ . As a model of the weight function, the Gaussian basis functions centered at the sample points are used.

**RuLSIF100-W**: the same as RuLSIF-W but with 100 randomly subsampled basis functions (Yamada et al., 2013) for reducing the time- and space-complexities.

See Section 6 for more explanations of those methods. For KMM-W, KLIEP-W, and RuLSIF-W, we used the implementations from *Awesome Domain Adaptation Python Toolbox (ADAPT)* (de Mathelin et al., 2021). All the computations were performed on the cluster, *Grid5000* (Balouek et al., 2013). For the methods using Gaussian basis functions (KLIEP-W, KLIEP100-W, RuLSIF-W, RuLSIF100-W), we use 5-fold cross-validation for choosing the Gaussian bandwidth from  $\{0.001, 0.01, 0.1, 1, 10\}$ . KMM-W does not offer a way to do cross-validation, and we fixed to 1. More details are in the supplementary material.

Furthermore, we also report the results for the following baseline method and ideal method.

**NoCorrection**: the method that takes the average  $\frac{1}{n} \sum_{i=1}^n h(X_i, Y_i)$  only using the source sample  $(X_i, Y_i)_{i=1}^n \sim P^m$ , ignoring the target sample.

**OracleY**: the result for taking the average  $\frac{1}{m} \sum_{i=1}^m h(X_i^*, Y_i^\circ)$  using a sample  $(X_i^*, Y_i^\circ)_{i=1}^m \sim Q^m$ . Note that  $Y_i^\circ$  are not available in practical scenarios of our interest and made invisible to other methods.

We conduct experiments in three setups, detailed below, with different sample sizes  $n (= m)$  and data dimensionalities  $d$ :  $(n, d) \in \{50, 100, 500, 1000, 5000, 10000\} \times \{5, 10\}$ . Each experiment is repeated 50 times with different random seeds.

**Setup of Experiment E1 (mean estimation with synthetic data)**: The task here is to estimate  $Q(h) = \iint y P_{Y|X}(dy | x) Q_X(dx)$  under the following setup. We define  $h$  by  $h(x, y) = y$ ,  $P_X$  as the uniform distribution over  $[-1, 1]^d$ ,  $Q_X$  as that over  $[0, 1] \times [-1, 1]^{d-1}$ , and  $P_{Y|X=x}$  as the normal distribution with mean  $x$  and variance 0.1. Figure 7a in Appendix G shows an illustration of the setup. In this setup, we have  $P(h) = \iint y P_{Y|X}(dy | x) P_X(dx) = 0$  while  $Q(h) = \iint y P_{Y|X}(dy | x) Q_X(dx) = 0.5$ . Because of this difference, covariate shift adaptation is essential for correctly estimating  $Q(h)$ .

**Comparison of estimation errors for Experiment E1**: The results are presented in Figure 1. First, the errors for NoCorrection are not decreasing as the sample sizes increase, ending up with large errors in all cases, because of the bias due to the covariate shift. Other methods with covariate-shift adaptation had always smaller errors than that of this baseline. Excluding OracleY, an ideal method unavailable in practice, KLIEP100-W, KMM-W, 1-NN-CSA, and  $\log n$ -NN-CSA were among the best for smaller dimensionalities  $d \in \{1, 2\}$  (Figures 1a and 1b). For the larger dimensionalities  $d \in \{5, 10\}$ , KMM-W and 1-NN-CSA outperformed other methods. In particular, 1-NN-CSA gave outstanding performances in many cases except  $d = 5$  and  $n \in \{100, 500, 1000\}$ , for which KMM-W was even better. The errors of most methods roughly follow power laws, where the slope of a line corresponds to the power of the convergence rate (steeper is better). 1-NN-CSA and  $\log n$ -NN-CSA seem to have the steepest slopes for  $d = 10$ , although comparison is difficult for the lower dimensionalities.

**Comparison of running times in Experiment E1**: Figure 2 shows the comparison in running times. 1- and  $\log n$ -NN-CSA were much faster than other methods in all cases except for  $(d, n) = (10, 50)$ . Their advantage is most pronounced for larger sample sizes. For instance, 1- and  $\log n$ -NN-CSA were at least 100 times faster than other methods for  $(n, d) = (10000, 1)$  (Figure 2a).

**Setup of Experiment E2 (risk estimation with synthetic data)**: In this experiment, we compare the methods in the context of risk estimation of a fixed function  $f_0$ . Setting  $f_0: x \mapsto -x_1$ , where  $x_1$  is the first coordinate of  $x$ , we estimate the expected loss (i.e., risk) of  $f_0(X)$  with the square loss in predicting the response  $Y$  when  $(X, Y)$  follows  $Q$ . In other words, we set  $h$  as  $h(x, y) := (y - f_0(x))^2$ , and the goal is to estimate the risk  $Q(h) = \int (y - f_0(x))^2 Q(dx, dy)$ . We use the uniform distribution over  $[-1, 1]^d$  for  $P_X$  and that over  $[0, 1] \times [0, 1]^{d-1}$  for  $Q_X$ . The conditional distribution  $P_{Y|X=x}$  is the normal distribution with mean  $|x_1|$  and variance 0.1 for any  $x \in \mathcal{X} := [-1, 1]^d$ . Under this setup, the function  $f_0$  performs poorly on the support of  $Q_X$  and should incur a large risk. See Figure 7b for an illustration of the setup. In this setup, the risks under  $P_X$  and  $Q_X$  largely differ because  $f_0$  fits  $(X, Y)$  well in a half of the support of  $P_X$  but not in that of  $Q_X$ .

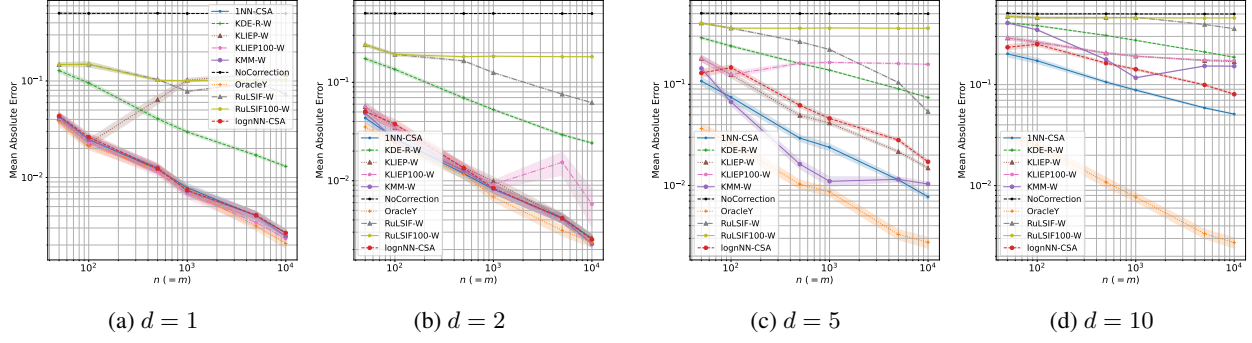


Figure 1: Mean Squared Errors (MSE) for Experiment E1 (estimation of  $\int y Q(dy)$ ). The horizontal axis is for the sample sizes  $n (= m)$ , and the vertical axis is for the mean absolute error of each estimate. The four figures are for different data dimensionalities.




Figure 2: Running times for Experiment E1. The horizontal axis is for the sample sizes  $n (= m)$ , and the vertical axis is for the mean running time of each method. The four figures are for different data dimensionalities.

**Comparison of estimation errors for Experiment E2:** We present the estimation errors for Experiment E2 in Figure 3. KMM-W, 1-NN-CSA,  $\log n$ -NN-CSA gave similar results, almost matching those of OracleY, while KMM-W and 1-NN-CSA were advantageous for  $d = 5$ , and 1-NN-CSA outperformed other methods for  $d = 10$ . We can notice that KDE-W, KMM-W, RuLSIF-W, and RuLSIF100-W did not always improve errors over NoCorrection (Figures 1c and 1d). Some methods such as KMM-W and KLIEP-W showed great performance in some cases while giving poor results in other cases. In contrast, 1-NN-CSA showed stable and often best performances in these experiments.

**Comparison of running times in Experiment E2:** The running times in Experiments E2 (Figure 4) were very similar to those in Experiment E1 (Figure 4), but we can see more clearly that 1- and  $\log n$ -NN-CSA outperform other methods even for the smallest sample size.




Figure 3: Estimation errors for Experiment E2 (estimation of  $\int (y - f_0(x))^2 Q(dx, dy)$ )




Figure 4: Running times in Experiment E2

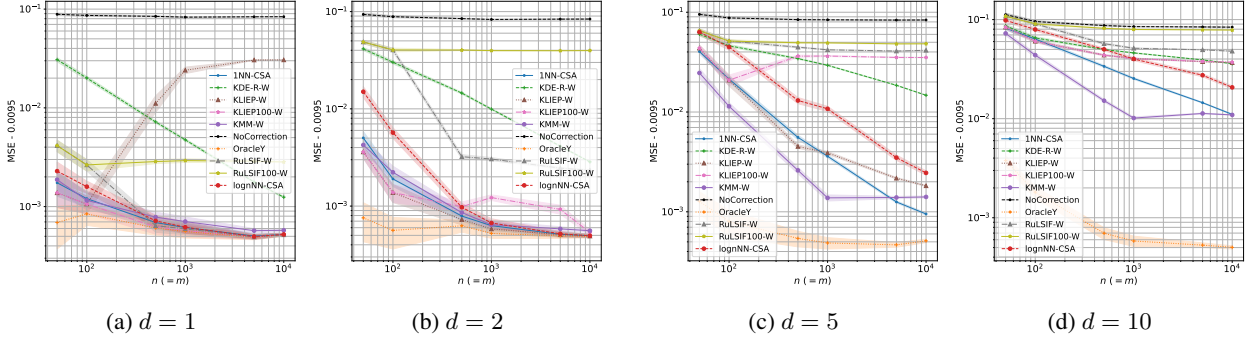


Figure 5: Mean Squared Errors (MSE) (subtracted by 0.0095) for Experiment E3 (linear regression)

**Setup of Experiment E3 (linear regression with synthetic data):** Next, we present experiments of linear regression. Using samples from the same source and test distributions as in Experiment E2, we perform the ordinary least squares after covariate adaptation. More precisely, we aim to optimize the parameters  $\theta \in \mathbb{R}^d$  of the model  $f_\theta: \mathbb{R}^d \rightarrow \mathbb{R}, x \mapsto \theta^\top x$  so that the Mean Squared Error (MSE)  $\mathbb{E}[(Y^* - f_\theta(X^*))^2]$  in the target domain will be minimized. To do so, we minimize the MSE estimated by each covariate shift adaptation method.

**Comparison of estimation errors for Experiment E3:** The results are summarized in Figures 5.<sup>3</sup> KMM-W performed better than any other methods for the higher dimensions  $d \in \{5, 10\}$  and the small-to-moderate sample sizes  $50 \leq n \leq 500$ , 1-NN-CSA being the second best. For  $n = 10000$ , 1-NN-CSA showed performance better than or comparable to KMM-W.

**Comparison of running times in Experiment E3:** As in Experiments E1–E2, 1-NN-CSA and  $\log n$ -NN-CSA finished their computations faster than the other adaptation methods by large margins (Figure 6).

In Experiments E1–E3, the proposed methods, 1- and  $\log n$ -NN-CSA were able to finish computation much faster than other adaptation methods without compromising on the statistical performance.  $\log n$ -NN-CSA did not show advantages in accuracy, with increased computation costs. We can conclude that 1-NN-CSA is preferred over  $\log n$ -NN-CSA. A reason that we were not able to conduct experiments with larger sample sizes than 10000 is that the existing adaptation methods have too demanding computational requirements. For instance, the running times of RuLSIF-W in Figure 6c grows about 100 times as the sample size increases by 10 times, taking more than  $10^3$  seconds for  $n = 10000$ . For  $n = 10^5$ , we would need at least  $10^3 \times 100$  seconds, that is 27 hours of compute for a single run. In contrast, the time complexity of 1-NN-CSA being  $\mathcal{O}(n \log n)$  and its running time less than one second for  $n = 10^4$ , we can estimate its running time for  $n = 10^5$  as  $1 \times (10^5 \log 10^5) / (10^4 \log 10^4) = 12.5$  seconds. 1-NN-CSA would stay feasible in applications of even larger scales.

The previous methods construct the distance matrix between pairs of data points, which takes running time and memory space quadratic in the sample size. Additionally, RuLSIF-W computes the inverse of the distance matrix, taking cubic running time. KMM-W and KLIEP-W solve convex optimization problems with iterative procedures, for which the

<sup>3</sup>We plot the MSEs subtracted by 0.0095 to better present the curves in the region close to the minimum population MSE 0.01 while keeping values positive.



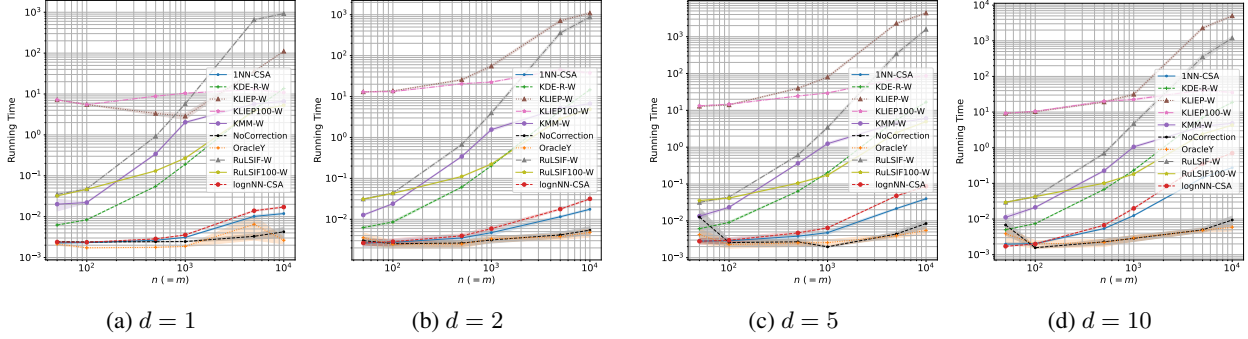


Figure 6: Computation times (seconds) in Experiment E3

Table 1: MSE/accuracy for regression/classification benchmark datasets. We repeat the experiment using 50 different random subsamples and calculate the average scores (and standard errors). The results comparable to the best in terms of Wilcoxon’s signed rank test with significance level 1% are shown in bold fonts.

	Regression (MSE)		Classification (accuracy)	
	diabetes	california	breast_cancer	twonorm
1NN-CSA	3470 (35)	<b>0.146</b> (0.001)	<b>0.9633</b> (0.002)	0.9327 (0.002)
lognNN-CSA	3605 (40)	0.150 (0.001)	0.9595 (0.002)	0.9293 (0.002)
KDE-R-W	3673 (52)	3.864 (1.067)	0.9596 (0.002)	0.5260 (0.009)
KMM-W	3831 (60)	3.702 (1.160)	0.9594 (0.002)	<b>0.9583</b> (0.001)
KLIEP-W	<b>3221</b> (31)	2.896 (0.798)	<b>0.9648</b> (0.002)	0.9482 (0.001)
KLIEP100-W	<b>3223</b> (31)	3.034 (0.843)	<b>0.9648</b> (0.002)	0.9480 (0.001)
RuLSIF-W	3235 (31)	3.039 (0.843)	0.7794 (0.015)	0.9512 (0.001)
RuLSIF100-W	3238 (31)	3.045 (0.844)	0.7794 (0.015)	0.9539 (0.001)

implementations from de Mathelin et al. (2021) use stopping criteria based on objective function values. This resulted in good accuracy and milder growth in running time in our experiments. However, tuning the solvers can be involved in practice. In contrast,  $k$ -NN-CSA does not have such subtle issues around optimization solvers: we only have to perform nearest neighbor search.

In all cases, we can observe that 1-NN-CSA showed clear power-law, with nearly straight lines in the logarithmic scales. This is a significant advantage in predicting returns when one invests on increasing the sample size.

**Experiment E4 (linear regression and logistic regression with benchmark datasets):** We use regression benchmark datasets, *diabetes*<sup>4</sup>, *california* (Pace and Barry, 1997)<sup>5</sup> and classification datasets, *twonorm* (Breiman, 1996)<sup>6</sup> and *breast\_cancer*<sup>4</sup>. We apply the ridge regression and the logistic regression, respectively. The evaluation metric is the mean squared error for the regression tasks and the classification accuracy for classification tasks. We synthetically introduce covariate shift by subsampling test data. See Appendix H for more details.

**Remark:** For fair comparison, the benchmark experiments presented in this paper follow the standard protocol used in the literature as similarly done in previous research (Gretton et al., 2008; Kanamori et al., 2009; Yamada et al., 2013; Sugiyama et al., 2007, 2008): we apply biased resampling to synthetically simulate a target dataset under covariate shift. It is thus important to note that they are not completely real-world data. Nevertheless, this ensures that the methods are tested in isolation from other types of distribution shifts while using real data for the source covariate distribution as well as the conditional distributions.

**Results for Experiment E4:** Table 1 shows the obtained MSEs and classification accuracies. 1-NN-CSA and  $\log n$ -NN-CSA gave the best performance for *california* and performances comparable to the best for *breast\_cancer*.

<sup>4</sup>Available at <https://archive.ics.uci.edu/ml/index.php>.

<sup>5</sup>Available at [https://www.dcc.fc.up.pt/~ltorgo/Regression/cal\\_housing.html](https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html).

<sup>6</sup>Available at <https://www.cs.utoronto.ca/~delve/data/datasets.html>.

Table 2: Total running times in seconds spent for the training including the hyper-parameter tuning (if any) for benchmark datasets. We repeat the experiment using 50 different random subsamples dataset and calculate the average running times (and standard errors). The results comparable to the best in terms of Wilcoxon’s signed rank test with significance level 1% are shown in bold.

	diabetes	california	breast_cancer	twonorm
1NN-CSA	<b>0.0015</b> (0.0000)	<b>0.0084</b> (0.0001)	<b>0.0036</b> (0.0000)	<b>0.0051</b> (0.0000)
lognNN-CSA	0.0016 (0.0000)	0.0128 (0.0001)	0.0037 (0.0000)	0.0052 (0.0000)
KDE-R-W	0.0078 (0.0000)	0.2121 (0.0008)	0.0117 (0.0000)	0.0124 (0.0000)
KMM-W	0.0373 (0.0015)	0.4067 (0.0038)	0.0542 (0.0014)	0.0220 (0.0006)
KLIEP-W	7.602 (0.051)	29.98 (0.34)	8.67 (0.07)	8.86 (0.16)
KLIEP100-W	7.501 (0.045)	16.91 (0.07)	8.68 (0.07)	8.26 (0.10)
RuLSIF-W	0.0575 (0.0014)	1.686 (0.011)	0.0529 (0.0020)	0.2014 (0.0016)
RuLSIF100-W	0.0401 (0.0007)	0.1237 (0.0004)	0.0454 (0.0014)	0.0391 (0.0002)

For the other datasets, different methods performed the best depending on the dataset. On the other hand, in terms of running time, 1NN-CSA was consistently faster than the previous methods (Table 2).

Our experiments show that the proposed method is almost always faster than the previous methods and gives great accuracy in many cases, even though it is not always the best. 1-NN-CSA is highly recommended as an off-the-shelf method applicable even in larger scales, although the previous methods such as KMM-W, KLIEP-W, and RuLSIF-W should not be neglected, as far as the computational budget allows. The times spent for adaptation are summarized in Table 2, showing that the proposed methods 1-NN-CSA and  $\log n$ -NN-CSA are much faster than other methods.

## 9 Conclusion

We proposed a  $k$ -NN-based covariate shift adaptation method. We provided error bounds, which suggest setting  $k = 1$  is among the best choices. This resulted in a scalable non-parametric method with no hyper-parameter. For future research directions, one could complete our results for the parametric inference on the target domain, in particular for finding the asymptotic distribution of  $M$ -estimators. For the average treatment effect, Abadie and Imbens (2006) derived asymptotic normality of their estimator and it could be interesting to get a similar result in our context. Investigating non-parametric estimation on the target domain could be also an interesting direction. However, non-parametric estimators computed with the source sample can be already optimal when the ratio of densities is bounded. See for instance Ma et al. (2023) in the reproducing kernel Hilbert space framework. It could be then interesting to extend our result to cases with unbounded density ratios. Finally, it may be interesting to extend our approach with approximate nearest neighbor methods for further scalability.

## Acknowledgement

Experiments presented in this paper were carried out using the Grid’5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>). IY was supported by the *Allocation d’Installation Scientifique (AIS) 2023* from Rennes Métropole.

## References

- Abadie, A. and G. W. Imbens (2006). Large sample properties of matching estimators for average treatment effects. *econometrica* 74(1), 235–267.
- Aminian, G., M. Abroshan, M. Mahdi Khalili, L. Toni, and M. Rodrigues (2022, 28–30 Mar). An information-theoretical approach to semi-supervised learning under covariate-shift. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, Volume 151 of *Proceedings of Machine Learning Research*, pp. 7433–7449. PMLR.
- Archibald, A. M. (2008). cKDTree.
- Balouek, D., A. Carpen Amarie, G. Charrier, F. Desprez, E. Jeannot, E. Jeanvoine, A. Lèbre, D. Margery, N. Niclausse, L. Nussbaum, O. Richard, C. Pérez, F. Quesnel, C. Rohr, and L. Sarzyniec (2013). Adding virtualization capabilities to the Grid’5000 testbed. In I. I. Ivanov, M. van Sinderen, F. Leymann, and T. Shan (Eds.), *Cloud Computing and Services Science*, Volume 367 of *Communications in Computer and Information Science*, pp. 3–20. Springer International Publishing.
- Bentley, J. L. (1975). Multidimensional binary search trees used for associative searching. *Communications of the ACM* 18(9), 509–517.
- Berrett, T. B., R. J. Samworth, and M. Yuan (2019). Efficient multivariate entropy estimation via k-nearest neighbour distances. *The Annals of Statistics* 47(1), 288–318.
- Biau, G. and L. Devroye (2015). *Lectures on the nearest neighbor method*, Volume 246. Springer.
- Blanchet, J., H. Chen, Y. Lu, and L. Ying (2024). When can regression-adjusted control variate help? rare events, sobolev embedding and minimax optimality. *Advances in Neural Information Processing Systems* 36.
- Bogachev, V. I. and M. A. S. Ruas (2007). *Measure theory*, Volume 2. Springer Science & Business Media.
- Breiman, L. (1996). Bias, variance, and arcing classifiers.
- Cannings, T. I., T. B. Berrett, and R. J. Samworth (2020). Local nearest neighbour classification with applications to semi-supervised learning. *The Annals of Statistics* 48(3), 1789–1814.
- Chen, L., M. Zaharia, and J. Y. Zou (2022). Estimating and explaining model performance when both covariates and labels shift. In *Advances in Neural Information Processing Systems*, Volume 35, pp. 11467–11479. Curran Associates, Inc.
- Cléménçon, S. and F. Portier (2018). Beating monte carlo integration: A nonasymptotic study of kernel smoothing methods. In *International Conference on Artificial Intelligence and Statistics*, pp. 548–556. PMLR.
- Dasgupta, S. and S. Kpotufe (2014). Optimal rates for k-nn density and mode estimation. In *Advances in Neural Information Processing Systems*, Volume 27.
- de Mathelin, A., M. Atiq, G. Richard, A. de la Concha, M. Yachouti, F. Deheeger, M. Mougeot, and N. Vayatis (2021). ADAPT : Awesome Domain Adaptation Python Toolbox. arXiv:2107.03049 [cs.LG].
- Delyon, B. and F. Portier (2016). Integral approximation by kernel smoothing. *Bernoulli* 22(4), 2177–2208.
- Devroye, L., L. Györfi, A. Krzyżak, and G. Lugosi (1994). On the strong universal consistency of nearest neighbor regression function estimates. *Ann. Statist.* 22(3), 1371–1385.
- Devroye, L., L. Györfi, and G. Lugosi (2013). *A probabilistic theory of pattern recognition*, Volume 31. Springer Science & Business Media.
- Dua, D. and C. Graff (2017). UCI machine learning repository.
- Friedman, J. H., J. L. Bentley, and R. A. Finkel (1977). An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software (TOMS)* 3(3), 209–226.
- Gadat, S., T. Klein, and C. Marteau (2016). Classification in general finite dimensional spaces with the  $k$ -nearest neighbor rule. *Ann. Statist.* 44(3), 982–1009.
- Gretton, A., A. Smola, J. Huang, M. Schmittfull, K. Borgwardt, and B. Schölkopf (2008, December). Covariate Shift by Kernel Mean Matching. In *Dataset Shift in Machine Learning*, pp. 131–160. The MIT Press.
- Györfi, L., M. Kohler, A. Krzyżak, and H. Walk (2006). *A Distribution-Free Theory of Nonparametric Regression*. Springer Science & Business Media.
- Hernan, M. A. and J. M. Robins (2023). *Causal Inference: What If*. Chapman & Hall/CRC Monographs on Statistics & Applied Probab. CRC Press.

- Huang, J., A. Gretton, K. Borgwardt, B. Schölkopf, and A. Smola (2006). Correcting sample selection bias by unlabeled data. In *Advances in Neural Information Processing Systems*, Volume 19. MIT Press.
- Jiang, H. (2019). Non-asymptotic uniform rates of consistency for  $k$ -NN regression. In *AAAI Proceedings*, Volume 33, pp. 3999–4006.
- Kanamori, T., S. Hido, and M. Sugiyama (2009). A least-squares approach to direct importance estimation. *Journal of Machine Learning Research* 10(48), 1391–1445.
- Kpotufe, S. and G. Martinet (2021). Marginal singularity and the benefits of labels in covariate-shift. *The Annals of Statistics* 49(6), 3299–3323.
- Le, Q., T. Sarló, A. Smola, et al. (2013). Fastfood—approximating kernel expansions in loglinear time. In *Proceedings of the 30th International Conference on Machine Learning*, Volume 28.
- Lee, D.-H. (2013). Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on Challenges in Representation Learning, ICML*.
- Leluc, R., F. Portier, J. Segers, and A. Zhuman (2023). Speeding up monte carlo integration: Control neighbors for optimal convergence. *To appear in Bernoulli*. [ArXiv:2305.06151](https://arxiv.org/abs/2305.06151).
- Lin, Z., P. Ding, and F. Han (2023). Estimation based on nearest neighbor matching: from density ratio to average treatment effect. *Econometrica* 91(6), 2187–2217.
- Ma, C., R. Pathak, and M. J. Wainwright (2023). Optimally tackling covariate shift in rkhs-based nonparametric regression. *The Annals of Statistics* 51(2), 738–761.
- Pace, R. K. and R. Barry (1997). Sparse spatial autoregressions. *Statistics & Probability Letters* 33(3), 291–297.
- Portier, F. (2021). Nearest neighbor process: weak convergence and non-asymptotic bound. *To appear in Bernoulli*. [ArXiv:2110.15083](https://arxiv.org/abs/2110.15083).
- Rosenbaum, P. R. (1995). *Observational Studies*. Springer.
- Shimodaira, H. (2000). Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference* 90(2), 227–244.
- Sugiyama, M., S. Nakajima, H. Kashima, P. v. Büna, and M. Kawanabe (2007). Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems 20, NIPS 2007*, pp. 1433–1440.
- Sugiyama, M., T. Suzuki, S. Nakajima, H. Kashima, P. von Büna, and M. Kawanabe (2008, December). Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics* 60(4), 699–746.
- Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics* 12, 389–434.
- Truquet, L. (2011). On a nonparametric resampling scheme for markov random fields. *Electronic Journal of Statistics* 5, 1503–1536.
- Van der Vaart, A. W. (2000). *Asymptotic Statistics*, Volume 3. Cambridge University Press.
- Virtanen, P., R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors (2020). SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17, 261–272.
- Wang, K. (2023, March). Pseudo-Labeling for Kernel Ridge Regression under Covariate Shift. [arXiv:2302.10160](https://arxiv.org/abs/2302.10160) [cs, math, stat].
- Wendel, J. G. (1948). Note on the gamma function. *The American Mathematical Monthly* 55(9), 563.
- Williams, C. K. I. and M. Seeger (2000). Using the Nyström Method to Speed up Kernel Machines. In *Advances in Neural Information Processing Systems*, Volume 13 of *NIPS 2000*, pp. 661–667. MIT Press.
- Yamada, M., T. Suzuki, T. Kanamori, H. Hachiya, and M. Sugiyama (2013, May). Relative Density-Ratio Estimation for Robust Distribution Comparison. *Neural Computation* 25(5), 1324–1370.
- Zhang, T., I. Yamane, N. Lu, and M. Sugiyama (2021, June). A One-Step Approach to Covariate Shift Adaptation. *SN Computer Science* 2(4), 319.

## A Preliminary results

The first preliminary result is concerned about the order of magnitude of  $P_X(B(x, \tau))$  for which we obtain a lower bound and an upper bound.

**Lemma 1.** *Under (X1), (X2), and (X3), it holds, for every  $x \in S_X$  and  $\tau \in [0, T]$ ,*

$$M_{1,d}\tau^d \leq P_X(B(x, \tau)) \leq M_{2,d}\tau^d,$$

with  $M_{1,d} = cb_X V_d$  and  $M_{2,d} = U_X V_d$ .

*Proof.* The proof of the lower bound follows from

$$P_X(B(x, \tau)) = \int_{B(x, \tau) \cap S_X} p_X(y) dy \geq b_X \int_{B(x, \tau) \cap S_X} dy \geq b_X c \int_{B(x, \tau)} dy,$$

where we have used (X3) to get the first inequality and then (X2) to obtain the second one. We conclude by change of variable  $y = x + \tau u$ . The proof of the upper bound is similar:

$$P_X(B(x, \tau)) = \int_{B(x, \tau) \cap S_X} p_X(y) dy \leq U_X \int_{B(x, \tau) \cap S_X} dy \leq U_X \int_{B(x, \tau)} dy.$$

□

The same type of result can be obtained for  $P_X(B(x_1, \tau_1) \cup B(x_2, \tau_2))$  as follows.

**Lemma 2.** *Under (X1), (X2), and (X3), it holds, for every  $(x_1, x_2) \in S_X \times S_X$  and  $(\tau_1, \tau_2) \in [0, T]^2$ ,*

$$\frac{1}{2} M_{1,d}(\tau_1^d + \tau_2^d) \leq P_X(B(x_1, \tau_1) \cup B(x_2, \tau_2)) \leq M_{2,d}(\tau_1^d + \tau_2^d),$$

with  $M_{1,d} = cb_X V_d$  and  $M_{2,d} = U_X V_d$ .

*Proof.* The proof of the upper bound follows from the union bound and Lemma 1. For the lower bound, start noting that for any events  $A$  and  $B$ ,  $1_{A \cup B} \geq (1_A + 1_B)/2$ . Then the conclusion follows from Lemma 1.

□

Based on the previous results, an upper and a lower bound are obtained on the moments of the nearest neighbor radius  $\hat{\tau}_{n,k,x}$ . A similar upper bound is stated as Lemma 3 in Leluc et al. (2023)).

**Lemma 3.** *Let  $q$  be a positive real number. Under (X1), (X2), and (X3), there exist two positive real numbers  $c_{q,d}$  and  $C_{q,d}$ , depending on  $q$ ,  $d$  and on the distribution of  $X$  such that*

$$c_{q,d} \frac{k^{q/d}}{(n+1)^{q/d}} \leq \mathbb{E} \hat{\tau}_{n,k,x}^q \leq C_{q,d} \frac{k^{q/d}}{(n+1)^{q/d}}. \quad (6)$$

A more precise expression of the two constants are

$$c_{q,d} = \frac{M_{2,d}^{-q/d}}{2\Gamma(1 + [q/d])}, \quad C_{q,d} = 2\Gamma(1 + [q/d]) M_{1,d}^{-q/d},$$

where  $[x]$  denotes the integer part of the real number  $x$ .

*Proof.* We have  $\hat{\tau}_{n,k,x} = Z_{(k)}(x)$  the  $k$ th-order statistics of  $Z_i(x) = |x - X_i|$ . Moreover, for any measurable and non-negative function  $f$ ,

$$\mathbb{E} f(Z_{(k)}(x)) = \mathbb{E} f \circ F_x^{-1}(U_{(k)}),$$

where  $U_{(k)}$  is the  $k$ th-order statistics of a  $n$  sample of uniform random variables and since  $F_x(z) = \mathbb{P}(X_1 \in B(x, z)) \in [M_{1,d}z^d, M_{2,d}z^d]$ ,

$$F_x^{-1}(u) = \inf \{z \in \mathbb{R} : F_x(z) \geq u\} \in \left[ \frac{u^{1/d}}{M_{2,d}^{1/d}}, \frac{u^{1/d}}{M_{1,d}^{1/d}} \right].$$

Note that the range of  $Z_i(x)$  is  $[0, \text{diam}(S_X)]$  and we use a constant  $c$  in the definition of  $M_{1,d} = cb_X V_d$  such that

$$\inf_{x \in S_X} \lambda_d(B(x, z)) \geq cV_d z^d \text{ for } 0 \leq z \leq \text{diam}(S_X).$$

If  $z = u^{1/d}/M_{1,d}^{1/d} \geq \text{diam}(S_X)$ , we have  $F_x(z) = 1 \geq u$  and we still have  $F_x^{-1}(u) \leq z$ .

Moreover, if  $g$  is measurable and nonnegative,

$$\mathbb{E}g(U_{(k)}) = n! \int g(u_k) \mathbf{1}_{0 \leq u_1 \leq \dots \leq u_n \leq 1} du_1 \cdots du_n = \frac{\Gamma(n+1)}{\Gamma(k)\Gamma(n-k+1)} \int_0^1 g(u) u^{k-1} (1-u)^{n-k} du.$$

When  $f(z) = z^q$  for some  $q > 0$ , we get

$$\begin{aligned} \mathbb{E}[Z_{(k)}^q] &\leq M_{1,d}^{-q/d} \mathbb{E}[U_{(k)}^{q/d}] = M_{1,d}^{-q/d} \frac{\Gamma(n+1)\Gamma(k+q/d)}{\Gamma(k)\Gamma(n+q/d+1)}, \\ \mathbb{E}[Z_{(k)}^q] &\geq M_{2,d}^{-q/d} \mathbb{E}[U_{(k)}^{q/d}] = M_{2,d}^{-q/d} \frac{\Gamma(n+1)\Gamma(k+q/d)}{\Gamma(k)\Gamma(n+q/d+1)}. \end{aligned}$$

For  $x \geq 1$  and  $s > 0$ , let  $N_{1,s} = \inf_{x \geq 1} \frac{\Gamma(x+s)}{x^s \Gamma(x)}$  and  $N_{2,s} = \sup_{x \geq 1} \frac{\Gamma(x+s)}{x^s \Gamma(x)}$ . We then get

$$M_{2,d}^{-q/d} \frac{N_{1,q/d}}{N_{2,q/d}} \frac{k^{q/d}}{(n+1)^{q/d}} \leq \mathbb{E}[Z_{(k)}(x)^q] \leq M_{1,d}^{-q/d} \frac{N_{2,q/d}}{N_{1,q/d}} \frac{k^{q/d}}{(n+1)^{q/d}}.$$

By Wendel's inequality (Wendel, 1948), for  $s \in (0, 1)$ , we have  $N_{1,s} \geq \inf_{x \geq 1} \left(\frac{x}{x+s}\right)^{1-s} \geq 1/2$  and  $N_{2,s} \leq 1$ . For  $s \geq 1$ , using the equality  $\Gamma(z+1) = z\Gamma(z)$ , one can deduce that

$$N_{1,s} \geq 1/2, \quad N_{2,s} \leq \Gamma(2 + [s]).$$

Indeed if  $s = s' + \ell$  with  $\ell \in \mathbb{N}$  and  $0 \leq s' < 1$ ,

$$\frac{\Gamma(x+s)}{x^s \Gamma(x)} = \prod_{j=1}^{\ell} \left(1 + \frac{j+s'-1}{x}\right) \frac{\Gamma(x+s')}{x^{s'} \Gamma(x)}$$

and

$$1 \leq \prod_{j=1}^{\ell} \left(1 + \frac{j+s'-1}{x}\right) \leq \prod_{j=1}^{\ell} (1+j) = \Gamma(\ell+2).$$

This completes the proof of Lemma 3.  $\square$

## B Proofs of the results on the marginal sampling error (Section 4.2)

### B.1 Proof of Proposition 1

The proof relies on the Lindeberg central limit theorem as given in Proposition 2.27 in Van der Vaart (2000) conditionally to  $\mathcal{F}_n$ . We need to show the two properties:

$$\begin{aligned} m^{-1} \sum_{i=1}^m \mathbb{E}[(h(X_i^*, Y_{n,i}^*) - \mathbb{E}[h(X_i^*, Y_{n,i}^*) | \mathcal{F}_n])^2 | \mathcal{F}_n] &\rightarrow V, \\ m^{-1} \sum_{i=1}^m \mathbb{E}[h(X_i^*, Y_{n,i}^*)^2 \mathbf{1}_{\{|h(X_i^*, Y_{n,i}^*)| > \epsilon\sqrt{n}\}} | \mathcal{F}_n] &\rightarrow 0, \end{aligned}$$

where each convergence needs to happen with probability 1. Equivalently, using that  $(X_i^*, Y_{n,i}^*)_{i=1, \dots, m}$  is identically distributed according to  $\hat{Q}$ , we need to show that

$$\begin{aligned} \hat{Q}(h^2) - \hat{Q}(h)^2 &\rightarrow V, \\ \hat{Q}(h^2 \mathbf{1}_{\{|h| > \epsilon\sqrt{n}\}}) &\rightarrow 0 \text{ for each } \epsilon > 0. \end{aligned}$$

The first result is a direct consequence of the assumption. Fix  $M > 0$ . For all  $n$  sufficiently large, we have  $M \leq \epsilon\sqrt{n}$ , implying that

$$\hat{Q}(h^2 \mathbf{1}_{\{|h| > \epsilon\sqrt{n}\}}) \leq \hat{Q}(h^2 \mathbf{1}_{\{|h| > M\}}),$$

which converges to  $Q(h^2 \mathbf{1}_{\{|h| > M\}})$  by assumption. Since  $Q(h^2)$  is finite, one can choose  $M$  large enough to make  $Q(h^2 \mathbf{1}_{\{|h| > M\}})$  arbitrarily small.

## B.2 Proof of Proposition 2

Set  $Z_{n,i}^* = h(Y_{n,i}^*, X_i^*) - \hat{Q}(h)$ . We have  $|Z_{n,i}^*| \leq 2U_h$  and  $\text{Var}(Z_{n,i}^* | \mathcal{F}_n) = \hat{v}_n$ . Note that  $\hat{Q}^*(h) - \hat{Q}(h) = \frac{1}{m} \sum_{i=1}^m Z_{n,i}^*$ . Bernstein's concentration inequality leads to

$$\mathbb{P}\left(\left|\hat{Q}^*h - \hat{Q}h\right| > u \mid \mathcal{F}_n\right) \leq \exp\left(-\frac{1/2u^2m}{\hat{Q}(h^2) - (\hat{Q}(h))^2 + 2/3U_hu}\right).$$

Then setting

$$\hat{u}_h(\delta) = \frac{4/3U_h}{m} \log(2/\delta) + \sqrt{2\frac{\hat{Q}h^2 - (\hat{Q}h)^2}{m} \log(2/\delta)},$$

we get

$$\mathbb{P}\left(\left|\hat{Q}^*h - \hat{Q}h\right| > \hat{u}_h(\delta) \mid \mathcal{F}_n\right) \leq \delta$$

and then integrate both sides to obtain

$$\mathbb{P}\left(\left|\hat{Q}^*h - \hat{Q}h\right| > \hat{u}_h(\delta)\right) \leq \delta,$$

which leads to the stated bound.

## C Proofs of the results on the $k$ -NN conditional sampling error (Section 4.3)

Here, we give proofs of the results on the  $k$ -NN conditional sampling error appearing in Section 4.3.

### C.1 Proof of Proposition 3

We start with a useful bias-variance decomposition. Introduce

$$\begin{aligned} \epsilon_i(x) &= h(Y_i, x) - \int h(y, x) P_{Y|X}(dy | X_i), \\ \Delta(x, X_i) &= \int h(y, x) (P_{Y|X}(dy | X_i) - P_{Y|X}(dy | x)). \end{aligned}$$

We have

$$\begin{aligned} & \int h(y, x) (\hat{P}_{Y|X}(dy | x) - P_{Y|X}(dy | x)) \\ &= k^{-1} \sum_{i=1}^n \epsilon_i(x) 1_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i) + k^{-1} \sum_{i=1}^n \Delta(x, X_i) 1_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i). \end{aligned}$$

Integrating with respect to  $Q_X(dx)$ , we obtain

$$(\hat{Q} - Q)(h) = B_h + S_h \tag{7}$$

with

$$\begin{aligned} B_h &= k^{-1} \sum_{i=1}^n \int \Delta(x, X_i) 1_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i) Q_X(dx), \\ S_h &= k^{-1} \sum_{i=1}^n \int \epsilon_i(x) 1_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i) Q_X(dx). \end{aligned}$$

The term  $B_h$  is a bias term and the term  $S_h$  (which has mean 0) is a variance term.

The proof is divided into 3 steps. The first step takes care of bounding the bias term. The second step deals with the variance upper-bound. The third step is concerned with the variance lower bound.

**The bias.** First (H1) gives that for any  $X_i \in S_X \cap B(x, \hat{\tau}_{n,k,x})$  and  $x \in S_X$ ,

$$|\Delta(x, X_i)| \leq g_h(x) \hat{\tau}_{n,k,x}.$$

Consequently, using (X4) and the fact that  $\sum_{i=1}^n \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} = k$ , we have

$$\left| k^{-1} \sum_{i=1}^n \int \Delta(x, X_i) \mathbf{1}_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i) Q_X(dx) \right| \leq \int \hat{\tau}_{n,k,x} g_h(x) Q_X(dx)$$

and from Jensen inequality

$$\left| k^{-1} \sum_{i=1}^n \int \Delta(x, X_i) \mathbf{1}_{\{B(x, \hat{\tau}_{n,k,x})\}}(X_i) Q_X(dx) \right|^2 \leq \int \hat{\tau}_{n,k,x}^2 g_h^2(x) Q_X(dx)$$

From Lemma 3, we have  $\sup_{x \in S_X} \mathbb{E}[\hat{\tau}_{n,k,x}^2] \leq C_{2,d} k^{2/d} (n+1)^{-2/d}$  and the control of the bias is given by

$$\mathbb{E}[|B_h^2|] \leq C_{2,d} \int g_h^2(x) Q_X(dx) \cdot \frac{k^{2/d}}{(n+1)^{2/d}}.$$

**The variance upper-bound.** For the proof, we assume that  $1 \leq k < n/2$ . We have for each  $(x, x') \in S_X^2$  and  $(i, j) \in \{1, \dots, n\}^2$ ,

$$\mathbb{E}[\epsilon_i(x) \epsilon_j(x') \mid X_1, \dots, X_n] = \begin{cases} 0 & \text{if } i \neq j, \\ \mathbb{E}[\epsilon_i(x) \epsilon_i(x') \mid X_i] \leq \sqrt{\mathbb{E}[\epsilon_i(x)^2 \mid X_i] \mathbb{E}[\epsilon_i(x')^2 \mid X_i]} \leq \sigma_+^2 & \text{if } i = j. \end{cases}$$

For the second case, we used (H2) and the Cauchy-Schwarz inequality. As a consequence, the variance is given by

$$\begin{aligned} \mathbb{E}[S_h^2] &= \mathbb{E} \left[ \left( k^{-1} \sum_{i=1}^n \int \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} \epsilon_i(x) Q_X(dx) \right)^2 \right] \\ &= k^{-2} \sum_{i,j}^n \mathbb{E} \left[ \left( \int \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} \epsilon_i(x) Q_X(dx) \right) \left( \int \mathbf{1}_{\|X_j - x'\| \leq \hat{\tau}_{n,k,x'}} \epsilon_j(x') Q_X(dx') \right) \right] \\ &= k^{-2} \sum_{i,j}^n \mathbb{E} \left[ \left( \int \int \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} \mathbf{1}_{\|X_j - x'\| \leq \hat{\tau}_{n,k,x'}} \mathbb{E}[\epsilon_i(x) \epsilon_j(x') \mid X_1, \dots, X_n] Q_X(dx) Q_X(dx') \right) \right] \\ &\leq k^{-2} \sigma_+^2 \sum_{i=1}^n \mathbb{E} \left[ \int \int \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} \mathbf{1}_{\|X_i - x'\| \leq \hat{\tau}_{n,k,x'}} Q_X(dx) Q_X(dx') \right] \\ &= k^{-2} \sigma_+^2 \sum_{i=1}^n \mathbb{E}[\hat{Y}_i^2] \end{aligned}$$

with  $\hat{Y}_i = \int \mathbf{1}_{\|X_i - x\| \leq \hat{\tau}_{n,k,x}} Q_X(dx)$ . Let  $Z_i(x) = \|x - X_i\|$  and  $Z_{(k)}(x)$ , the  $k$ -th order statistics of the sample  $(Z_i(x))_{1 \leq i \leq n}$ . One can observe that

$$Z_i(x) \leq Z_{(k)}(x) \iff Z_i(x) < Z_{(k)}^{-i}(x),$$

where  $Z_{(k)}^{-i}(x)$  is the  $k$ -th order statistics of the sample  $(Z_j(x))_{1 \leq j \neq i \leq n}$ . Note that the two sigma fields generated respectively by  $\{Z_i(x) : x \in S_X\}$  and  $\{Z_{(k)}^{-i}(x) : x \in S_X\}$  are independent. For one mapping  $\rho : S_X \rightarrow \text{diam}(S_X)$ ,



we first bound

$$\begin{aligned}
I &= \mathbb{E} \left| \int \mathbf{1}_{Z_1(x) \leq \rho(x)} Q(dx) \right|^2 \\
&= \int \int \mathbb{P}(Z_1(x) \leq \rho(x), Z_1(y) \leq \rho(y)) Q(dx)Q(dy) \\
&= 2 \int \int \mathbf{1}_{\{\rho(x) \leq \rho(y)\}} \mathbb{P}(Z_1(x) \leq \rho(x), Z_1(y) \leq \rho(y)) Q(dx)Q(dy) \\
&= 2 \int \int \mathbf{1}_{\{\rho(x) \leq \rho(y), \|x - y\| \leq 2\rho(y)\}} \mathbb{P}(Z_1(x) \leq \rho(x), Z_1(y) \leq \rho(y)) Q(dx)Q(dy) \\
&\leq 2M_{2,d} \int \int \mathbf{1}_{\{\rho(x) \leq \rho(y), \|x - y\| \leq 2\rho(y)\}} \rho(y)^d Q(dx)Q(dy) \\
&\leq 2M_{2,d} \int Q(B(y, 2\rho(y))) \rho(y)^d Q(dy) \\
&\leq 2^{d+1} M_{2,d}^2 \int \rho(y)^{2d} Q(dy).
\end{aligned}$$

The fourth inequality is due to the fact that when  $\|x - y\| > 2\rho(y)$ , the two balls  $B(x, \rho(x))$  and  $B(y, \rho(y))$  do not intersect. We then get using Lemma 3,

$$E\hat{Y}_i^2 \leq 2^{d+1} M_{2,d}^2 \int \mathbb{E} \left\{ Z_{(k)}^{-i}(y) \right\}^{2d} Q(dy) \leq \frac{2^{d+3} M_{2,d}^2 k^2}{M_{1,d}^2 n^2}.$$

This leads to the variance upper-bound.

**The variance lower-bound.** Here we assume the  $\sigma_-^2 := \inf_{x \in S_X} \text{Var}(h(Y)|X = x) > 0$ . From previous computations, we have

$$\text{Var}(\hat{Q}h) \geq \frac{\sigma_-^2}{k^2} \sum_{i=1}^n \mathbb{E}[\hat{Y}_i^2].$$

We have  $\mathbb{E}\hat{Y}_i^2 \geq \mathbb{E}^2\hat{Y}_i$  and we have to find a lower bound for  $\mathbb{E}\hat{Y}_i$ . But from the arguments used in the proof for the upper-bound, we have

$$\mathbb{E}\hat{Y}_i = \int_{S_X} \mathbb{P}(Z_i(x) \leq Z_{(k)}^{-i}(x)) Q(dx) \geq M_{1,d} \int_{S_X} \mathbb{E} \left| Z_{(k)}^{-i}(x) \right|^d Q(dx) \geq \frac{k}{2nM_{2,d}},$$

where the last inequality follows from Lemma 3. This shows the lower-bound and the proof of Proposition 3 is now complete.

## C.2 Proof of Proposition 4

Define

$$\bar{\tau}_{n,k} = \left( \frac{2k}{nM_{1,d}} \right)^{1/d}.$$

The following Lemma (Portier, 2021, Lemma 4) controls the size of the  $k$ -NN balls uniformly over all  $x \in S_X$ .

**Lemma 4** (Portier (2021, Lemma 4)). *Suppose that (X1) (X2) and (X3) hold true. Then for all  $n \geq 1$ ,  $\delta \in (0, 1)$  and  $1 \leq k \leq n$  such that  $24d \log(12n/\delta) \leq k \leq T^d n b_X c V_d / 2$ , it holds, with probability at least  $1 - \delta$ :*

$$\sup_{x \in S_X} \hat{\tau}_{n,k,x} \leq \bar{\tau}_{n,k}.$$

We now deal with the variance term of our estimator. The variance term of the nearest-neighbors estimator is given by  $V_n = k^{-1} \sum_{i=1}^n \hat{Y}_i$ , where

$$\hat{Y}_i = \int \epsilon_i(x) \mathbf{1}_{B(x, \hat{\tau}_{n,k,x})}(X_i) Q_X(dx).$$

and Set  $\hat{\tau}_k = \sup_{x \in S_X} \hat{\tau}_{n,k,x}$ . From our assumptions and Jensen's inequality, we have

$$|\hat{Y}_i| \leq 2U_h M_{2,d} \hat{\tau}_k, \quad \text{Var}(\hat{Y}_i | X_1, \dots, X_n) \leq \sigma_+^2 M_{2,d}^2 \hat{\tau}_k^{2d}.$$

Applying Bernstein's inequality for i.i.d. random variables (we recall that the  $Y_i$ 's are independent conditionally on the  $X_i$ 's), we get for  $t > 0$ ,

$$\mathbb{P}(|V_n| > t | X_1, \dots, X_n) \leq 2 \exp\left(-\frac{\frac{1}{2}k^2 t^2}{n\sigma_+^2 M_{2,d}^2 \hat{\tau}_k^{2d} + \frac{2}{3}U_h M_{2,d} \hat{\tau}_k^d k}\right).$$

This leads to

$$\mathbb{P}(|V_n| > t, \hat{\tau}_k \leq \bar{\tau}_{n,k} | X_1, \dots, X_n) \leq 2 \exp\left(-\frac{\frac{1}{2}t^2 n}{L_1 \sigma_+^2 + L_2 t}\right).$$

Note that this upper-bound is not random and we get

$$\mathbb{P}(|V_n| > t, \hat{\tau}_k \leq \bar{\tau}_{n,k}) \leq 2 \exp\left(-\frac{\frac{1}{2}t^2 n}{L_1 \sigma_+^2 + L_2 t}\right). \quad (8)$$

Setting

$$\tilde{t}_n(\delta, h) = \frac{L_2}{n} \log(2/\delta) + \sqrt{\frac{L_2^2}{n^2} \log^2(2/\delta) + \frac{2L_1 \sigma_+^2}{n} \log(2/\delta)},$$

which is smaller than

$$t_n(\delta, h) = L_0 \left(\frac{k}{n}\right)^{1/d} + \frac{2L_2}{n} \log(2/\delta) + \sqrt{\frac{2L_1 \sigma_+^2}{n} \log(2/\delta)}$$

we then get

$$\begin{aligned} \mathbb{P}(|V_n| > t_n(\delta, h)) &\leq \mathbb{P}(|V_n| > \tilde{t}_n(\delta, h), \hat{\tau}_k \leq \bar{\tau}_{n,k}) + P(\hat{\tau}_k > \bar{\tau}_{n,k}) \\ &\leq 2\delta, \end{aligned}$$

where the last inequality is a consequence of (8) and Lemma 4. Moreover, from the proof of Theorem 3 and Lemma 4, the bias part can be dominated by  $L_0(k/n)^{1/d}$  with probability at least  $1 - \delta$ . This concludes the proof.

## D Proof of Theorem 1

First, note that the boundedness of  $\sup_{x \in S_X} \mathbb{E}[h^2(Y, x) | X]$  entails **H2**. Setting  $A_{m,n} = \hat{Q}^*(h) - \hat{Q}(h)$  and  $B_n = \hat{Q}(h) - Q(h)$ , Proposition 3 guarantees that

$$\mathbb{E}B_n^2 \leq C_1 \left\{ \frac{1}{n^{2/d}} + \frac{1}{n} \right\},$$

for some  $C_1 > 0$  only depending on the distribution of  $(X, Y)$ ,  $X^*$  and on  $h$ . It remains to show that the same bound can be obtained for  $A_{m,n}$ . Since,

$$\mathbb{E}[A_{m,n}^2 | \mathcal{F}_n] \leq \frac{\hat{Q}(h^2) - \hat{Q}(h)^2}{m} \leq \frac{\hat{Q}(h^2)}{m}.$$

It only remains to show that  $\mathbb{E}\hat{Q}(h^2)$  is bounded with respect to  $n$ . The approach is similar to the control of the variance term for  $k = 1$  studied in the proof of Proposition 3. If  $L$  is an upper-bound for  $\sup_{x \in S_X} \mathbb{E}[h^2(Y, x) | X]$ , we have

$$\begin{aligned} \mathbb{E}\hat{Q}(h^2) &= \sum_{i=1}^n \int_{S_X} Q_X(dx) \mathbb{E}[h^2(Y_1, x) \mathbb{1}_{\|X_1 - x\| \leq \min_{2 \leq j \leq n} \|X_j - x\|}] \\ &\leq nL \int_{S_X} \mathbb{P}\left(\|X_1 - x\| \leq \min_{2 \leq j \leq n} \|X_j - x\|\right) Q_X(dx) \\ &\leq nLM_{2,d} \int_{S_X} \mathbb{E} \min_{2 \leq j \leq n} \|X_j - x\|^d Q_X(dx) \\ &\leq \frac{2LM_{2,d}}{M_{1,d}}. \end{aligned}$$

The last upper bound is obtained from Lemma 3 using the fact that  $\min_{2 \leq j \leq n} \|X_j - x\|$  has the same probability distribution as  $\hat{\tau}_{n-1,1,x}$ . We deduce the result taking  $C$  as the maximum between  $C_1$  and  $2LM_{2,d}/M_{1,d}$ .  $\blacksquare$

## E A corollary bounding the sampling error for $k$ -NN sampling

**Corollary 1.** *Suppose that (X1) (X2) (X3) (X4) are fulfilled and that (H1) and (H2) are fulfilled for both  $h$  and  $h^2$  and  $h$  is bounded. Let  $k = k_n$  satisfying the condition of Lemma 4 and  $\delta \in (0, 1/7)$  and set  $s^2(h) = Qh^2 - (Qh)^2$  and  $\hat{s}^2(h) = \hat{Q}h^2 - (\hat{Q}h)^2$ . Let  $\delta \in (0, 1/7)$ . Then with probability greater than  $1 - 7\delta$ ,*

$$\hat{Q}^*h - \hat{Q}h \leq \frac{4/3U_h}{m} \log(2/\delta) + \sqrt{2 \frac{v_n^2(\delta, h)}{m} \log(2/\delta)} + \sqrt{2 \frac{s^2(h)}{m} \log(2/\delta)}, \quad (9)$$

where

$$v_n^2(\delta, h) = t_n(\delta, h^2) + t_n(\delta, h)^2 + 2t_n(\delta, h)Q|h|$$

and  $t_n(\delta, h)$  is defined in the statement of Proposition 4.

*Proof.* We first use the result of Proposition 2. In particular, setting

$$\hat{u}_h(\delta) = \frac{4/3U_h}{m} \log(2/\delta) + \sqrt{2 \frac{\hat{Q}h^2 - (\hat{Q}h)^2}{m} \log(2/\delta)},$$

we have

$$\mathbb{P} \left( \left| \hat{Q}^*h - \hat{Q}h \right| > \hat{u}_h(\delta) \right) \leq \delta,$$

We then use the decomposition

$$\hat{Q}h^2 - (\hat{Q}h)^2 = s^2(h) + \hat{Q}h^2 - Qh^2 - \left( \hat{Q}h - Qh \right)^2 - 2Qh \left( \hat{Q}h - Qh \right).$$

From Proposition 4, we know that

$$\left| \hat{Q}h^2 - Qh^2 \right| \leq t_n(\delta, h^2)$$

with probability greater than  $1 - 3\delta$  and

$$\left| \hat{Q}h - Qh \right| \leq t_n(\delta, h)$$

with probability greater than  $1 - 3\delta$ . Collecting these three bounds, we easily obtain the conclusion of the second point of Corollary 1.  $\square$

## F Proofs of the results on the empirical risk minimization (Section 5)

Here, we present proofs of the result on the application to empirical risk minimization.

### F.1 Proof of Theorem 2.

From (A1),  $\sup_{\theta \in \Theta} |m_\theta(Y_1^*, X_1^*)|$  is integrable and  $\theta \mapsto \mathcal{R}^*(\theta)$  is continuous over the compact set  $\Theta$ . As a consequence, weak consistency will follow from Theorem 5.7 in Van der Vaart (2000) if we show that

$$\sup_{\theta \in \Theta} \left| \mathcal{R}_{m,n}^*(\theta) - \mathcal{R}^*(\theta) \right| = o_{\mathbb{P}}(1). \quad (10)$$

Pointwise convergence holds true from assumptions (A1), (A2) as each mapping  $m_\theta$  satisfies Assumptions (H1), (H2). One can then apply Theorem 1 and the Markov inequality to get for any  $\theta \in \Theta$ ,

$$\mathcal{R}_{m,n}^*(\theta) - \mathcal{R}^*(\theta) = o_{\mathbb{P}}(1), \quad m, n \rightarrow \infty.$$

We now prove uniform convergence. Let  $\delta > 0$ . One can cover the compact set  $\Theta$  with finitely many balls  $B(\theta_i, \delta)$ ,  $1 \leq i \leq k$ . For  $\theta \in \Theta \cap B(\theta_i, \delta)$ , we have

$$\left| \mathcal{R}_{m,n}^*(\theta) - \mathcal{R}_{m,n}^*(\theta_i) \right| \leq \eta(\delta) \frac{1}{m} \sum_{i=1}^m h(Y_{n,i}^*).$$

Moreover, from Assumptions (A2) and Theorem 1 with the Markov inequality, we know that

$$\frac{1}{m} \sum_{i=1}^m h(Y_{n,i}^*) \rightarrow \mathbb{E}h(Y_1^*), \quad m, n \rightarrow \infty,$$

in probability. We also have

$$|\mathcal{R}^*(\theta) - \mathcal{R}^*(\theta_i)| \leq \eta(\delta) \mathbb{E}h(Y_1^*).$$

Finally, one can use the bound

$$\begin{aligned} \sup_{\theta \in \Theta} |\mathcal{R}_{m,n}^*(\theta) - \mathcal{R}^*(\theta)| &\leq \max_{1 \leq i \leq k} |\mathcal{R}_{m,n}^*(\theta_i) - \mathcal{R}^*(\theta_i)| + \eta(\delta) \left\{ \frac{1}{m} \sum_{i=1}^m h(Y_{n,i}^*) + \mathbb{E}h(Y_1^*) \right\} \\ &= 2\eta(\delta) \mathbb{E}h(Y_1^*) + o_{\mathbb{P}}(1). \end{aligned}$$

Given that  $\delta$  is arbitrary, the above implies (10) and the weak consistency of  $\hat{\theta}^*$  follows. The second assertion about the excess risk then follows easily using that

$$\mathcal{R}^*(\hat{\theta}^*) - \mathcal{R}^*(\theta^*) \leq 2 \sup_{\theta \in \Theta} |\mathcal{R}_{m,n}^*(\theta) - \mathcal{R}^*(\theta)|.$$

■

## F.2 Proof of Theorem 3.

Let  $Z_i^* = \Gamma^{-1/2} X_i^*$  and define

$$\Sigma_m = \frac{1}{m} \sum_{i=1}^m Z_i^* Z_i^{*T}, \quad N_m = \frac{1}{m} \sum_{i=1}^m Z_i^* [Y_{n,i}^* - X_i^{*T} \theta^*].$$

The proof first requires some analysis of the smallest eigenvalues of  $\Sigma_m$ . From the matrix Chernoff inequality given in Tropp (2012), see Corollary 5.2 and Remark 5.3, we have

$$P(\lambda_{\min}(\Sigma_m) \leq 1 - \eta) \leq d \exp(-\eta^2 m / 2B)$$

where  $B = \lambda_{\min}(\Gamma)^{-1} d \max_{x \in S_X} |x|_{\infty}^2$  is defined so as to satisfy  $\|Z\|_2^2 \leq \lambda_{\min}(\Gamma)^{-1} \|X\|_2^2 \leq B$ , with probability 1. Inverting the previous we obtain that with probability at least  $1 - \delta$ ,

$$\lambda_{\min}(\Sigma_m) > 1 - \sqrt{(2B/m \log(d\delta^{-1}))}$$

and therefore as soon as  $(8B/n) \log(d\delta^{-1}) \leq 1$ , we have that  $\lambda_{\min}(\Sigma_m) \geq 1/2$ . On the previous event, we have that

$$\Gamma^{1/2}(\hat{\theta}^* - \theta^*) = \Sigma_m^{-1} N_m$$

It follows that

$$\mathcal{R}^*(\hat{\theta}^*) - \mathcal{R}^*(\theta^*) = \|\Gamma^{1/2}(\hat{\theta}^* - \theta^*)\|_2^2 = \|\Sigma_m^{-1} N_m\|_2^2 \leq 2\|N_m\|_2^2.$$

We conclude using Theorem 1 with  $h(y, x) = \Gamma^{-1/2} x (y - x^T \theta^*)$ . Note that  $\mathbb{E}[h(Y^*, X^*)] = 0$  by definition of  $\theta^*$ .

■

## G Illustration for Experiments E1–E3

Illustrations of data used in Experiments E1–E3 can be found in Figure 7b.

## H Details of the benchmark data experiments

We use the following datasets.

- **california**: Regression dataset called ‘‘California Housing’’ available from [https://www.dcc.fc.up.pt/~ltorgo/Regression/cal\\_housing.html](https://www.dcc.fc.up.pt/~ltorgo/Regression/cal_housing.html).
- **diabetes**: Regression dataset available from <https://archive.ics.uci.edu/ml/index.php> (Dua and Graff, 2017).
- **breast cancer**: Classification dataset available from <https://archive.ics.uci.edu/ml/index.php> (Dua and Graff, 2017).
- **twonorm**: Classification dataset available from <https://www.cs.utoronto.ca/~delve/data/datasets.html>.

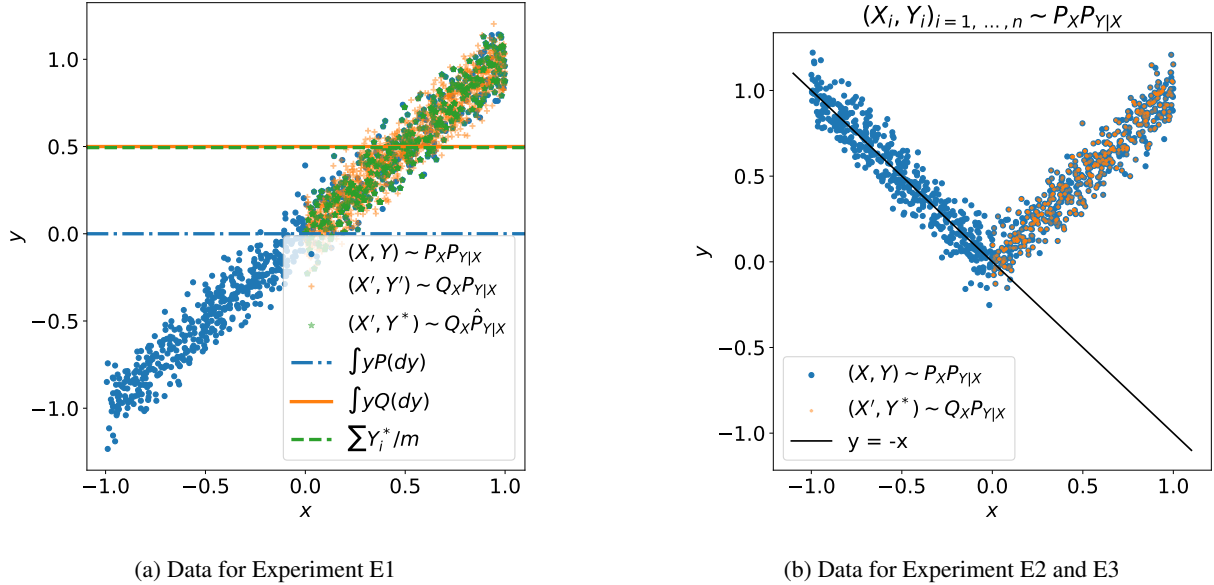


Figure 7: Visualization of data for Experiments E1–E3

**Data splitting and sampling bias simulation** We split the original to the training and test set and simulate covariate shift by rejection sampling from the test set with rejection probability determined according to the value of a covariate. For `california`, `twonorm`, `breast cancer`, we follow the procedure of Sugiyama et al. (2007): we include each target data point  $X_i$  to the target set with probability  $\min(1, 4X_{i,c}^2)$  or reject it otherwise, where  $X_{i,c}$  is the  $c$ -th attribute of  $X_i$ . For `diabetes`, we used a different biasing procedure for this data set because the technique of Sugiyama et al. (2007) rejects too many data points to perform our experiment for this dataset. We instead use the procedure of an example from the ADAPT package de Mathelin et al. (2021)<sup>7</sup> for `diabetes`: for each data point  $X_i$ , we accept it with probability proportional to  $\exp(-20 \times |X_{i,\text{age}} + 0.06|)$ , where  $X_{i,\text{age}}$  is the age attribute of  $X_i$  and reject (i.e., exclude) otherwise.

**Pre-processing** We use the hot-encoding for all categorical features. We center and normalize all the data using the mean and the dimension-wise standard deviation of the source set. We do the same centering and normalization for the output variables for regression datasets.

After training and prediction, we post-process the output using the inverse operation. Table 3 shows basic information about the datasets after the bias-sampling and pre-processing.

Table 3: Basic information of the datasets

	california	twonorm	diabetes	breast cancer
Input dimension $d$	8	20	10	9
source sample size $n$	1000	100	150	200
Target sample size $m$	1000	500	150	100

<sup>7</sup>[https://adapt-python.github.io/adapt/examples/Sample\\_bias\\_example.html](https://adapt-python.github.io/adapt/examples/Sample_bias_example.html)