
Correcting for Selection Bias and Missing Response in Regression using Privileged Information

Philip Boeken^{1,2}

Noud de Kroon¹

Mathijs de Jong²

Joris M. Mooij¹

Onno Zoeter²

¹Korteweg-de Vries Institute for Mathematics, University of Amsterdam, The Netherlands

²Booking.com, The Netherlands

Abstract

When estimating a regression model, we might have data where some labels are missing, or our data might be biased by a selection mechanism. When the response or selection mechanism is *ignorable* (i.e., independent of the response variable given the features) one can use off-the-shelf regression methods; in the *nonignorable* case one typically has to adjust for bias. We observe that *privileged information* (i.e. information that is only available during training) might render a nonignorable selection mechanism ignorable, and we refer to this scenario as *Privilegedly Missing at Random* (PMAR). We propose a novel imputation-based regression method, named *repeated regression*, that is suitable for PMAR. We also consider an importance weighted regression method, and a doubly robust combination of the two. The proposed methods are easy to implement with most popular out-of-the-box regression algorithms. We empirically assess the performance of the proposed methods with extensive simulated experiments and on a synthetically augmented real-world dataset. We conclude that repeated regression can appropriately correct for bias, and can have considerable advantage over weighted regression, especially when extrapolating to regions of the feature space where response is never observed.

1 INTRODUCTION

Regression is a primary technique in data science, machine learning and statistics. When presented with data $(X_1, Y_1), \dots, (X_n, Y_n)$ sampled from the distribution $\mathbb{P}(X, Y)$, the goal is to find the conditional expectation of Y given X , i.e. estimate the function $\mu(x) = \mathbb{E}[Y|X = x]$. Practitioners are often presented with either incomplete data

(where some values are missing) or data that is not representative of the population (drawn from some other $\tilde{\mathbb{P}}(X, Y)$ with $\tilde{\mathbb{P}} \neq \mathbb{P}$), and hence have to correct for the bias that is present in their training data. The discrepancy between $\tilde{\mathbb{P}}$ and \mathbb{P} can for example be induced by a selection mechanism.

Consider the situation where we want to predict whether a loan applicant will default. Let X be the digital record of a person applying for a loan, Y whether the borrower defaults on the loan, \hat{Y} a current algorithm's prediction of default, and Z (costly) expert advice on whether the applicant will default, which correlates with Y through information that is unavailable to us. We want to reject any applicants that will default on the loan, so we have the issue of the loan S , as a weighted combination of \hat{Y} and Z . Only when we issue the loan, Y will be observed. If we want to re-train our current prediction model $\hat{Y} = \hat{\mathbb{E}}[Y|X]$ we have to take into account the bias induced by Z during training of the model without explicitly adding Z to the covariates X . In this work, we demonstrate how one can incorporate such *privileged information* in a regression model, to correct for any bias that it induces in the data generating process.

An important line of work on selection bias is by Pearl [2012] and Bareinboim et al. [2014], who consider the problem of estimating $\mathbb{P}(Y|X)$ from a potentially biased dataset by leveraging knowledge of the underlying causal graph. They derive an expression of $\mathbb{P}(Y|X)$ as an integral of quantities that can be estimated from readily available biased and additional ('external') unbiased data. Although identification of a conditional distribution might be useful in certain scenarios, this does not tell the practitioner how to estimate a regression model, especially when dealing with continuous variables. In this work we address this problem, keeping in mind the applicability of the proposed methods.

Missingness problems are often characterised by the presence of certain conditional independencies in the data. Typically these independencies are untestable, making it unclear whether these conditional independence assumptions are appropriate for the data at hand. Drawing a causal graph of

the data generating process is often helpful for gaining a better understanding of the problem. Moreover, the causal graph identifies conditional independencies in the data, and thereby allows the practitioner to motivate these conditional independence assumptions. As posed by Rubin [1976]:

“The inescapable conclusion seems to be that when dealing with real data, the practising statistician should explicitly consider the process that causes missing data far more often than he does. However, to do so, he needs models for this process and these have not received much attention in the statistical literature.”

Contrary to what is often seen in the literature we do not need a full probabilistic model of the missingness mechanism, as the identification of certain conditional independencies can already be sufficient for estimating a regression model.

Our contributions are as follows. First, this paper serves a pedagogical purpose by reviewing literature of missingness and selection bias, as addressed in Section 2. We address what the consequences of certain conditional independencies are for the practitioner. While doing so, we motivate the use of *privileged information* [Vapnik and Vashist, 2009] when missingness or selection is nonignorable, and introduce the *Privilegedly Missing at Random* setting (PMAR). Secondly, in Section 3 we formulate a novel imputation-based estimator for the PMAR setting, which we refer to as *repeated regression*. We point out that importance weighted regression is also suitable for the PMAR setting, and combine the two methods into a doubly robust estimator. These estimators are formulated such that they are easy to implement using out-of-the box regression algorithms. We then address the intricacies of evaluation under selection bias and missingness in Section 4. We warn the practitioner that, to our knowledge, there is no appropriate way of evaluating a regression method on a finite, biased dataset without relying on auxiliary models. Lastly, in Section 5 we assess the performance of the formulated methods on extensive simulated experiments and on synthetically augmented real-world data. We observe that repeated regression has considerable advantage over importance weighting, especially when extrapolating.

2 MISSING DATA, SELECTION BIAS AND PRIVILEGED INFORMATION

2.1 MISSING DATA MECHANISMS

A very general framework for handling missing data is proposed by Rubin [1976]. When modelling the distribution of a set of random variables X_1, \dots, X_m , this framework takes into account potential missing values of any of the X_i by considering a random vector of *response indicators*: binary variables S_1, \dots, S_m , where $S_i = 1$ indicates that variable

X_i is observed. The process that determines whether we observe X_i can be explicitly modelled as $\mathbb{P}(S_i | X_1, \dots, X_m)$. We then define the *observed* random vector X_1^o, \dots, X_m^o where $X_i^o = X_i$ if $S_i = 1$ and $X_i = ?$ if $S_i = 0$, where ‘?’ denotes the value that is missing in the dataset. We focus on a specific missingness problem where we have covariates X , response variable Y , and where S is a response indicator for Y , so only Y can have missing values.

Rubin [1976] proposed the following classification of missing data mechanisms, which approximately categorises the difficulty of many inference problems.

Definition 1. [Rubin, 1976] Given variables X, Y , and response indicator S for Y , we say that Y is

- Missing Completely at Random (MCAR) if the missingness mechanism is independent of all other observed variables, i.e. $S \perp\!\!\!\perp X, Y$;
- Missing at Random (MAR) if the missingness mechanism is independent of the missing variable given all other fully observed variables, i.e. $Y \perp\!\!\!\perp S | X$;
- Missing Not at Random (MNAR) if it is neither MCAR nor MAR.

A large body of literature has been written about inference under missingness. Important works are the EM algorithm [Dempster et al., 1977], Nobel prize winning work by Heckman [1979] on correcting for selection bias in linear regression, Rosenbaum and Rubin [1984] which laid the foundations of propensity score based methods in causal inference, and a series of generalised estimating equations (GEE) based methods for inference under missingness [Robins et al., 1994, Robins and Rotnitzky, 1995, Rotnitzky et al., 1998, Scharfstein et al., 1999]. Graphical modelling of missingness mechanisms and related independence testing and identification problems have been investigated in Daniel et al. [2012], Thoemmes and Mohan [2015], Nabi et al. [2020], Mohan and Pearl [2021], Goel et al. [2021]. For an overview of the field, see Little and Rubin [2019]. In this work we focus on correcting for bias in regression using privileged data, which is not treated in the aforementioned literature.

S-recoverability As proposed by Pearl [2012] and further developed in Bareinboim et al. [2014], *s-recoverability* is a method to deal with selection bias. Consider again the case of modelling the distribution of X_1, \dots, X_m . Instead of having one response indicator for each X_i as in the missingness framework, Pearl [2012] considers one selection variable S where $S = 1$ indicates that all variables are observed (*selected*), and $S = 0$ indicates that no variable is observed, i.e. there is no row for this observation in our dataset. Any data that we observe is drawn from the distribution $\mathbb{P}(X_1, \dots, X_m | S = 1)$.

Bareinboim et al. [2014] prove that for discrete variables and under certain positivity assumptions, $\mathbb{P}(Y|X)$ is recoverable from $\mathbb{P}(X, Y|S = 1)$ if and only if $Y \perp\!\!\!\perp S | X$. The ‘if’ part is straightforward, since the conditional independence implies $\mathbb{P}(Y|X) = \mathbb{P}(Y|X, S = 1)$, and the right-hand-side can be estimated from the data. When this conditional independence is not satisfied, Bareinboim et al. [2014] consider joint measurement of X, Y and some other variable Z in the biased dataset (so availability of $\mathbb{P}(X, Y, Z|S = 1)$), satisfying the conditional independence $Y \perp\!\!\!\perp S | X, Z$. Additionally, they assume availability of unbiased measurements of (X, Z) , i.e. sampled from $\mathbb{P}(X, Z)$, which they refer to as ‘external data’. For discrete Z , the quantity $\mathbb{P}(Y|X)$ is then identified with

$$\begin{aligned} \mathbb{P}(Y|X) &= \sum_z \mathbb{P}(Y|X, Z = z)\mathbb{P}(Z = z|X) \\ &= \sum_z \mathbb{P}(Y|X, Z = z, S = 1)\mathbb{P}(Z = z|X) \end{aligned} \quad (1)$$

provided we have $\mathbb{P}(S = 1|X, Z) > 0$, $\mathbb{P}(X, Z)$ -almost surely, which is for example satisfied when $\text{supp}[\mathbb{P}(X, Z|S = 1)] = \text{supp}[\mathbb{P}(X, Z)]$.¹ One can straightforwardly replace the sum with an integral when Z is continuous. However, when the domain of Z is countably infinite or continuous, estimating this quantity is not straightforward. The repeated regression estimator proposed in Section 3.1 is a solution to this problem.

Our proposed methods can be applied to the selection bias and missingness settings. In either case, we require availability of i.i.d. observations of $(X, Y, Z) \sim \mathbb{P}(X, Y, Z|S = 1)$ whose index set we denote with \mathcal{S} , and i.i.d. observations of $(X, Z) \sim \mathbb{P}(X, Z)$ whose index set we denote with \mathcal{D} . In the missingness setting, both samples are readily available and we have $\mathcal{S} \subseteq \mathcal{D}$. In the selection bias setting the sample \mathcal{D} consists of ‘external’ data and we typically have $\mathcal{S} \cap \mathcal{D} = \emptyset$. A schematic display of the assumed available data under missingness and selection bias is provided in the supplements. For the methods of Sections 3.2 and 3.3 we additionally require knowledge of the selection probability $\mathbb{P}(S = 1|X, Z)$, which is directly estimable in the missingness setting, but which has to be assumed under selection bias. Throughout this paper, any distinctions between missingness and selection bias will be pointed out when necessary. Otherwise, either setting can be assumed.

2.2 REGRESSION UNDER DIFFERENT SELECTION MECHANISMS

Suppose we are interested in estimating $\mathbb{E}[Y|X]$ with continuous or discrete (ordinal) Y , and arbitrary X .² We might be confronted with a dataset with missing values of Y , or we might suspect that some selection mechanism is in play

¹We let $\text{supp}[\mathbb{P}]$ denote the support of \mathbb{P} .

²Note that this setting includes binary classification.




Figure 1: Examples of missingness or selection bias settings, where S indicates whether Y is observed or not.

which makes certain (X, Y) pairs unobserved. In this section, we investigate whether there is need for any bias correction. As this investigation is based on conditional independence assumptions, we first elaborate how such assumptions can be motivated.

Causal modelling Missingness mechanisms can be characterised by independencies. This is everything we need: all proposed methods will only require certain conditional independencies in the data, and no causal assumptions. However, the conditional independencies that are assumed are typically untestable. For example, the independence $Y \perp\!\!\!\perp S$ is not testable, as we have not observed Y for $S = 0$. To motivate such an independence assumption, one could model the data generating process with a graphical causal model, and infer from d-separations in the graph that there must be certain independencies in the data [Pearl, 2009]. In the missingness setting where only Y can be missing (with indicator S), we can draw a simplified graph that discards the Y^o variable, as in Figure 1a. In the graphical framework of s-recoverability, S is implicitly required to be a sink node (i.e. a node without children) [Bareinboim et al., 2014]. We do not adopt this convention.

Ignorable missingness/selection bias When the data is MCAR or MAR, we have the conditional independence $Y \perp\!\!\!\perp S | X$ and thus

$$\mathbb{E}[Y|X] = \mathbb{E}[Y|X, S = 1], \quad (2)$$

in which case we call the missingness/selection mechanism *ignorable* for estimating $\mathbb{E}[Y|X]$. An example of an ignorable (MAR) missingness or selection mechanism is given in Figure 1a. When this mechanism is ignorable, any correctly specified model could directly be learned as *Empirical Risk Minimization* (ERM) is consistent [Sugiyama et al., 2007], and hence, there is no need for bias correction. Note that the right-hand side of equation (2) is only defined for $X \in \text{supp}[\mathbb{P}(X|S = 1)]$, and one has to be mindful when extrapolating to values $X \in \text{supp}[\mathbb{P}(X)] \setminus \text{supp}[\mathbb{P}(X|S = 1)]$ [Martius and Lampert, 2016].

Although in this case the missingness or selection mechanism can in principle be ignored, one should beware that under model misspecification, ERM is not consistent any-

more. Also, efficiency of the estimation procedure can be affected by MCAR and MAR. Zadrozny [2004] and Wei Fan et al. [2005] investigate the performance of multiple popular classification algorithms under covariate shift. Robins and Rotnitzky [1995] investigate the asymptotic efficiency of semi-parametric estimation of the regression function via Generalised Estimating Equations (GEE). Note that in the missingness setting, we can use the values of X where Y is unobserved as additional (unlabelled) input of a *semi-supervised learning* algorithm. When doing *anticausal* prediction (i.e., when Y is the cause of X), these additional samples could benefit performance [Schölkopf et al., 2012].

Privilegedly ignorable missingness/selection bias It might be the case that the missingness or selection mechanism is nonignorable. For example, we may suspect that S and Y are confounded by some variable that is not contained in X , as depicted in Figure 1b. In this case we typically have $\mathbb{E}[Y|X, S = 1] \neq \mathbb{E}[Y|X]$, in which case naive regression on the available data yields a biased model. However, we might be able to identify the latent confounder, or more generally, any variable (e.g. mediator or confounder of Y and S) that induces a dependence between Y and S after conditioning on X .

More specifically, we are looking for a set of variables Z such that $Y \perp\!\!\!\perp S | X, Z$, such as depicted in Figure 1c. Having identified Z , one might include it into the features of the model and deploy the unbiased model $\hat{\mathbb{E}}[Y|X, Z, S = 1]$. However, in practice it might be costly or impossible to measure Z during deployment, e.g. when Z can only be measured after making a prediction with the regression model. In these settings, it is undesirable to incorporate Z into the features of the deployed model. When Z is only available during training it is referred to as *privileged information*, following Vapnik and Vashist [2009]. The specific setting that we consider in this work is formally defined as follows:

Definition 2 (PMAR). Given features X and label Y , response indicator S for Y , and privilegedly observed variable Z , we say that Y is *Privilegedly Missing at Random (PMAR)* if the response indicator is independent of the missing label given all other fully or privilegedly observed variables, i.e. $Y \perp\!\!\!\perp S | X, Z$.

When the data used for regression is PMAR, we refer to the missingness/selection mechanism as *privilegedly ignorable*.

Examples of PMAR A causal graph related to the bank loan problem from the introduction is given in Figure 2a, from which we deduce that indeed $Y \perp\!\!\!\perp S | X, Z$. The bank loan problem is an example of a more general PMAR setting where individuals or items are *selectively labelled* [Guerdan et al., 2023], based on predictions from an existing model $\hat{Y} = \hat{\mathbb{E}}[Y|X]$ and additional data Z that is not part of the features X .




Figure 2: Causal graphs relating to examples of PMAR.

The PMAR setting is also encountered when we want to replace the feature set Z of a current selection model $\hat{\mathbb{E}}[Y|Z]$ with a new feature set X . One could start measuring X when the old algorithm is still in use, to generate data on which the new prediction algorithm $\hat{\mathbb{E}}[Y|X]$ can be learned; this data generating process is depicted in Figure 2b.

Another example that fits the PMAR signature would be the selection of patients for costly CT scans, with X digitally available measurements of the patient, Z all additional information (besides X) that the doctor uses to decide whether a CT scan should be made of the patient (e.g. pale skin, slurred speech, pain indicated by patient), Y whether the patient has a certain disease (as measured by the CT scan), and S the doctors decision of making a CT scan of the patient. If we want to estimate $\mathbb{E}[Y|X]$, e.g. to assist the doctor in the future in their decision making process, then we can only unbiasedly estimate this if we have measured Z , which requires the doctor to manually register all information that they use for making the decision. This can be a costly process, hence we might only want to have to measure this at train time, and not at test time.

The previous examples are of missing response, where the distributions $\mathbb{P}(X, Y, Z | S = 1)$ and $\mathbb{P}(X, Z)$ are readily available. As an example of selection bias with the PMAR conditional independence, we consider a dataset of patients from the University Hospital of Caracas, Venezuela, who have been tested for the presence of cervical cancer through a biopsy or colposcopy (Y).³ In this dataset we have demographic and medical information of these patients (X), from which we might want to estimate $\mathbb{E}[Y|X]$ for automated screening of the population. However, it might be that these patients are self-selected based on any symptoms, denoted by Z . A possible causal graph of the data generating process is depicted in Figure 2c. Naive regression on this dataset would yield the biased model $\hat{\mathbb{E}}[Y|X, S = 1]$. As mentioned earlier, to correct for this bias we require external data \mathcal{D} from distribution $\mathbb{P}(X, Z)$, i.e. an unbiased sample of covariates X and symptoms Z from the population. This could for example be sampled through a questionnaire. One could compare this with the approach of sampling (X, Y) from

³Available at [https://archive.ics.uci.edu/ml/datasets/Cervical+cancer+\(Risk+Factors\)](https://archive.ics.uci.edu/ml/datasets/Cervical+cancer+(Risk+Factors)).

the population to directly estimate $\mathbb{E}[Y|X]$, which requires a costly and intrusive intervention (biopsy or colposcopy) on randomly sampled subjects, hence it would be preferred to only sample (X, Z) from the population.

In all these cases, the variable Z induces bias that should be corrected for when estimating $\mathbb{E}[Y|X]$. In the following section, we demonstrate how this can be achieved.

3 IMPUTATION, WEIGHTING, AND DOUBLE ROBUSTNESS UNDER PMAR

We propose three estimation procedures for training a regression model $\mathbb{E}[Y|X]$ in the PMAR setting: a repeated regression method, a weighted regression method, and a doubly robust regression method. Throughout this section, we consider the following example.

Example 1. Consider the data generating process

$$\begin{aligned} X &= \varepsilon_X \\ Z &= 3 \sin(X) + \varepsilon_Z \\ Y &= \frac{1}{2}X + Z + \varepsilon_Y \\ S &\sim \text{Bernoulli}(p_S(X, Z)) \end{aligned} \quad (3)$$

with independent Gaussian random variables $\varepsilon_X, \varepsilon_Z$ and ε_Y . The selection probability is defined as $p_S(x, z) := \sigma(x)\sigma(z)$ with sigmoid $\sigma(x) = 1/(1 + e^x)$.

Figure 3 displays the graphical model related to this data generating process. Note that we can indeed read off that $Y \perp\!\!\!\perp S | X, Z$, so if Z is only available during training, the missingness mechanism is PMAR. Figure 4 depicts $n = 400$ draws from the generating model (3). In this example, we have $\#\{S = 1\} = 123$. The black dots indicate (X, Y) pairs where Y is observed ($S = 1$), and the grey dots indicate (X, Y) pairs where Y is unobserved ($S = 0$). The green line shows the true regression line $\mathbb{E}[Y|X = x] = \frac{1}{2}x + 3 \sin(x)$, and the black line shows the regression line of the naive, biased estimate $\hat{\mathbb{E}}[Y|X, S = 1]$. In this example and throughout this paper we use thin plate splines regression [Duchon, 1977, Wood, 2003] as implemented in the `mgcv` package [Wood, 2015]. Note that, additional to what is shown in Figure 4, we have all (X, Z) pairs available. Our goal is to fit a regression model that is close to the green line, $\mathbb{E}[Y|X]$. \triangle




Figure 3




Figure 4: Dataset from Example 1, with observed (black) and unobserved data (grey); the true regression line $\mathbb{E}[Y|X]$ (green) and the naively estimated regression line $\hat{\mathbb{E}}[Y|X, S = 1]$ (black).

run into extrapolation issues already when $\mathbb{P}(S = 1|X, Z)$ is close to zero. In the example of Figure 4 we have no positivity for large values of X ; we will see that extrapolation is still possible by using privileged data Z that is highly predictive of Y .

3.1 REPEATED REGRESSION

By the law of total expectation and the PMAR conditional independence $Y \perp\!\!\!\perp S | X, Z$ we can write the conditional expectation $\mathbb{E}[Y|X]$ as follows

$$\begin{aligned} \mathbb{E}[Y|X] &= \mathbb{E}[\mathbb{E}[Y|X, Z]|X] \\ &= \mathbb{E}[\mathbb{E}[Y|X, Z, S = 1]|X], \end{aligned} \quad (4)$$

similar to equation (1). For this equality to hold we require $\mathbb{P}(S = 1|X, Z) > 0$, $\mathbb{P}(X, Z)$ -almost surely. We formulate an estimation procedure based on this expression by estimating each conditional expectation with a regression model. That is, we first regress Y on X and Z using the data S , which results in the regression model $\hat{\mathbb{E}}[Y|X, Z, S = 1]$. Using the unbiased data (x_i, z_i) for $i \in \mathcal{D}$ we construct *pseudo-labels* $\tilde{Y}_i := \hat{\mathbb{E}}[Y|X = x_i, Z_i = z_i, S = 1]$. Now, we regress \tilde{Y} on X using the data $(X_i, \tilde{Y}_i)_{i \in \mathcal{D}}$, which yields the final model:

$$\hat{\mu}_{RR}(x) = \hat{\mathbb{E}}[\tilde{Y}|X = x]. \quad (5)$$

Note that this method requires datasets S and \mathcal{D} , and assumes that $Y \perp\!\!\!\perp S | X, Z$. So, if this conditional independence is satisfied, it can be directly applied in both the missingness setting, and in the selection bias setting where we have ‘external data’ $\mathbb{P}(X, Z)$.

Extrapolation and positivity Many missingness methods assume *positivity* of the selection probability, i.e. that $\mathbb{P}(S = 1|X, Z) > 0$ holds $\mathbb{P}(X, Z)$ -almost surely, as is for example required for the identification equation (1). This assumption ensures that asymptotically, we do not run into the problem of extrapolation. For finite samples, we might

Example 1 (continued). Recall that the data for Y is generated by $Y = \frac{1}{2}X + Z + \varepsilon_Y$, so for fitting a model $\hat{\mathbb{E}}[Y|X, Z, S = 1]$, although we can only fit this model

on a small part of the data (123 out of 400 observations) it is a relatively easy (linear) model to fit. We compute the pseudo-labels $\tilde{Y}_i := \mathbb{E}[Y|X = x_i, Z = z_i, S = 1]$, as depicted with orange crosses in Figure 5. Regressing these imputed values on X yields the final model $\hat{\mu}_{RR}$ (orange line in Figure 5). \triangle




Figure 5: Repeated regression. Orange crosses indicate imputed values, and the orange line is the regression line resulting from repeated regression.

Regression imputation is explicitly used for estimating the mean under missingness via $\mathbb{E}[Y] = \mathbb{E}[\mathbb{E}[Y|X, S = 1]]$ by Bang and Robins [2005]. Similarly, the *average causal effect* can be estimated by $\mathbb{E}[Y \text{ do}(X)] = \mathbb{E}[\mathbb{E}[Y|X, Z]]$ for a valid *adjustment set* Z ; this method is known as *standardization* [Hernán and Robins, 2021]. For these estimators, the inner expectation is estimated with a regression model, and the outer expectation is calculated by taking the mean of the pseudo-labels. The novelty in estimator (5) lies in the fact that the outer expectation is estimated by regression instead of the empirical mean, and the observation that this estimator is applicable to the PMAR setting. Note that in general, automatically generated confidence intervals resulting from the outer regression $\hat{\mathbb{E}}[\tilde{Y}|X]$ are not valid. To this end, *multiple imputation* methods are often used. As this requires the modelling of the full distribution $\mathbb{P}(Y|X, Z)$, we do not consider this. For an overview of multiple imputation methods in GEE regression, see Dithlhong et al. [2018].

3.2 IMPORTANCE WEIGHTING

For estimating the parameter β^* in the regression model $\mathbb{E}[Y|X] = g(X; \beta^*)$ we often specify a loss function ℓ and perform empirical risk minimisation (6) as an approximation of the parameter that is optimal in terms of the true risk (7).

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} \ell(g(X_i; \beta), Y_i) \quad (6)$$

$$\beta^* = \arg \min_{\beta} \mathbb{E}[\ell(g(X; \beta), Y)] \quad (7)$$

Writing $f(x, y) := \ell(g(x; \beta), y)$, we can express the risk in terms of the distribution conditional on $S = 1$ using importance weighting:

$$\mathbb{E}[f(X, Y)] = \mathbb{E}[w(X, Z)f(X, Y)|S = 1], \quad (8)$$

with *importance weights* $w(x, z) := \mathbb{P}(S = 1)/\mathbb{P}(S = 1|X = x, Z = z)$, provided we have that $\mathbb{P}(S = 1|X, Z) > 0$ holds $\mathbb{P}(X, Z)$ -almost surely. A derivation of these importance weights can be found in the supplementary material. Since $\beta^* = \arg \min_{\beta} \mathbb{E}[w(X, Z)\ell(g(X; \beta), Y)|S = 1]$, when we have observations $(X_i, Z_i, Y_i) \sim \mathbb{P}(X, Z, Y|S = 1)$ for $i \in \mathcal{S}$, we can directly perform empirical risk minimization on this dataset using the weighted loss:

$$\hat{\beta}_w := \arg \min_{\beta} \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} w(X_i, Z_i)\ell(g(X_i; \beta), Y_i), \quad (9)$$

and use $\hat{\beta}_w$ as an estimate of β^* . Many implementations of regression algorithms allow the user to specify such sample weights. When the weights are not known in the missingness setting, they can be estimated [Cortes et al., 2008]. Practically, when the selection probability $\mathbb{P}(S = 1|x_i, z_i)$ is nearly zero for particular drawn values x_i, z_i , these probabilities are ‘clipped’ (i.e. transformed to remain lower bounded by some pre-specified strictly positive value) to reduce variance. In the experiments, we consider a linear map of the array of selection probabilities to $[1/20, 1]$, as this yields the best performance among different clipping strategies.

Example 1 (continued). The fitted IW regression model is depicted in Figure 6. Here, we used the true weights $w(x_i, z_i)$. The size of the observed points indicate the associated weight. Note that the IW regression only uses the black points (i.e. the (x_i, y_i) pairs for which $S_i = 1$). Comparing to repeated regression (Figure 5) we see that IW extrapolates poorly, possibly due to the effect of importance weighting on regularization [Xu et al., 2021]. On the other hand, IW interpolates better than the naive model. \triangle




Figure 6: IW regression with true weights.

The idea of weighting observations stems from the Horvitz-Thompson estimator for the population mean [Horvitz and

Thompson, 1952]. The use of importance weights in GEE regression in the MAR setting is analysed by [Robins et al., 1994, Robins and Rotnitzky, 1995], and its MNAR counterpart is described by Scharfstein et al. [1999]. When weighting is performed to correct for confounding bias in causal effect estimation, this is often called *inverse propensity weighting* [Hernán and Robins, 2021]. For weighted SVM- and kernel regression under covariate shift and target shift, see Zhang et al. [2013].

3.3 DOUBLY ROBUST REGRESSION

We now follow a relatively standard procedure to combine the two previous models into a doubly robust model [Bang and Robins, 2005, Kang and Schafer, 2007, Dudík et al., 2014]. First, we consider the repeated regression method $\hat{\mu}_{RR}$ from Section 3.1. We calculate the residuals of this method on the available Y values: we set $R_i := Y_i - \hat{\mu}_{RR}(X_i)$ for all $i \in \mathcal{S}$. We model the conditional expectation of the residuals given X using the IW regression method from Section 3.2, as

$$\hat{r}_{IW}(x) := g(X; \hat{\beta}_w) \approx \mathbb{E}[R|X] \quad (10)$$

where $\hat{\beta}_w = \arg \min_{\beta} |\mathcal{S}|^{-1} \sum_{i \in \mathcal{S}} w(X_i, Z_i) \ell(g(X_i; \beta), R_i)$ and where we use the same importance weights w as in Section 3.2. Then we define the *doubly robust* model

$$\hat{\mu}_{DR}(x) := \hat{\mu}_{RR}(x) + \hat{r}_{IW}(x). \quad (11)$$

Here, double robustness refers to the fact that for this model to be consistent, only one of $\hat{\mu}_{RR}$ and \hat{r}_{IW} has to be consistent.

Example 1 (continued). For sake of exposition, we apply the doubly robust method with a misspecified model $\hat{\mu}_{RR}(x)$. Recall that $\hat{\mu}_{RR}(x) = \hat{\mathbb{E}}[\tilde{Y}|X = x]$, where \tilde{Y} are the values that are imputed with the model $\hat{\mathbb{E}}[Y|X, Z, S = 1]$. To make the RR model deliberately misspecified, we use the same \tilde{Y} as in Section 3.1 (more specifically, as in Figure 5), but for the outer regression $\hat{\mathbb{E}}[\tilde{Y}|X]$ we employ polynomial regression with degree 5. The resulting models are depicted in Figure 7. We see that $\hat{\mu}_{DR}$ extrapolates poorly, but it interpolates better than the misspecified model $\hat{\mu}_{RR}$. \triangle

For overview papers on doubly robust estimation of the mean, see Bang and Robins [2005] and Kang and Schafer [2007]. For doubly robust regression in the MNAR setting, see Rotnitzky et al. [1998] and Scharfstein et al. [1999]. Doubly robust estimation of causal effects is treated by Coston et al. [2020] and Bhattacharya et al. [2022], among others.

4 EVALUATION

Consider the setup where we split the available data $\mathcal{D} = \mathcal{D}'' \cup \mathcal{D}'$ into train- and test test respectively, and




Figure 7: Doubly robust regression. At the top the (X, Y) pairs are plotted, with the true model (green) and misspecified RR model (black). The residuals of the misspecified model are plotted at the bottom. The IW regression model for the residuals (yellow) is then added to the misspecified model, which yields the doubly robust model (orange).

where we split \mathcal{S} accordingly. When all values of Y are known, performance is typically assessed using $\text{MSE} := |\mathcal{D}'|^{-1} \sum_{i \in \mathcal{D}'} (\hat{y}_i - y_i)^2$. Under missingness or selection bias, this quantity cannot be calculated as it requires the unobserved values y_i for $i \notin \mathcal{S}'$. Quantities that *can* be calculated on the test set are as follows:

$$\text{MSE-n} := \frac{1}{|\mathcal{S}'|} \sum_{i \in \mathcal{S}'} (\hat{y}_i - y_i)^2 \quad (12)$$

$$\text{MSE-w} := \frac{1}{|\mathcal{S}'|} \sum_{i \in \mathcal{S}'} w(x_i, z_i) (\hat{y}_i - y_i)^2 \quad (13)$$

$$\text{MSE-}\tilde{y} := \frac{1}{|\mathcal{D}'|} \sum_{i \in \mathcal{D}'} (\hat{y}_i - \tilde{y}_i)^2, \quad (14)$$

which are respectively the naively calculated MSE, the re-weighted MSE using the true weights, and the MSE calculated as if the imputed values \tilde{y} were true. We define $\text{MSE-}\hat{w}$ similarly, in terms of the estimated weights.

Example 1 (continued). The IW and DR models are learned both with true and estimated weights (denoted with -t and -e respectively), and all weights are clipped. We calculate the different mean squared error metrics for the estimated models on 500 independently drawn test sets of $n = 400$ samples. The results are provided in Table 1. We emphasize that MSE is not calculable in practice, and note that MSE-n , MSE-w and $\text{MSE-}\hat{w}$ yield qualitatively different results than the oracle MSE. $\text{MSE-}\tilde{y}$ provides the same ordering as MSE, but in practice the reliance on the imputation model can be undesirable. We observe that RR appropriately corrects for bias, as its MSE lies close to the MSE of the true function $\mathbb{E}[Y|X]$. Additional to the oracle MSE, we calculate the MSEs specifically for data points where the regression model is interpolating (between the minimum and maximum values of X for which $S = 1$) and extrapolating

	MSE	MSE-n	MSE- \tilde{y}	MSE- w	MSE- \hat{w}
Naive	24.07	7.56	19.54	16.19	19.79
RR	8.81	8.13	4.44	8.50	9.34
IW-t	40.42	8.40	35.58	22.22	24.67
IW-e	38.44	8.28	33.60	21.49	23.08
DR-t	25.91	8.13	21.27	14.82	16.65
DR-e	24.09	8.00	19.43	13.96	15.52
True	8.03	7.78	4.63	7.87	8.87

Table 1: Mean squared errors of different regression methods applied to Example 1, averaged over 500 test sets.

(the complement of these values), as depicted in Table 2. We observe that IW extrapolates poorly, especially compared to RR. \triangle

	MSE	MSE-interp.	MSE-extrap.
RR	8.81	8.43	10.17
IW-t	40.42	12.00	138.22
IW-e	38.44	11.77	129.41
DR-t	25.91	10.47	78.48
DR-e	24.09	10.23	70.42

Table 2: Inter- and extrapolation results for Example 1.

When the true weights are not known, one can only evaluate the proposed bias correction methods on biased data by relying on an auxiliary model, i.e. either the imputation model $\hat{\mathbb{E}}[Y|X, Z, S = 1]$ or $\hat{\mathbb{P}}(S = 1|X, Z)$. In general, if the practitioner is able to evaluate on an unbiased random sample of (X, Y) , this would be much more reliable than evaluating on a biased sample.

5 EXPERIMENTS

The example that is used throughout this paper is quite extreme in the sense that the model $\mathbb{E}[Y|X, Z]$ is relatively simple and $\mathbb{E}[Y|X]$ is rather complex, which gives RR a clear advantage over IW. To investigate the proposed methods in a more general setting, we assess performance on extensive simulations and on the Boston Housing dataset.⁴

5.1 SIMULATIONS

To assess the performance of the proposed methods empirically, we first identify which graphs with variables X, Y, Z and S satisfy the pattern of d-separations $X \not\perp\!\!\!\perp Y$ (such that regression $\mathbb{E}[Y|X]$ makes sense), $Y \not\perp\!\!\!\perp S|X$ and $Y \perp\!\!\!\perp S|X, Z$ (PMAR). There are 550 *Acyclic Directed*

⁴Code for the experiments is publicly available at https://github.com/philipboeken/debiased_regression.

Mixed Graphs (ADMGs) that fit this pattern. For each ADMG we simulate 50 datasets, and for each dataset we draw 2000 samples according to an additive noise structural equation model, where all equations for X, Y, Z are random functions, independently drawn from a Gaussian process, and the additive noise is the pushforward of a $\text{Unif}[0, 1]$ distribution through a random function that is also drawn from a Gaussian process [Mooij et al., 2016]. We let the selection probability be $p(x, z) = \sigma(x)\sigma(z)$, when X and Z are parents of S in the ADMG. We make a 50/50 train-test split for training and evaluation. For a complete description of the simulation setup, we refer to the supplements. On each dataset, we fit a ‘naive’ model $\hat{\mathbb{E}}[Y|X, S = 1]$ and the models $\hat{\mu}_{RR}, \hat{\mu}_{IW}$ and $\hat{\mu}_{DR}$ from Section 3. We also fit a ‘true’ model $\hat{\mathbb{E}}[Y|X]$ for comparison. In Table 3 we report

	MSE	MSE- \tilde{y}	MSE- w	MSE- \hat{w}
Naive	3.11 (20.6)	2.81 (18.0)	0.91 (1.3)	0.90 (0.6)
RR	2.01 (2.6)	0.60 (0.9)	1.25 (1.5)	1.20 (1.3)
IW-t	4.91 (33.6)	4.57 (30.9)	0.92 (1.5)	0.92 (0.6)
IW-e	4.18 (23.2)	3.83 (20.7)	0.92 (1.2)	0.91 (0.6)
DR-t	4.98 (29.8)	4.97 (31.1)	0.77 (0.4)	0.81 (0.4)
DR-e	4.51 (45.1)	4.47 (49.0)	0.80 (0.4)	0.79 (0.4)
True	0.98 (0.3)	1.63 (3.0)	1.03 (0.6)	1.00 (0.5)

Table 3: Results over 27,500 simulated datasets.

the different mean squared error metrics of the methods, averaged over all 27,500 tests sets (standard deviations are shown in parentheses). We see that the average MSE of RR is half that of the best IW method. The median MSEs follow the relation $RR < \text{Naive} < \text{IW-t}$, confirmed with respective p-values $1.6 \cdot 10^{-110}$ and $3.2 \cdot 10^{-39}$ of Wilcoxon’s signed rank test. In ADMGs where X is a parent of S there is a clear region of the support of X with no positivity of S , hence a clear region where the regression model is extrapolating, similar to Example 1. When singling out these ADMGs we calculate the average MSE, and additionally the MSEs on the parts where the regression model is inter- and extrapolating, as shown in the table below. From Table 4 we see that IW and DR do not extrapolate well, while RR does.

	MSE	MSE-interp.	MSE-extrap.
RR	2.13 (2.7)	1.46 (1.8)	2.89 (4.2)
IW-t	7.29 (34.0)	1.29 (0.7)	13.75 (53.8)
IW-e	5.82 (16.4)	1.27 (0.6)	10.81 (32.6)
DR-t	7.73 (45.7)	1.26 (0.6)	14.37 (70.1)
DR-e	6.93 (72.3)	1.24 (0.6)	12.34 (94.9)

Table 4: Inter- and extrapolation results for simulated data.

We hypothesize that IW and DR extrapolate badly because of the interplay between the smoothness assumption and reg-

ularization of thin plate spline regression, and large weights near the edge of the support of $\mathbb{P}(X|S = 1)$, causing the model to diverge away from the true $\mathbb{E}[Y|X]$ when extrapolating (as is for example the case in Figure 6). We investigate this by running the experiments with RR, IW and DR implemented with regression trees [Breiman et al., 1984] instead of thin plate regression. As regression trees extrapolate flatly, this can allow for better extrapolation performance than thin plate regression. The results of these experiments are provided in the supplements. They are indeed numerically less extreme than for thin plate regression, but qualitatively the same.

5.2 BOSTON HOUSING DATA

In absence of a real-world dataset where there is missingness *and* the missing value is known, we consider the Boston Housing dataset [Harrison Jr and Rubinfeld, 1978] and simulate a selection mechanism ourselves. We consider the problem of predicting the value of owner-occupied homes in US Dollars (Y) from the number of rooms per dwelling (X). We let selection depend on X , and on the ‘percentage of people of lower status of the population in the town where the house is situated’ (Z), which strongly correlates with Y . Selection is simulated by setting $\mathbb{P}(S = 1|X, Z) := \sigma(f_1(X))\sigma(f_2(Z))$, where f_1, f_2 are random functions drawn from a Gaussian process. For each dataset, the 506 available observations are randomly split into evenly sized train and test sets. The average MSEs are provided in Table 5, and show the same qualitative results as the simulations. When calculating the MSE of the inter- and extrapolation regions, we observe that RR performs better than IW at both tasks (table shown in supplements).

	MSE	MSE- \tilde{y}	MSE- w	MSE- \hat{w}
Naive	1.23 (2.5)	0.84 (2.3)	0.48 (0.7)	1.79 (5.9)
RR	0.71 (0.3)	0.26 (0.3)	0.44 (0.4)	1.63 (4.8)
IW-t	2.18 (4.9)	1.77 (4.9)	0.56 (1.5)	2.40 (9.2)
IW-e	1.75 (4.4)	1.40 (4.7)	0.50 (0.9)	1.98 (7.4)
DR-t	1.92 (3.7)	1.65 (3.4)	0.23 (0.2)	0.39 (1.2)
DR-e	2.43 (5.4)	2.20 (5.5)	0.25 (0.2)	0.21 (0.3)
True	0.45 (0.1)	0.54 (0.4)	0.36 (0.3)	0.85 (2.2)

Table 5: Results of 500 biased instantiations of the Boston Housing dataset.

6 DISCUSSION AND CONCLUSION

In this work, we have motivated the use of *privileged information* for estimating a regression model when selection or missingness is nonignorable, and introduced the *Privilegedly Missing at Random* (PMAR) setting. We formulated

the repeated regression method, the IW regression method, and the doubly robust combination of the two. We note that evaluation of regression methods on biased data is not straightforward and relies on auxiliary models. Experiments show that repeated regression can appropriately correct for bias, and with considerable advantage over IW regression. In particular, in repeated regression, extrapolation is facilitated by privileged data Z that is predictive of Y . IW does not have such a property, and extrapolates worse than RR.

Further research can be done on the statistical properties of the proposed methods, e.g. on bounding the regret of DR in terms of the regret of RR and IW (possibly under misspecification), or on the interplay of importance weighting, regularization, and extrapolation.

Additionally, further research on the sensitivity of these methods with respect to the conditional independence assumption would be valuable, e.g. to determine whether it’s always better to correct for bias with some privileged information Z , even when the independence $Y \perp\!\!\!\perp S | X, Z$ is not met.

Finally, it would be interesting to see whether the improved performance of repeated regression over weighted regression translates to their counterparts for estimating causal effects, i.e. whether standardization should be preferred over inverse propensity weighting.

Acknowledgements

This work is supported by Booking.com. We thank Kees Jan de Vries for fruitful discussions and Christina Katsimerou, Nils Skotara, Kostas Tokis and anonymous reviewers for their feedback on the manuscript.

References

- H. Bang and J. M. Robins. Doubly Robust Estimation in Missing Data and Causal Inference Models. *Biometrics*, 61(4):962–972, 2005. ISSN 0006-341X. URL <https://www.jstor.org/stable/3695907>.
- E. Bareinboim, J. Tian, and J. Pearl. Recovering from Selection Bias in Causal and Statistical Inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, page 12, 2014.
- R. Bhattacharya, R. Nabi, and I. Shpitser. Semiparametric Inference For Causal Effects In Graphical Models With Hidden Variables. *Journal of Machine Learning Research*, 23(295):1–76, 2022. ISSN 1533-7928. URL <http://jmlr.org/papers/v23/20-296.html>.
- L. Breiman, J. Friedman, C. Stone, and R. Olshen. *Classification and Regression Trees*. CRC Press, 1984.

- C. Cortes, M. Mohri, M. Riley, and A. Rostamizadeh. Sample Selection Bias Correction Theory. In Y. Freund, L. Györfi, G. Turán, and T. Zeugmann, editors, *Algorithmic Learning Theory*, volume 5254, pages 38–53. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008. ISBN 978-3-540-87986-2 978-3-540-87987-9. doi: 10.1007/978-3-540-87987-9_8. URL http://link.springer.com/10.1007/978-3-540-87987-9_8.
- A. Coston, E. Kennedy, and A. Chouldechova. Counterfactual Predictions under Runtime Confounding. In *Advances in Neural Information Processing Systems*, volume 33, pages 4150–4162. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/2b64c2f19d868305aa8bbc2d72902cc5-Abstract.html>.
- R. M. Daniel, M. G. Kenward, S. N. Cousens, and B. L. De Stavola. Using causal diagrams to guide analysis in missing data problems. *Statistical Methods in Medical Research*, 21(3):243–256, June 2012. ISSN 0962-2802, 1477-0334. doi: 10.1177/0962280210394469. URL <http://journals.sagepub.com/doi/10.1177/0962280210394469>.
- A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38, 1977. ISSN 0035-9246. URL <https://www.jstor.org/stable/2984875>.
- K. E. Dithlong, O. O. Ngesa, and A. Y. Kombo. A Comparative Analysis of Generalized Estimating Equations Methods for Incomplete Longitudinal Ordinal Data with Ignorable Dropouts. *Open Journal of Statistics*, 8(5):770–792, Sept. 2018. doi: 10.4236/ojs.2018.85051. URL <http://www.scirp.org/Journal/Paperabs.aspx?paperid=87364>.
- J. Duchon. Splines minimizing rotation-invariant seminorms in Sobolev spaces. In W. Schempp and K. Zeller, editors, *Constructive Theory of Functions of Several Variables*, Lecture Notes in Mathematics, pages 85–100. Berlin, Heidelberg, 1977. Springer. ISBN 978-3-540-37496-1. doi: 10.1007/BFb0086566.
- M. Dudík, D. Erhan, J. Langford, and L. Li. Doubly Robust Policy Evaluation and Optimization. *Statistical Science*, 29(4), Nov. 2014. ISSN 0883-4237. doi: 10.1214/14-STS500. URL <http://arxiv.org/abs/1503.02834>.
- N. Goel, A. Amayuelas, A. Deshpande, and A. Sharma. The Importance of Modeling Data Missingness in Algorithmic Fairness: A Causal Perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9):7564–7573, May 2021. ISSN 2374-3468. doi: 10.1609/aaai.v35i9.16926. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16926>.
- L. Guerdan, A. Coston, Z. S. Wu, and K. Holstein. Ground(less) Truth: A Causal Framework for Proxy Labels in Human-Algorithm Decision-Making, Feb. 2023. URL <http://arxiv.org/abs/2302.06503>.
- D. Harrison Jr and D. L. Rubinfeld. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 5(1):81–102, 1978.
- J. J. Heckman. Sample Selection Bias as a Specification Error. *Econometrica*, 47(1):153, Jan. 1979. ISSN 00129682. doi: 10.2307/1912352. URL <https://www.jstor.org/stable/1912352?origin=crossref>.
- M. Hernán and J. M. Robins. *Causal Inference*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Chapman & Hall/CRC, Boca Raton, 2021. ISBN 978-1-4200-7616-5.
- D. G. Horvitz and D. J. Thompson. A Generalization of Sampling Without Replacement From a Finite Universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952. ISSN 0162-1459. doi: 10.2307/2280784. URL <https://www.jstor.org/stable/2280784>.
- J. D. Y. Kang and J. L. Schafer. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. *Statistical Science*, 22(4), Nov. 2007. ISSN 0883-4237. doi: 10.1214/07-STS227. URL <http://arxiv.org/abs/0804.2958>.
- R. J. Little and D. B. Rubin. *Statistical analysis with missing data*. John Wiley & Sons, 2019.
- G. Martius and C. H. Lampert. Extrapolation and learning equations, Oct. 2016. URL <http://arxiv.org/abs/1610.02995>.
- K. Mohan and J. Pearl. Graphical Models for Processing Missing Data. *Journal of the American Statistical Association*, 116(534):1023–1037, Apr. 2021. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2021.1874961. URL <https://www.tandfonline.com/doi/full/10.1080/01621459.2021.1874961>.
- J. M. Mooij, J. Peters, D. Janzing, J. Zscheischler, and B. Schölkopf. Distinguishing Cause from Effect Using Observational Data: Methods and Benchmarks. *Journal of Machine Learning Research*, 17(32):1–102, 2016. ISSN 1533-7928. URL <http://jmlr.org/papers/v17/14-518.html>.

- R. Nabi, R. Bhattacharya, and I. Shpitser. Full Law Identification in Graphical Models of Missing Data: Completeness Results. In *Proceedings of the 37th International Conference on Machine Learning*, pages 7153–7163. PMLR, Nov. 2020. URL <https://proceedings.mlr.press/v119/nabi20a.html>.
- J. Pearl. *Causality*. Cambridge university press, 2009.
- J. Pearl. A solution to a class of selection-bias problems, 2012.
- J. M. Robins and A. Rotnitzky. Semiparametric Efficiency in Multivariate Regression Models with Missing Data. *Journal of the American Statistical Association*, 90(429):122–129, 1995. ISSN 0162-1459. doi: 10.2307/2291135. URL <https://www.jstor.org/stable/2291135>.
- J. M. Robins, A. Rotnitzky, and L. P. Zhao. Estimation of Regression Coefficients When Some Regressors Are Not Always Observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994. ISSN 0162-1459. doi: 10.2307/2290910. URL <https://www.jstor.org/stable/2290910>.
- P. R. Rosenbaum and D. B. Rubin. Reducing Bias in Observational Studies Using Subclassification on the Propensity Score. *Journal of the American Statistical Association*, 79(387):516–524, 1984. ISSN 0162-1459. doi: 10.2307/2288398. URL <https://www.jstor.org/stable/2288398>.
- A. Rotnitzky, J. M. Robins, and D. O. Scharfstein. Semiparametric Regression for Repeated Outcomes with Nonignorable Nonresponse. *Journal of the American Statistical Association*, 93(444):1321–1339, 1998. ISSN 0162-1459. doi: 10.2307/2670049. URL <https://www.jstor.org/stable/2670049>.
- D. B. Rubin. Inference and Missing Data. *Biometrika*, 63(3):581–592, 1976. ISSN 0006-3444. doi: 10.2307/2335739. URL <https://www.jstor.org/stable/2335739>.
- D. O. Scharfstein, A. Rotnitzky, and J. M. Robins. Adjusting for Nonignorable Drop-Out Using Semiparametric Nonresponse Models. *Journal of the American Statistical Association*, 94(448):1096–1120, 1999. ISSN 0162-1459. doi: 10.2307/2669923. URL <https://www.jstor.org/stable/2669923>.
- B. Schölkopf, D. Janzing, J. Peters, E. Sgouritsa, K. Zhang, and J. Mooij. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning, ICML’12*, pages 459–466, Madison, WI, USA, June 2012. Omnipress. ISBN 978-1-4503-1285-1.
- M. Sugiyama, M. Krauledat, and K.-R. Müller. Covariate Shift Adaptation by Importance Weighted Cross Validation. *The Journal of Machine Learning Research*, 8: 985–1005, Dec. 2007. ISSN 1532-4435.
- F. Thoemmes and K. Mohan. Graphical Representation of Missing Data Problems. *Structural Equation Modeling: A Multidisciplinary Journal*, 22(4):631–642, Oct. 2015. ISSN 1070-5511, 1532-8007. doi: 10.1080/10705511.2014.937378. URL <http://www.tandfonline.com/doi/full/10.1080/10705511.2014.937378>.
- V. Vapnik and A. Vashist. A new learning paradigm: Learning using privileged information. *Neural Networks*, 22(5):544–557, July 2009. ISSN 0893-6080. doi: 10.1016/j.neunet.2009.06.042. URL <https://www.sciencedirect.com/science/article/pii/S0893608009001130>.
- Wei Fan, I. Davidson, B. Zadrozny, and Philip S. Yu. An Improved Categorization of Classifier’s Sensitivity on Sample Selection Bias. In *Fifth IEEE International Conference on Data Mining (ICDM’05)*, pages 605–608, Houston, TX, USA, 2005. IEEE. ISBN 978-0-7695-2278-4. doi: 10.1109/ICDM.2005.24. URL <http://ieeexplore.ieee.org/document/1565737/>.
- S. Wood. Package ‘mgcv’. *R package version*, 1(29):729, 2015.
- S. N. Wood. Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65(1):95–114, 2003. ISSN 1467-9868. doi: 10.1111/1467-9868.00374. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/1467-9868.00374>.
- D. Xu, Y. Ye, and C. Ruan. Understanding the role of importance weighting for deep learning, Mar. 2021. URL <http://arxiv.org/abs/2103.15209>.
- B. Zadrozny. Learning and evaluating classifiers under sample selection bias. In *Proceedings of the twenty-first international conference on Machine learning*, page 114, 2004.
- K. Zhang, B. Schölkopf, K. Muandet, and Z. Wang. Domain Adaptation under Target and Conditional Shift. In *Proceedings of the 30th International Conference on Machine Learning*, pages 819–827. PMLR, May 2013. URL <https://proceedings.mlr.press/v28/zhang13d.html>.

Correcting for Selection Bias and Missing Response in Regression using Privileged Information (Supplementary Material)

Philip Boeken^{1,2}

Noud de Kroon¹

Mathijs de Jong²

Joris M. Mooij¹

Onno Zoeter²

¹Korteweg-de Vries Institute for Mathematics, University of Amsterdam, The Netherlands

²Booking.com, The Netherlands

1 SELECTION BIAS AND MISSINGNESS

A schematic display of the available data under missingness and under selection bias with external data is provided in Figure 1. Note that under missingness we can estimate $\mathbb{P}(S = 1|X, Z)$ from the dataset \mathcal{D} ; under selection bias this is not possible.

		X	Z	S	Y	
{	S	x_1	z_1	1	y_1	$\mathbb{P}(X, Y, Z S = 1)$
		\vdots				
		x_m	z_m	1	y_m	
		x_{m+1}	z_{m+1}	0	y_{m+1}	
		x_n	z_n	0	y_n	
		$\mathbb{P}(X, Z, S)$				

(a) Missing response

		X	Z	S	Y	
{	S	x_1	z_1	1	y_1	$\mathbb{P}(X, Y, Z S = 1)$
		\vdots				
		x_m	z_m	1	y_m	
		x_{m+1}	z_{m+1}	0	y_{m+1}	
		x_n	z_n	0	y_n	
		$\mathbb{P}(X, Z, S)$				

		X	Z	
{	\mathcal{D}	x_1	z_1	$\mathbb{P}(X, Z)$
		\vdots	\vdots	
		\vdots	\vdots	
		\vdots	\vdots	
		x_n	z_n	
		$\mathbb{P}(X, Z)$		

(b) Selection bias with external data

Figure 1: Available data under missingness and under selection bias with external data. Grayed-out areas indicate unobserved data.

In the PMAR setting we have $Y \perp\!\!\!\perp S | X, Z$ and we are only given values of X at test time; our target is to estimate the function $\mathbb{E}[Y|X]$.

2 IMPORTANCE WEIGHTING

For estimating the parameter β^* in the regression model $\mathbb{E}[Y|X] = g(X; \beta^*)$ we often specify a loss function ℓ and perform empirical risk minimisation (1) as an approximation of the optimal parameter in terms of the true risk (2).

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{n} \sum_{i=1}^n \ell(g(X_i; \beta), Y_i) \quad (1)$$

$$\beta^* = \arg \min_{\beta} \mathbb{E}[\ell(g(X; \beta), Y)] \quad (2)$$

Writing $f(x, y) := \ell(g(x; \beta), y)$, we can express the risk in terms of the distribution conditional on $S = 1$ using importance weighting:

$$\begin{aligned} \mathbb{E}[f(X, Y)] &= \int f(x, y) p(x, y, z) d(x, y, z) \\ &= \int f(x, y) p(x, y, z) \frac{p(x, y, z | S = 1)}{p(x, y, z | S = 1)} d(x, y, z) \\ &= \int f(x, y) \frac{p(x, y, z) \mathbb{P}(S = 1)}{p(x, y, z, S = 1)} p(x, y, z | S = 1) d(x, y, z) \\ &= \int f(x, y) \frac{\mathbb{P}(S = 1)}{\mathbb{P}(S = 1 | x, y, z)} p(x, y, z | S = 1) d(x, y, z) \\ &= \int f(x, y) \frac{\mathbb{P}(S = 1)}{\mathbb{P}(S = 1 | x, z)} p(x, y, z | S = 1) d(x, y, z) \\ &= \int f(x, y) w(x, z) p(x, y, z | S = 1) d(x, y, z) \\ &= \mathbb{E}[w(X, Z) f(X, Y) | S = 1], \end{aligned} \quad (3)$$

where in the fifth equation we use the conditional independence $Y \perp\!\!\!\perp S | X, Z$, and where we define the *importance weights*

$$w(x, z) := \frac{\mathbb{P}(S = 1)}{\mathbb{P}(S = 1 | x, z)}. \quad (4)$$

Since $\beta^* = \arg \min_{\beta} \mathbb{E}[w(X, Z) \ell(g(X; \beta), Y) | S = 1]$, when we have observations $(X_1, Z_1, Y_1), \dots, (X_n, Z_n, Y_n) \sim \mathbb{P}(X, Z, Y | S = 1)$, we can directly perform empirical risk minimization on this dataset using the weighted loss:

$$\hat{\beta} = \arg \min_{\beta} \frac{1}{n} \sum_{i=1}^n w(x_i, z_i) \ell(g(x_i; \beta), y_i). \quad (5)$$

3 SIMULATIONS

We performed a brute-force search of all *Acyclic Directed Mixed Graphs* (ADMGs) that satisfy $X \not\perp\!\!\!\perp Y$ and the PMAR pattern of d-separations $Y \not\perp\!\!\!\perp S, Y \not\perp\!\!\!\perp S | X$, and $Y \perp\!\!\!\perp S | X, Z$, using the `pcalg` package [Kalisch et al., 2012], resulting in 550 ADMGs. For reference, all 55 DAGs are depicted in figures 2 and 3, where the graphs are categorised by whether S has any children or not. The remaining 495 ADMGs are not depicted here.

For each of the 550 ADMGs, we simulate 50 datasets according to the following procedure. Throughout, let V, S, X etc. denote vertices in the graph, and let x_V, x_S, x_X denote vectors of dimension $n = 2000$ for the simulated values.

- First, we replace any bidirected edge pointing to variables V by a variable U_V , and let V be the children of U_V . This turns the ADMG into a DAG.
- Then, we calculate a topological order of the DAG.
- For every variable V in the topological order, we simulate 2000 observations as follows:
 - If V has no parents, then
 - * if $V = S$, draw $x_S \sim \text{Bernoulli}(1/3)$;

- * otherwise, $V \neq S$ and we draw $x_V \sim \mathcal{RD}$, as defined below.
- Otherwise, denote the parents of V with \mathbf{Pa} and their value $\mathbf{x}_{\mathbf{Pa}}$, and then
 - * if $V = S$, draw $(x_S)_i \sim \text{Bernoulli}(p((x_{\mathbf{Pa}})_i))$ where

$$p((x_{\mathbf{Pa}})_i) := \prod_{v \in \mathbf{Pa}} \sigma((x_v)_i) \quad (6)$$

and $\sigma(x) := (1 + e^{20x})^{-1}$;

- * otherwise, $V \neq S$ and

- draw a random function f_V from a Gaussian process on $\mathbb{R}^{|\mathbf{Pa} \setminus \{S\}|}$ as $f_V \sim \mathcal{GP}(0, K_M)$;
- draw noise $\varepsilon_V \sim \mathcal{RD}$ and set

$$x_V := f_V(\mathbf{x}_{\mathbf{Pa} \setminus \{S\}}) + \frac{1}{2}\varepsilon_V; \quad (7)$$

- if $S \in \mathbf{Pa}$, calculate the empirical standard deviation $c := \text{sd}(x_V)$ and set $(x_V)_i := (x_V)_i - 3c$ for all i where $(x_S)_i = 1$;
- then, standardize x_V .

Drawing from a Gaussian process The kernels used for calculating the covariance matrix for drawing from a Gaussian process are the Matérn kernel and the squared exponential kernel, being

$$K_M(x, y) := (1 + \sqrt{5}\|x - y\| + \frac{5}{3}\|x - y\|^2)e^{-\sqrt{5}} \quad (8)$$

$$K_{SE}(x, y) := \frac{1}{4}e^{\frac{2}{5}\|x - y\|^2} \quad (9)$$

respectively, where $\|\cdot\|$ denotes the Euclidean norm. Then, given input $x_i \in \mathbb{R}^d$ for $i = 1, \dots, n$ and kernel $K : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, denote $\mathbf{x} := (x_i)_{i=1}^n$ and draw $f(\mathbf{x}) \sim \mathcal{N}(0, (K(x_i, x_j))_{i,j})$, where i and j run over $\{1, \dots, n\}$ in the kernel matrix $(K(x_i, x_j))_{i,j}$. This draw is denoted with $f \sim \mathcal{GP}(0, K)$.

Drawing from a random distribution Drawing noise from a random distribution, denoted with $\varepsilon \sim \mathcal{RD}$, is done as follows:

- First, draw 2000 i.i.d. samples $U \sim \text{Unif}[0, 1]$.
- Then, draw a random function $f_\varepsilon \sim \mathcal{GP}(0, K_{SE})$
- Set $\varepsilon := f_\varepsilon(U)$, and standardize.




Figure 2: PMAR DAGs where S is a sink node.




Figure 3: PMAR DAGs where S is not a sink node.

3.1 SIMULATION EXPERIMENTS WITH REGRESSION TREES

In the main paper, we hypothesize that the bad extrapolation performance of weighted regression is caused by a detrimental effect of the weights on the regularization of the regression method. In the main paper we use thin plate spline regression, which doesn't necessarily extrapolate flatly and can diverge away from the true $\mathbb{E}[Y|X]$, yielding large MSE values for IW and DR. Tables 1 and 2 show the simulation results when using regression trees as implemented in the `rpart` package, instead of thin plate regression. The results are numerically less extreme, but qualitatively the same as for thin plate regression, as depicted in the main paper.

	MSE	MSE- \tilde{y}	MSE- w	MSE- \hat{w}
Naive	1.53 (0.6)	0.62 (0.8)	0.96 (0.6)	0.96 (0.5)
RR	1.38 (0.6)	0.27 (0.3)	0.97 (0.6)	0.94 (0.5)
IW-t	1.54 (0.6)	0.63 (0.8)	0.97 (0.6)	0.98 (0.5)
IW-e	1.52 (0.6)	0.61 (0.7)	0.97 (0.6)	0.97 (0.5)
DR-t	1.48 (0.7)	0.58 (0.7)	0.66 (0.3)	0.73 (0.4)
DR-e	1.46 (0.6)	0.55 (0.7)	0.71 (0.4)	0.69 (0.3)
True	1.00 (0.2)	0.69 (0.7)	1.05 (0.6)	1.02 (0.5)

Table 1: Results over 27.500 simulated datasets.

	MSE	MSE-interp.	MSE-extrap.
Naive	1.58 (0.6)	1.37 (0.6)	1.82 (1.0)
RR	1.44 (0.6)	1.22 (0.6)	1.68 (1.0)
IW-t	1.60 (0.6)	1.38 (0.6)	1.85 (1.1)
IW-e	1.58 (0.6)	1.37 (0.6)	1.82 (1.0)
DR-t	1.55 (0.7)	1.30 (0.6)	1.84 (1.1)
DR-e	1.52 (0.6)	1.28 (0.6)	1.80 (1.0)
True	1.01 (0.2)	1.03 (0.3)	0.98 (0.3)

Table 2: Interpolation and extrapolation results of regression trees for simulated data, on graphs with $X \rightarrow S$.

4 BOSTON HOUSING DATA

Considering the Boston Housing Dataset [Harrison Jr and Rubinfeld, 1978], let the variables X , Y and Z be ‘the number of rooms per dwelling’, ‘the value of owner-occupied homes in US Dollars’ and ‘percentage of people of lower status of the population’ respectively. We sample the selection probability

$$p(X, Z) := \sigma(f_1(X))\sigma(f_2(Z)) \quad (10)$$

where $f_1, f_2 \sim \mathcal{GP}(0, K_{SE})$ independent. Then we draw $U_i \sim \text{Unif}[0, 1]$ for $i = 1, \dots, 506$, and set

$$S_i := \mathbb{1}\{p(X, Z)_i < U_i\} \quad (11)$$

for all i . The dataset is resampled when $\#\{S = 1\} < 120$. One realisation of such a dataset with an overview of all regression methods is provided in Figure 4.

The MSE and the interpolated and extrapolated variants, as calculated on the Boston Housing dataset, are shown in Table 3. We observe that RR performs better than IW and DR on all three metrics.

References

- D. Harrison Jr and D. L. Rubinfeld. Hedonic housing prices and the demand for clean air. *Journal of environmental economics and management*, 5(1):81–102, 1978.
- M. Kalisch, M. Mächler, D. Colombo, M. H. Maathuis, and P. Bühlmann. Causal inference using graphical models with the R package `pcalg`. *Journal of Statistical Software*, 47(11):1–26, 2012. doi: 10.18637/jss.v047.i11.




Figure 4: An instantiation of the biased Boston Housing dataset.

	MSE	MSE-interp.	MSE-extrap.
Naive	1.23 (2.5)	0.57 (0.5)	3.88 (9.1)
RR	0.71 (0.3)	0.47 (0.2)	2.05 (4.3)
IW-t	2.18 (4.9)	0.62 (0.8)	7.68 (17.7)
IW-e	1.75 (4.4)	0.58 (0.6)	5.96 (16.1)
DR-t	1.92 (3.7)	0.48 (0.3)	7.37 (16.4)
DR-e	2.43 (5.4)	0.52 (0.4)	9.42 (22.1)

Table 3