

Class Prior Estimation under Covariate Shift: No Problem?

Dirk Tasche^[0000–0002–2750–2970]

Independent researcher
`dirk.tasche@gmx.net`

Abstract. We show that in the context of classification the property of source and target distributions to be related by covariate shift may be lost if the information content captured in the covariates is reduced, for instance by dropping components or mapping into a lower-dimensional or finite space. As a consequence, under covariate shift simple approaches to class prior estimation in the style of classify and count with or without adjustment are infeasible. We prove that transformations of the covariates that preserve the covariate shift property are necessarily sufficient in the statistical sense for the full set of covariates. A probing algorithm as alternative approach to class prior estimation under covariate shift is proposed.

Keywords: Covariate shift · Prior probability shift · Quantification · Class prior estimation · Prevalence estimation · Sufficiency

1 Introduction

Class prior estimation (also known as quantification, class distribution estimation, prevalence estimation etc.) may be considered one of the tasks referred to under the general term domain adaptation.

Domain adaptation means adapting algorithms designed for a source (training) dataset (also distribution or domain) to a target (test) dataset. The source and target distributions may be different, a phenomenon which is called *dataset shift*. In this paper, attention is restricted to ‘unsupervised’ domain adaptation. This term refers on the one hand to the situation where under the source distribution all events and realisations of random variables – including the target (label) variable – are observable such that in principle the whole distribution can be estimated. On the other hand under the target distribution only the marginal distribution of the covariates (features) can be observed, via realisations of the covariates. The target distribution class labels cannot be observed at all or only with delay.

Moreno-Torres et al. [21] proposed the following popular taxonomy of types of dataset shift:

- Covariate shift: Source and target posterior class probabilities are the same but source and target covariate distributions may be different.

- Prior probability shift (label shift [20], global drift [14]): Source and target class-conditional covariate distributions are the same but source and target prior class probabilities may be different.
- Concept shift: Source and target covariate distributions are the same but source and target posterior class probabilities may be different, or source and target prior class probabilities are the same but source and target class-conditional covariate distributions may be different.
- Other shift: Any dataset shift not captured by the previous types.

Covariate shift and prior probability shift are described in constructive terms. Based on their defining properties source and target distributions are fully specified. For this reason, a host of focussed literature is available for these two types of dataset shift. In contrast, it is hardly possible to make specific statements about the two other types of shift such that the literature on these types is much more diverse and hard to capture.

In this paper we focus on covariate shift and a classification setting. We choose a measure-theoretic approach that is particularly suitable for this context as it facilitates a rigorous joint treatment of continuous and discrete random variables, or covariates and class labels more specifically. We work in the same binary classification setting as Ben-David et al. [3] and Johansson et al. [16]. Like Ben-David et al. and Johansson et al., we focus on the binary case but the results are easily generalised to the multi-class case.

Prior probability shift is robust in the following sense: If the set of covariates is transformed in a way that reduces the information reflected by them (e.g. by dropping components or mapping into a lower-dimensional or finite space) then the resulting source and target joint distributions of covariates and labels are still related by prior probability shift. As a consequence, simple approaches to class prior estimation under prior probability shift can be designed which avoid the need to estimate the full class-conditional covariate distributions.¹ The primary example for such an approach is the ‘confusion matrix method’ (Gart and Buck [9]; Saerens et al. [23]; ‘adjusted count’ in Forman [8]).

We show by examples and by theoretical analysis that such robustness is not displayed by covariate shift. Under the condition that the target distribution is absolutely continuous with respect to the source distribution, we prove that a set of covariates passes on the covariate shift property if and only if the transformed set of covariates is ‘sufficient’ in the sense of Adraghi and Cook [1] and Tasche [26] for the untransformed set under the source distribution. The result refines an observation of Johansson et al. [16] who found that covariate shift was inherited “only if” the transformation was invertible.

An important consequence of this finding is that in general for class prior estimation under covariate shift, simplification in the sense of reducing the complexity of the covariate set is not a viable path because the covariate shift property of identical posterior class probabilities between source and target distributions

¹ The simplification may come at a cost of increased variance of the estimator (Tasche [27]).

might get lost. We point to a potential alternative approach, based on the so-called ‘probing’ method of Langford and Zadrozny [19].

The plan of this paper is as follows: We introduce the assumptions and the notation for this paper in Section 2. In Section 3 we give examples of how loss of information may affect the covariate shift property. The main result (Theorem 1 of this paper) is presented in Section 4 while Section 5 provides some comments on the result. A proposal for applying ‘probing’ to class prior estimation is made in Section 6. The paper concludes with a short summary in Section 7.

2 Assumptions and Notation

In this paper, we work only at population (distribution) level as this level is appropriate for the design of estimators and predictors as well as the study of their fundamental properties. A detailed treatment of the intricacies of sample properties is not needed.

We follow the example of Scott [24] who introduced consistent concepts and notation for appropriately dealing with the classification setting we need. As the concept of information plays a more important role in this paper than in Scott’s, we dive somewhat deeper into the measure-theoretic details of the setting than Scott.

2.1 Setting for Binary Classification in the Presence of Dataset Shift

We introduce a measure-theoretic setting, expanding the setting of Scott [24] and adapting the approach of Holzmann and Eulert [15] and Tasche [26]. Phrasing the context in measure theory terms is particularly efficient when random variables with continuous and discrete distributions are studied together like in the case of binary or multi-class classification. Moreover, the measure-theoretic notion of σ -algebras allows for the convenient description of differences in available information.

We use the following population-level description of the binary classification problem in terms of measure theory. See standard textbooks on probability theory like Billingsley [4] or Klenke [17] for formal definitions and background of the notions introduced in Assumption 1.

Assumption 1 *(Ω, \mathcal{A}) is a measurable space. The source distribution P and the target distribution Q are probability measures on (Ω, \mathcal{A}) . An event $A_1 \in \mathcal{A}$ with $0 < P[A_1] < 1$ and a sub- σ -algebra $\mathcal{H} \subset \mathcal{A}$ with $A_1 \notin \mathcal{H}$ are fixed. $A_0 = \Omega \setminus A_1$ is the complementary event of A_1 in Ω .*

In the literature, P is also called ‘training distribution’ while Q is also referred to as ‘test distribution’.

Interpretation. The elements ω of Ω are objects (or instances) with class (label) and covariate (feature) attributes. $\omega \in A_1$ means that ω belongs to class

1 (or the positive class). $\omega \in A_0$ means that ω belongs to class 0 (or the negative class).

The σ -algebra \mathcal{A} of events $M \in \mathcal{A}$ is a collection of subsets M of Ω with the property that they can be assigned probabilities $P[M]$ and $Q[M]$ in a logically consistent way. In the literature, thanks to their role of reflecting the available information, σ -algebras are sometimes also called *information set* (Holzmann and Eulert [15]). In the following, we use both terms exchangeably.

Binary classification problem. The sub- σ -algebra $\mathcal{H} \subset \mathcal{A}$ contains the events which are observable at the time when the class label of an object ω has to be predicted. Since $A_1 \notin \mathcal{H}$, then the class of an object may not yet be known. It can only be predicted on the basis of the events $H \in \mathcal{H}$ which are assumed to reflect the features of the object.

Dataset shift. We denote by \mathcal{H}_A the minimal sub- σ -algebra of \mathcal{A} containing both \mathcal{H} and $\sigma(\{A_1\}) = \{\emptyset, A_1, A_0, \Omega\}$, i.e. $\mathcal{H}_A = \sigma(\mathcal{H} \cup \sigma(\{A_1\}))$. The σ -algebra \mathcal{H}_A can be represented as

$$\mathcal{H}_A = \{(A_1 \cap H_1) \cup (A_0 \cap H_0) : H_1, H_0 \in \mathcal{H}\}. \quad (1)$$

A standard assumption in machine learning is that source and target distribution are the same, i.e. $P = Q$. The situation where $P[M] \neq Q[M]$ holds for at least one $M \in \mathcal{H}_A$ is called *dataset shift* (Moreno-Torres et al. [21], Definition 1).

Class prior estimation. Under dataset shift as defined above, typically the prior probabilities $P[A_1]$ of the positive class in the source distribution (assumed to be observable) and $Q[A_1]$ in the target distribution (assumed to be unknown or known with delay only) are different. Class prior estimation in the binary classification context of Assumption 1 is the task to estimate $Q[A_1]$, based on observations from P (the entire source distribution) and from² $Q|\mathcal{H}$ (the target distribution of the covariates, also called features).

Notation. Denote by $\mathbf{1}_M$ the indicator function of an event M , i.e. $\mathbf{1}_M(\omega) = 1$ if $\omega \in M$ and $\mathbf{1}_M(\omega) = 0$ if $\omega \notin M$.

If X is a real-valued random variable on a probability space (Ω, \mathcal{F}, P) and $\mathcal{G} \subset \mathcal{F}$ is a sub- σ -algebra of \mathcal{F} , then a random variable Ψ is called *expectation of X conditional on \mathcal{G}* (see, e.g., Definition 8.11 of Klenke [17]) if it has the following two properties:

- (i) Ψ is \mathcal{G} -measurable.
- (ii) For all events $G \in \mathcal{G}$ it holds that $E_P[\mathbf{1}_G X] = E_P[\mathbf{1}_G \Psi]$.

In the following, we use the usual shorthand notation $\Psi = E_P[X | \mathcal{G}]$. In the case of an indicator function of an event $F \in \mathcal{F}$, the conditional expectation $E_P[\mathbf{1}_F | \mathcal{G}]$ is called *probability of F conditional on \mathcal{G}* and denoted by $P[F | \mathcal{G}]$.

2.2 Reconciliation of Machine Learning and Measure Theory Settings

The setting of Assumption 1 is similar to a standard setting for binary classification in the machine learning and pattern recognition literature (see e.g. Scott [24] or Devroye et al. [7]):

² $Q|\mathcal{H}$ stands for the measure Q with domain restricted to \mathcal{H} .

Typically a random vector (X, Y) is studied, where X stands for the covariates of an object and Y stands for its class. X is assumed to take values in a feature space \mathfrak{X} (often $\mathfrak{X} = \mathbb{R}^d$) while Y takes either the value 0 (or -1) or the value 1 (for the positive class).

Standard formulation of the binary classification problem: Predict the value of Y from X or make an informed decision on the occurrence or non-occurrence of the event $Y = 1$ despite only being able to observe the values of X .

This is captured by the measure-theoretic setting of Assumption 1: Assume that X and Y map Ω into \mathfrak{X} and $\{0, 1\}$ respectively. Choose $\mathcal{H} = \sigma(X)$ (the smallest sub- σ -algebra of \mathcal{A} such that X is measurable) and $A_1 = \{Y = 1\} = \{\omega \in \Omega : Y(\omega) = 1\}$.

In many machine learning papers, the image (or pushforward) measure of P (or Q if it refers to the target distribution) under the mapping (X, Y) (see Definition 1.98 of Klenke [17]) is denoted by $p(x, y)$.

Often no probability space is specified but only samples $(x_1, y_1), \dots, (x_n, y_n)$ of realisations of (X, Y) from the source distribution and x_{n+1}, \dots, x_{n+m} of realisations of X from the target distribution are assumed to be given. This context is sometimes called ‘unsupervised domain adaptation’. Usually the samples are assumed to have been generated through i.i.d. drawings from some population distributions which may be identified with (Ω, \mathcal{A}, P) and (Ω, \mathcal{A}, Q) as described above.

2.3 More on Dataset Shift

Arguably, the two most important special cases of dataset shift are the following, in the terms introduced in Assumption 1:

- *Covariate shift* (Moreno-Torres et al. [21], Definition 3; Storkey [25], Section 5):
In this case, $P|\mathcal{H} \neq Q|\mathcal{H}$ holds but $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$, i.e. the posterior probabilities under P and Q are the same but the covariate distributions may be different.
- *Prior probability shift* (Moreno-Torres et al. [21], Definition 2; Storkey [25], Section 6):
In this case, we have $P[A_1] \neq Q[A_1]$ but $P[H | A_i] = Q[H | A_i]$, $i \in \{0, 1\}$, for all $H \in \mathcal{H}$, i.e. the class-conditional covariate source and target distributions are the same but the unconditional class prior probabilities may be different.

Covariate shift and prior probability shift are similar in the sense that in both cases one of the conditional distributions (of A_1 conditional on \mathcal{H} and of \mathcal{H} conditional on $\sigma(\{A_1\})$ respectively) are invariant between P and Q , and at least one pair of the marginal distributions (of \mathcal{H} and $\sigma(\{A_1\})$ respectively) are different.

Thanks to the invariance assumptions on the conditional distributions in prior probability shift and in covariate shift, these two types of dataset shift are relatively easily amenable to mathematical treatment and, therefore, have

received considerable attention by researchers. See e.g. Quiñonero-Candela et al. [22] for covariate shift and Caelen [5] for prior probability shift, as well as the references therein.

Note that the definition of dataset shift in Section 2.1 explicitly mentions an associated set of covariates (features, represented through the sub- σ -algebra \mathcal{H}). In Section 3, we are going to look closer at the question whether or not the covariate shift property is preserved in the relationship between source and target distribution if the amount of information reflected by the set of covariates is reduced. Formally, the question is phrased as follows:

Under Assumption 1, if $\mathcal{G} \subset \mathcal{H}$ is another sub- σ -algebra of \mathcal{A} , does then $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$ imply $P[A_1 | \mathcal{G}] = Q[A_1 | \mathcal{G}]$?

3 Covariate Shift is Fragile

In theory, class prior estimation under covariate shift is straightforward. Assume that the source distribution P and the target distribution Q are related through covariate shift as defined in Section 2.3. Then by the law of total probability and the fact that $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$, the prior class probability $Q[A_1]$ of the positive class can be represented as

$$Q[A_1] = E_Q[P[A_1 | \mathcal{H}]]. \quad (2)$$

As mentioned in Section 2.1, both $P[A_1 | \mathcal{H}]$ and $Q[A_1 | \mathcal{H}]$ typically are observable at the time when $Q[A_1]$ is to be estimated such that $Q[A_1]$ in principle can be calculated by means of (2). In the literature on class prior estimation, the approach based on (2) is known as ‘probability estimation & average (P & A)’ (Bella et al. [2]) or ‘probabilistic classify & count (PCC)’ (González et al. [11]).

Unfortunately, (2) may not work well in practice:

- Card and Smith [6] observed that poor calibration of the estimates of the posterior class probabilities would entail poor results for the PCC prior probability estimates.
- At a more fundamental level, Storkey ([25], Section 5.1) pointed out that the probability masses of the covariates might be quite differently located under the source and target distributions. As a consequence, an estimate of $P[A_1 | \mathcal{H}]$ made under the source distribution P might turn out to be rather biased in those regions of the covariate space to which the target distribution Q attributes most mass. This problem can be mitigated by ‘importance weighting’ which, however, may significantly complicate the estimation procedure.

Due to these issues, it is tempting to try to avoid the potentially difficult estimation of the posterior class probability $P[A_1 | \mathcal{H}]$ which is conditioned on the full covariate information set \mathcal{H} , by mimicking the simplification achieved through the confusion matrix method (Saerens et al. [23]; also called ‘adjusted count’ in Forman [8]) under prior probability shift.

Adapting the confusion matrix method to covariate shift would work as follows: Fix some hard (i.e. taking either the value 0 or the value 1) classifier which is a function of the covariates and therefore \mathcal{H} -measurable. We can identify the classifier with an event $H \in \mathcal{H}$ which specifies the range of the covariates on which a positive class label is predicted. If the source distribution P and the target distribution Q are related by covariate shift for the simple information set $\mathcal{G} = \{\emptyset, H, \Omega \setminus H, \Omega\} \subset \mathcal{H}$ then the following special case of (2) applies:

$$Q[A_1] \stackrel{?}{=} Q[H] P[A_1 | H] + (1 - Q[H]) P[A_1 | (\Omega \setminus H)]. \quad (3)$$

Eq. (3) appears to suggest a simple and efficient approach to class prior estimation under covariate shift which avoids the potentially difficult problem to estimate $P[A_1 | \mathcal{H}]$ for the more complex information set \mathcal{H} .

But can we always find a classifier (observable event) H such that the following condition for covariate shift with respect to $\mathcal{G} = \{\emptyset, H, \Omega \setminus H, \Omega\}$ and, as a consequence, also (3) hold true?

$$Q[A_1 | H] = P[A_1 | H] \quad \text{and} \quad Q[A_1 | (\Omega \setminus H)] = P[A_1 | (\Omega \setminus H)]. \quad (4)$$

The question mark in (3) is meant to suggest that the answer is ‘no’. This is illustrated with the following example.

Example 1. We revisit the binormal model with equal variances as an example that fits into the setting of Assumption 1. The source distribution P is defined by specifying the marginal distribution of $Y = \begin{cases} 1, & \text{on } A_1 \\ 0, & \text{on } A_0 \end{cases}$, with $P[A_1] = p \in (0, 1)$, and defining the class-conditional distributions of the covariate X given Y as normal distributions with equal variances:

$$P[X \in \cdot | A_1] = \mathcal{N}(\nu, \sigma^2) \quad \text{and} \quad P[X \in \cdot | A_0] = \mathcal{N}(\mu, \sigma^2). \quad (5)$$

In (5), we assume that $\mu < \nu$ and $\sigma > 0$. The unconditional distribution of X then is a mixture with weight p of the two normal distributions.

The posterior class probability $P[A_1 | \mathcal{H}]$ for $\mathcal{H} = \sigma(X)$ in this setting is given by $P[A_1 | \mathcal{H}] = (1 + \exp(aX + b))^{-1}$, with $a = \frac{\mu - \nu}{\sigma^2} < 0$ and $b = \frac{\nu^2 - \mu^2}{2\sigma^2} + \log\left(\frac{1-p}{p}\right)$. For the target distribution Q , we only specify the marginal distribution of the covariate X as another normal distribution with mean $E_Q[X] = \tau$ and variance $\text{var}_Q[X] = \sigma^2 + p(1-p)(\mu - \nu)^2$ such that the variance of X under Q matches the variance of X under P .

Under covariate shift, then by (2) it holds for the target prior class probability $Q[A_1]$ that

$$Q[A_1] = E_Q \left[(1 + \exp(aX + b))^{-1} \right]. \quad (6)$$

To illustrate the effect of simplification as suggested by (3), we define a family of classifiers $H_x = \{X > x\}$ for thresholds $x \in \mathbb{R}$.




Fig. 1. Illustration of Example 1, with parameters $\mu = 0$, $\nu = 1.5$, $\sigma = 1$, $p = 0.3$ and $\tau = 2.5$ for the normal distributions and the mixture parameter p of the source distribution. Simplification of the covariate information set causes the covariate shift property to be lost as demonstrated by the failure to hit the true target prior $Q[A_1]$ with the simplified estimates according to (3).

Figure 1 shows the true target prior class probability $Q[A_1]$ according to (6) (constant, dashed line) and, for moving threshold x , ‘pseudo’ priors according to (3) (solid curve). As the pseudo priors do not match the true prior, the covariate shift property (4) must be violated for all information sets

$$\mathcal{G}_x = \{\emptyset, H_x, \Omega \setminus H_x, \Omega\} \subset \mathcal{H} = \sigma(X).$$

This is due to the loss of information compared to the full information set \mathcal{H} associated with the covariate X . \square

On the basis of Example 1, we can conclude that under Assumption 1, if $\mathcal{G} \subset \mathcal{H}$ is another sub- σ -algebra of \mathcal{A} , then $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$ does not always imply $P[A_1 | \mathcal{G}] = Q[A_1 | \mathcal{G}]$, i.e. the covariate property may get lost if the amount of information represented by the covariates is reduced.

Information loss and subsequent loss of the covariate shift property can also be the consequence of deploying ‘domain-invariant representations’ (Johansson et al. [16], Section 4.1).

4 Covariate Shift and Statistical Sufficiency

In the following we will identify sufficient and necessary conditions for simplifications of covariate shift like (4) to hold. We will see that indeed it is almost impossible for (4) to be true if the information set \mathcal{H} of Assumption 1 is large compared to the information set \mathcal{G} on which (4) is based.

Definition 1. Under Assumption 1, denote by $C_A(P, \mathcal{H})$ the set of all probability measures Q on (Ω, \mathcal{A}) such that P and Q are related by covariate shift, i.e.

$$C_A(P, \mathcal{H}) = \{Q \text{ probability measure on } (\Omega, \mathcal{A}) : P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]\}.$$

Denote by $C_A^*(P, \mathcal{H})$ the set of all probability measures Q on (Ω, \mathcal{A}) such that P and Q are related by covariate shift and Q is absolutely continuous³ with respect to P on \mathcal{H} , i.e.

$$C_A^*(P, \mathcal{H}) = \{Q \in C_A(P, \mathcal{H}) : Q|_{\mathcal{H}} \ll P|_{\mathcal{H}}\}.$$

The following example illustrates the definition of $C_A(P, \mathcal{H})$ in some simple special cases.

Example 2. Consider the following three special cases for \mathcal{H} :

- (i) If $A_1 \in \mathcal{H}$ then we have that $\mathcal{H} \supset \sigma(\{A\})$ and

$$C_A(P, \mathcal{H}) = \{\text{All probability measures on } (\Omega, \mathcal{A})\},$$

because in this case it holds that $P[A_1 | \mathcal{H}] = \mathbf{1}_{A_1} = Q[A_1 | \mathcal{H}]$.

- (ii) If A_1 and \mathcal{H} are independent under P and Q , it follows that

$$P[A_1 | \mathcal{H}] = P[A_1] \quad \text{and} \quad Q[A_1 | \mathcal{H}] = Q[A_1].$$

Hence we have $Q \in C_A(P, \mathcal{H})$ if and only if $P[A_1] = Q[A_1]$.

- (iii) If $\mathcal{H} = \{\emptyset, \Omega\}$ we are in a special case of (ii). This implies

$$C_A(P, \mathcal{H}) = \{Q \text{ probability measure on } (\Omega, \mathcal{A}) : P[A_1] = Q[A_1]\}. \quad \square$$

Note that in case (i) of Example 2, $P[A_1] \neq Q[A_1]$ is possible.

Remark 1. The set $C_A(P, \mathcal{H})$ contains all probability measures Q with the property that there is an event $\Omega_Q \in \mathcal{H}$ such that $Q[\Omega_Q] = 1$ and $P[\Omega_Q] = 0$, i.e. Q and P are mutually singular. Although in this case there is an \mathcal{H} -measurable random variable Ψ that is both a version of $P[A_1 | \mathcal{H}]$ and of $Q[A_1 | \mathcal{H}]$, it is impossible to completely learn Ψ from the source distribution P because no instances $\omega \in \Omega_Q$ can be sampled under P due to $P[\Omega_Q] = 0$. Hence the distributions Q which are singular to P are not of great theoretical interest. $C_A(P, \mathcal{H})$ may also contain probability measures Q with both absolutely continuous and singular components (with respect to P). In this case, it is not possible to completely learn Ψ from P either. Therefore, in the following the focus is on the distributions Q that are absolutely continuous with respect to P . \square

³ Q is absolutely continuous with respect to P on \mathcal{H} (expressed symbolically as $Q|_{\mathcal{H}} \ll P|_{\mathcal{H}}$) if $P[N] = 0$ for $N \in \mathcal{H}$ implies $Q[N] = 0$. By the Radon-Nikodym theorem (Klenke [17], Corollary 7.34) then there exists an \mathcal{H} -measurable non-negative function h such that $Q[H] = E_P[h \mathbf{1}_H]$ for all $H \in \mathcal{H}$. The function h is called density of Q with respect to P .

At first glance, one might guess that the tower property of conditional expectations (Klenke [17], Theorem 8.14) implies $C_A(P, \mathcal{H}) \subset C_A(P, \mathcal{G})$ if \mathcal{G} is a sub- σ -algebra of \mathcal{H} . However, the following example shows that this is not true in general.

Example 3. Assume that $\mathcal{H} = \sigma(\mathcal{F} \cup \mathcal{G})$ for sub- σ -algebras \mathcal{F}, \mathcal{G} of \mathcal{A} , with $A_1 \notin \mathcal{H}$. Assume further that \mathcal{G} and $\sigma(\{A_1\} \cup \mathcal{F})$ are independent under P and Q . Then it follows that $P[A_1 | \mathcal{H}] = P[A_1 | \mathcal{F}]$ and $Q[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{F}]$.

Hence we have $Q \in C_A(P, \mathcal{H})$ if and only if $Q \in C_A(P, \mathcal{F})$. By case (ii) of Example 2, we have $Q \in C_A(P, \mathcal{G})$ if and only if $P[A_1] = Q[A_1]$. Hence, if there is a $Q \in C_A(P, \mathcal{F})$ with $P[A_1] \neq Q[A_1]$, we have an example showing that $C_A(P, \mathcal{H}) \not\subset C_A(P, \mathcal{G})$ may happen despite $\mathcal{G} \subset \mathcal{H}$. \square

Example 3 demonstrates that the covariate shift property may get lost if components of the covariates are dropped. We continue with presenting sufficient criteria for covariate shift (Lemma 1) and inheritance of covariate shift (Proposition 1 below).

Lemma 1. *Under Assumption 1, assume further that Q is absolutely continuous with respect to P on \mathcal{H}_A and that there is an \mathcal{H} -measurable density h of $Q|_{\mathcal{H}_A}$ with respect to $P|_{\mathcal{H}_A}$. Then it follows that $Q \in C_A^*(P, \mathcal{H})$.*

Proof. Fix any $H \in \mathcal{H}$. Then we obtain that

$$\begin{aligned} E_Q[\mathbf{1}_H P[A_1 | \mathcal{H}]] &= E_P[h \mathbf{1}_H P[A_1 | \mathcal{H}]] \\ &= E_P[E_P[h \mathbf{1}_{A_1 \cap H} | \mathcal{H}]] = E_P[h \mathbf{1}_{A_1 \cap H}] = Q[A_1 \cap H]. \end{aligned}$$

This implies $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$. \square

We are now going to point out connections between the notion of covariate shift and the following two concepts that have been considered in the literature in other contexts:

- Covariate shift with posterior drift (Scott [24]): Under Assumption 1, dataset shift between P and Q is more specifically called *covariate shift with posterior drift* if there is an increasing function $f : [0, 1] \rightarrow \mathbb{R}$ such that it holds that

$$Q[A_1 | \mathcal{H}] = f(P[A_1 | \mathcal{H}]). \quad (7)$$

- Sufficiency [7, 1, 26]: Under Assumption 1, if $\mathcal{G} \subset \mathcal{H}$ is another sub- σ -algebra of \mathcal{A} then \mathcal{G} is called (statistically) *sufficient* for \mathcal{H} with respect to A_1 if $P[A_1 | \mathcal{G}] = P[A_1 | \mathcal{H}]$ holds true.

Proposition 1. *Under Assumption 1, let P and Q be related by covariate shift with posterior drift such that (7) holds for an increasing function f . Assume that \mathcal{G} is sufficient for \mathcal{H} with respect to A_1 under the source distribution P and that Q is absolutely continuous with respect to P on \mathcal{H} . Then $Q[A_1 | \mathcal{G}] = f(P[A_1 | \mathcal{G}])$ follows.*

Proof. Let $h \geq 0$ be a density of Q with respect to P on \mathcal{H} . Then, in particular, h is \mathcal{H} -measurable. For any $G \in \mathcal{G}$, we therefore obtain

$$\begin{aligned} E_Q[\mathbf{1}_G f(P[A_1 | \mathcal{G}])] &= E_P[h \mathbf{1}_G f(P[A_1 | \mathcal{G}])] = E_P[h \mathbf{1}_G f(P[A_1 | \mathcal{H}])] \\ &= E_Q[\mathbf{1}_G f(P[A_1 | \mathcal{H}])] = E_Q[\mathbf{1}_G Q[A_1 | \mathcal{H}]] = Q[A_1 \cap G]. \end{aligned}$$

This implies the assertion. \square

Based on Lemma 1 and Proposition 1, we are in a position to prove the main result of this paper. It states that an information subset inherits the covariate shift property from its information superset for all absolutely continuous target distributions if and only if the subset is statistically sufficient for the superset with respect to the positive class label under the source distribution.

Theorem 1. *Under Assumption 1, let $\mathcal{G} \subset \mathcal{H}$ be another sub- σ -algebra of \mathcal{A} . Then \mathcal{G} is sufficient for \mathcal{H} with respect to A_1 under the source distribution P if and only if $C_A^*(P, \mathcal{H}) \subset C_A^*(P, \mathcal{G})$ holds true.*

Proof. The ‘only if’ part of the assertion is implied by Proposition 1. By the definition of conditional probability, for the ‘if’ part we have to show that for each $H \in \mathcal{H}$ it holds that $P[A_1 \cap H] = E_P[\mathbf{1}_H P[A_1 | \mathcal{G}]]$. This is obvious for H with $P[H] = 0$. Hence fix an event $H \in \mathcal{H}$ and assume $P[H] > 0$.

Define the probability measure Q_H on (Ω, \mathcal{A}) as P conditional on H , i.e.

$$Q_H[M] = P[M | H] = \frac{P[M \cap H]}{P[H]}, \quad \text{for all } M \in \mathcal{A}.$$

This Q_H is absolutely continuous with respect to P on $\mathcal{A} \supset \mathcal{H}_A$, with \mathcal{H} -measurable density $\frac{1_H}{P[H]}$. Hence, by Lemma 1 we obtain $Q_H \in C_A^*(P, \mathcal{H})$. By assumption, this implies $Q_H \in C_A^*(P, \mathcal{G})$, and in particular $P[A_1 | \mathcal{G}] = Q_H[A_1 | \mathcal{G}]$. From this, it follows that

$$\begin{aligned} E_P[\mathbf{1}_H P[A_1 | \mathcal{G}]] &= P[H] E_{Q_H}[P[A_1 | \mathcal{G}]] \\ &= P[H] E_{Q_H}[Q_H[A_1 | \mathcal{G}]] = P[H] Q_H[A_1] = P[A_1 \cap H]. \end{aligned}$$

This completes the proof. \square

5 Discussion of Theorem 1

Can sufficiency of \mathcal{G} for \mathcal{H} with respect to A_1 be characterised in other ways than just requiring $P[A_1 | \mathcal{G}] = P[A_1 | \mathcal{H}]$?

- As observed by Devroye et al. [7] (Section 32), if $\mathcal{G} = \sigma(T)$ is generated by some random variable T , then \mathcal{G} is sufficient for \mathcal{H} if and only if there exists a measurable function g such that $P[A_1 | \mathcal{H}] = g(T)$.

- Primary examples for such T are transformations $T = f(P[A_1 | \mathcal{H}])$ of the posterior class probability which may emerge as scoring classifiers optimising the area under the Receiver Operating Characteristic (ROC) or the area under the Brier curve (Tasche [26], Section 5.3). The process to reengineer $P[A_1 | \mathcal{H}]$ from T is called ‘calibration’ (see Kull et al. [18] and the references therein).

Johansson et al. [16] wrote in Section 4.1: “One interpretation . . . is that covariate shift ([their] Assumption 1) need not hold with respect to the representation $Z = \phi(X)$, even if it does with respect to X . With $\phi^{-1}(z) = \{x : \phi(x) = z\}$,

$$p_t(Y | z) = \frac{\int_{x \in \phi^{-1}(z)} p_t(Y | x) p_t(x) dx}{\int_{x \in \phi^{-1}(z)} p_t(x) dx} \neq p_s(Y | z). \quad (8)$$

Equality holds for general p_s, p_t only if ϕ is invertible.” According to Section 2 of Johansson et al., p_s and p_t stand for the densities of the covariate X on the ‘source domain’ and ‘target domain’ respectively. By Theorem 1, with $\mathcal{G} = \sigma(Z)$, actually covariate shift holds under the transformation ϕ if \mathcal{G} is sufficient for $\mathcal{H} = \sigma(X)$ (in the setting of Johansson et al.). Sufficiency of \mathcal{G} is implied by invertibility of ϕ . Hence, Theorem 1 is a more general statement than the one by Johansson et al. [16].⁴

Under Assumption 1, a mapping (representation) $T : (\Omega, \mathcal{H}) \rightarrow (\Omega_T, \mathcal{H}_T)$ which is \mathcal{H}_T - \mathcal{H} -measurable is said to have ‘invariant components’ (Gong et al. [10]) if its distributions under the source and target distributions are the same, i.e. if

$$P[T \in M] = Q[T \in M], \quad \text{for all } M \in \mathcal{H}_T. \quad (9)$$

As \mathcal{H} reflects the covariates, T can be interpreted as a transformation of the covariates that makes their distributions undistinguishable under the source and target distributions. As Gong et al. [10] noted, (9) alone does not imply that the posterior probabilities under source and target distributions are the same or at least similar. He et al. [12] therefore defined the notion of ‘domain invariance’ by

$$P[A_1 | \sigma(T)] = P[A_1 | \mathcal{H}], \quad Q[A_1 | \sigma(T)] = Q[A_1 | \mathcal{H}], \quad \text{and} \quad (10a)$$

$$P[A_i \cap \{T \in M\}] = Q[A_i \cap \{T \in M\}], \quad i \in \{0, 1\}, M \in \mathcal{H}_T. \quad (10b)$$

He et al. [12] then observed that (10a) and (10b) together imply covariate shift with respect to the information set \mathcal{H} , i.e. $P[A_1 | \mathcal{H}] = Q[A_1 | \mathcal{H}]$.⁵ In a sense, the observation by He et al. can be considered complementary to Theorem 1 because

⁴ The derivation of (8) in [16] is somewhat sloppy. In Section 2.3 of [16], the assumption is made for $Z = \phi(X)$ that ‘ $p(Z)$ ’ is a density. This implies $\int_{x \in \phi^{-1}(z)} p_t(x) dx = 0$ which means that the denominator of the fraction in (8) is zero.

⁵ Actually, (10b) implies covariate shift with respect to $\sigma(T)$. From this, together with (10a), follows covariate shift with respect to \mathcal{H} . Hence the assumption of (10b) could be replaced by the weaker assumption of having covariate shift with respect to $\sigma(T)$.

Theorem 1 is about passing on covariate shift from a larger information set to a smaller one while the observation by He et al. is a statement about covariate shift on a smaller information set implying covariate shift on a larger one.

In unsupervised domain adaptation, the case of source and target distributions where part or all of the support of the target distribution is not covered by the support of the source distribution is of great interest [3,16]. In that case, the target distribution is at least partially singular to the source distribution. Has Theorem 1 any relevance for this situation? Arguably, representations of the covariates which do not work even in the plain-vanilla environment of target distributions which are absolutely continuous with respect to the source distribution, are rather questionable. Hence Theorem 1 may be considered useful for providing a kind of ‘fatal flaw’ test for representations.

There are situations when covariate shift for a given sub- σ -algebra \mathcal{G} can be forced. The most important example of such a situation is sample selection (Hein [13], ‘Class-Conditional Independent Selection’). Theorem 1 may not be relevant then.

However, if the rationale for the assumption of covariate shift is based on causality considerations (like e.g. in Storkey [25]), the set of covariates associated to the information set \mathcal{H} in the definition of covariate shift might turn out to be quite large, rendering tedious the task of estimating the posterior $P[A_1 | \mathcal{H}]$. Theorem 1 provides the condition under which the size (or dimension) of the set of covariates may be reduced without destroying the invariance of the posterior class probabilities between the source and arbitrary target distributions. This condition does not require any special properties of the target distributions Q but the harmless requirement of being absolutely continuous with respect to the source distribution P . Note however that Theorem 1 leaves open the possibility that the covariate shift property is inherited by a non-sufficient sub- σ -algebra for some (but not all) specific target distributions.

If the set of covariates generating \mathcal{H} contains at least one real-valued covariate which has a Lebesgue-density and is not independent of A_1 , then there is no sufficient four-elements sub- σ -algebra \mathcal{G} such that (4) holds. For sufficiency would imply that the range of the posterior class probability $P[A_1 | \mathcal{H}]$ consists of two values only – which is wrong for probabilities conditional on continuous random variables. Hence by Theorem 1 no radically simple approach to class prior estimation like (3) that would be applicable under all possible shifts of the covariate distribution is available in this case.

6 Probing for class prior estimation under covariate shift

To the author’s best knowledge, there is basically one approach to class prior estimation on the target dataset under covariate shift: Estimate the posterior probability of the positive class as a function of the covariates on the source dataset and then calculate its average on the target dataset, see (2). Card and Smith [6] discuss two variants of this approach, one of them with and the other

without proper calibration of the posterior probabilities – hence the concept in principle is the same in both variants.

Under prior probability shift, the simple ‘confusion matrix method’ can be deployed to achieve consistent class prior estimates [9,23]. As seen in Sections 3 and 4, no similarly simple approach based on merely making use of one classifier’s output works under covariate shift. However, averaging the counting results of a large ensemble of classifiers trained for a variety of cost-sensitive classification problems would work (‘probing’: Langford and Zadrozny [19]; Tasche [26]).

Sketch of class prior estimation with probing. Define the cost-sensitive (weighted) classification loss (with $0 \leq t \leq 1$) in the setting of Assumption 1:

$$L(H, t) = (1 - t)P[A_1 \cap (\Omega \setminus H)] + tP[A_0 \cap H], \quad H \in \mathcal{H}.$$

The probing algorithm adapted to class prior estimation then can be described as follows:

- 1) Choose an appropriately ‘dense’ set $0 = t_0 < t_1 < t_2 < \dots < t_n < 1$.
- 2) For each t_i , $i = 1, \dots, n$, find – with possibly different approaches – a nearly optimal minimising classifier⁶ $H(t_i)$ of $L(H, t_i)$, $H \in \mathcal{H}$.
- 3) Let $Z = \sum_{i=1}^n (t_i - t_{i-1}) \mathbf{1}_{H(t_i)}$.
- 4) For all j with $L(\{Z > t_j\}, t_j) < L(H(t_j), t_j)$, replace $H(t_j)$ with $\{Z > t_j\}$.
- 5) Repeat steps 3) and 4) until $L(\{Z > t_j\}, t_j) \geq L(H(t_j), t_j)$ for all j .
- 6) Calculate $\hat{q} = \sum_{i=1}^n (t_i - t_{i-1}) Q[H(t_i)]$ as estimate of the positive class prior probability $Q[A_1]$ under the target distribution.

7 Conclusions

We have shown that covariate shift is a fragile notion, in the sense that the invariance of the posterior class probabilities between source and target distributions may be lost if the set of covariates on which the posterior probabilities are conditioned is diminished. This observation implies that under covariate shift simple estimators of the target prior class probabilities are infeasible if they are designed in the style of the confusion matrix method (adjusted count) which is a popular quantifier under prior probability shift.

Valid methods for class prior estimation under covariate shift are the careful estimation of the posterior class probabilities conditioned on the full set or a sufficient subset of the covariates, combined with subsequently averaging them on the target dataset (probabilistic classify & count). The application of probing as described in Section 6 could also prove useful for class prior estimation under covariate shift. So far, probing for class prior estimation has not yet been thoroughly tested. This could be a subject for future research.

⁶ As before, we identify a set with its indicator function that gives the value 1 on the set and the value 0 on its complement.

Acknowledgements. The author is grateful to Juan José del Coz and Pablo González for drawing his attention to the subject of class prior estimation under covariate shift and to four anonymous reviewers whose comments redounded to significant improvements of the paper.

References

1. Adraghi, K., Cook, R.: Sufficient dimension reduction and prediction in regression. *Philosophical Transactions of the Royal Society A* **367**, 4385–4405 (2009)
2. Bella, A., Ferri, C., Hernandez-Orallo, J., Ramírez-Quintana, M.: Quantification via probability estimators. In: *Data Mining (ICDM), 2010 IEEE 10th International Conference on*. pp. 737–742. IEEE (2010)
3. Ben-David, S., Blitzer, J., Crammer, K., Pereira, F.: Analysis of Representations for Domain Adaptation. In: Schölkopf, B., Platt, J., Hoffman, T. (eds.) *Advances in Neural Information Processing Systems*. vol. 19, pp. 137–144. MIT Press (2006)
4. Billingsley, P.: *Probability and measure*. John Wiley & Sons, second edn. (1986)
5. Caelen, O.: *Quantification and learning algorithms to manage prior probability shift*. Master thesis, Institut de Statistique, Biostatistique et Sciences Actuarielles, Université catholique de Louvain (2017)
6. Card, D., Smith, N.: The Importance of Calibration for Estimating Proportions from Annotations. In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 1636–1646 (2018)
7. Devroye, L., Györfi, L., Lugosi, G.: *A Probabilistic Theory of Pattern Recognition*. Springer (1996)
8. Forman, G.: Counting Positives Accurately Despite Inaccurate Classification. In: *European Conference on Machine Learning (ECML 2005)*. pp. 564–575. Springer (2005)
9. Gart, J., Buck, A.: Comparison of a screening test and a reference test in epidemiologic studies. II. A probabilistic model for the comparison of diagnostic tests. *American Journal of Epidemiology* **83**(3), 593–602 (1966)
10. Gong, M., Zhang, K., Liu, T., Tao, D., Glymour, C., Schölkopf, B.: Domain Adaptation with Conditional Transferable Components. In: Balcan, M., Weinberger, K. (eds.) *Proceedings of The 33rd International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 48, pp. 2839–2848. PMLR, New York, New York, USA (20–22 Jun 2016)
11. González, P., Castaño, A., Chawla, N., Coz, J.D.: A Review on Quantification Learning. *ACM Comput. Surv.* **50**(5), 74:1–74:40 (2017)
12. He, H., Yang, Y., Wang, H.: Domain Adaptation with Factorizable Joint Shift. *arXiv preprint arXiv:2203.02902* (2022)
13. Hein, M.: Binary Classification under Sample Selection Bias. In: Quiñero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N. (eds.) *Dataset Shift in Machine Learning*, chap. 3, pp. 41–64. The MIT Press, Cambridge, Massachusetts (2009)
14. Hofer, V., Kreml, G.: Drift mining in data: A framework for addressing drift in classification. *Computational Statistics & Data Analysis* **57**(1), 377–391 (2013)
15. Holzmann, H., Eulert, M.: The role of the information set for forecasting – with applications to risk management. *The Annals of Applied Statistics* **8**(1), 595–621 (2014)

16. Johansson, F., Sontag, D., Ranganath, R.: Support and Invertibility in Domain-Invariant Representations. In: Chaudhuri, K., Sugiyama, M. (eds.) Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 89, pp. 527–536. PMLR (16–18 Apr 2019)
17. Klenke, A.: Probability Theory: A Comprehensive Course. Springer Science & Business Media (2013)
18. Kull, M., Silva Filho, T., Flach, P.: Beyond sigmoids: How to obtain well-calibrated probabilities from binary classifiers with beta calibration. *Electron. J. Statist.* **11**(2), 5052–5080 (2017)
19. Langford, J., Zadrozny, B.: Estimating Class Membership Probabilities using Classifier Learners. In: Cowell, R., Ghahramani, Z. (eds.) AISTATS 2005 – Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics. pp. 198–205. The Society for Artificial Intelligence and Statistics (2005)
20. Lipton, Z., Wang, Y.X., Smola, A.: Detecting and Correcting for Label Shift with Black Box Predictors. In: Dy, J., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 80, pp. 3122–3130. PMLR (10–15 Jul 2018)
21. Moreno-Torres, J., Raeder, T., Alaiz-Rodriguez, R., Chawla, N., Herrera, F.: A unifying view on dataset shift in classification. *Pattern Recognition* **45**(1), 521–530 (2012)
22. Quiñero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N. (eds.): Dataset Shift in Machine Learning. MIT Press (2008)
23. Saerens, M., Latinne, P., Decaestecker, C.: Adjusting the Outputs of a Classifier to New a Priori Probabilities: A Simple Procedure. *Neural Computation* **14**(1), 21–41 (2001)
24. Scott, C.: A Generalized Neyman-Pearson Criterion for Optimal Domain Adaptation. In: Proceedings of Machine Learning Research, 30th International Conference on Algorithmic Learning Theory. vol. 98, pp. 1–24 (2019)
25. Storkey, A.: When Training and Test Sets Are Different: Characterizing Learning Transfer. In: Quiñero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N. (eds.) Dataset Shift in Machine Learning, chap. 1, pp. 3–28. The MIT Press, Cambridge, Massachusetts (2009)
26. Tasche, D.: Calibrating sufficiently. *Statistics* **55**(6), 1356–1386 (2021)
27. Tasche, D.: Minimising quantifier variance under prior probability shift. In: Cong, G., Ramanath, M. (eds.) Proceedings of the CIKM 2021 Workshops (2021), first International Workshop on Learning to Quantify: Methods and Applications (LQ 2021)