

---

# Causal Transfer Learning

---

**Sara Magliacane**  
University of Amsterdam  
sara.magliacane@gmail.com

**Thijs van Ommen**  
University of Amsterdam  
thijsvanommen@gmail.com

**Tom Claassen**  
Radboud University Nijmegen  
tomc@cs.ru.nl

**Stephan Bongers**  
University of Amsterdam  
srbongers@gmail.com

**Philip Versteeg**  
University of Amsterdam  
p.j.j.p.versteeg@uva.nl

**Joris M. Mooij**  
University of Amsterdam  
j.m.mooij@uva.nl

## Abstract

An important goal in both transfer learning and causal inference is to make accurate predictions when the distribution of the test set and the training set(s) differ. Such a distribution shift may happen as a result of an external intervention on the data generating process, causing certain aspects of the distribution to change, and others to remain invariant. We consider a class of causal transfer learning problems, where multiple training sets are given that correspond to different external interventions, and the task is to predict the distribution of a target variable given measurements of other variables for a new (yet unseen) intervention on the system. We propose a method for solving these problems that exploits causal reasoning but does neither rely on prior knowledge of the causal graph, nor on the type of interventions and their targets. We evaluate the method on simulated and real world data and find that it outperforms a standard prediction method that ignores the distribution shift.

## 1 Introduction

Predicting unknown values based on observed training data is a problem central to many sciences, and well studied in statistics and machine learning. This problem becomes significantly harder if the training and test data do not come from the same distribution. Such a distribution shift will happen in practice whenever the circumstances under which the training data were gathered are different from those for which the predictions are to be made. A rich literature exists on this problem of *transfer learning*; see e.g. Quiñero-Candela et al. [2009], Pan and Yang [2010] for overviews.

When the setting changes, so do the relations between the different variables under consideration. While for some (sets of) variables  $\mathcal{A}$ , a function  $f : \mathcal{A} \rightarrow Y$  learned in one setting may continue to offer good predictions for  $Y$  in a different setting, this may not be true of other sets  $\mathcal{A}'$  of variables. *Causal graphs* [e.g., Pearl, 2009, Spirtes et al., 2000] allow us to reason about this in a principled way. Knowledge of the causal graph that describes the data generating mechanism, and which parts of the model are invariant under the different settings, allows one to transfer knowledge from one situation to the other [Spirtes et al., 2000, Storkey, 2009, Schölkopf et al., 2012, Bareinboim and Pearl, 2016].

In practice, the causal graph is often unknown, and estimating it from data is a challenging task. Therefore, practitioners often revert to standard prediction methods that ignore the distribution shift. The following example demonstrates an instance of a causal transfer learning problem where a feature

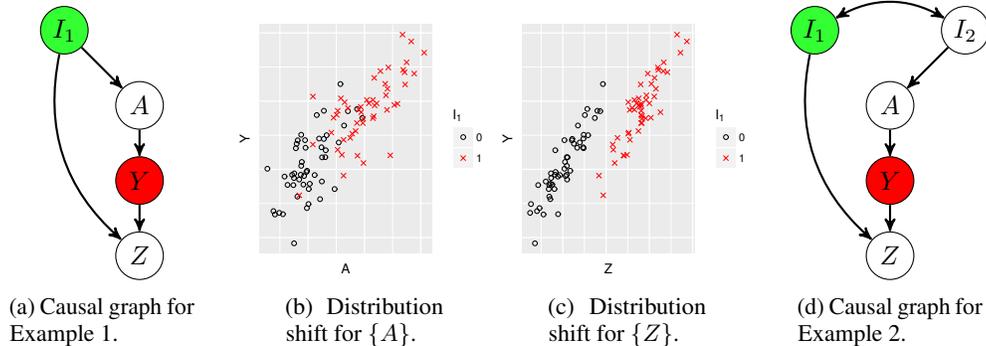


Figure 1: Two examples of a situation where the intervention  $I_1$  leads to a shift of distribution between training and test data. For graph (a), corresponding to Example 1, feature selection methods using only available training data ( $I_1 = 0$ ) will typically select  $\{Z\}$  or  $\{A, Z\}$  as optimal sets to predict  $Y$ . Because of the distribution shift, this may lead to arbitrarily bad predictions in the test setting ( $I_1 = 1$ ); plot (c) shows this for the case of feature set  $\{Z\}$  (black dots are drawn from the training distribution, red dots from the test distribution). Predictions based on  $\{A\}$  transfer much better to the test setting, as plot (b) shows: there may be a covariate shift (i.e.,  $\mathbb{P}(A | I_1 = 0) \neq \mathbb{P}(A | I_1 = 1)$ ), but  $\mathbb{P}(Y | A, I_1 = 0) = \mathbb{P}(Y | A, I_1 = 1)$ . A similar situation occurs in case (d), but there knowledge of the causal graph is not needed as it can be inferred from the data and certain assumptions that set  $\{A\}$  should be used to predict  $Y$  in the test setting, as discussed in Example 2.

selection method that is unaware of the causal structure will pick a set of features that does not generalize to the test setting, and may lead to arbitrarily bad predictions (even asymptotically). On the other hand, correctly taking into account the causal structure and the possible distribution shift from training to test setting allows to upper bound the prediction error in the test setting (as will be discussed in Section 3.2).

**Example 1.** Suppose we wish to predict the value of a variable  $Y$ , given observations of variables  $A$  and  $Z$ , under an intervention where  $I_1 = 1$ , making use of observational training data on  $(A, Y, Z)$  where  $I_1 = 0$ . As an example,  $A, Y, Z$  may be different blood cell phenotypes in mice, and the intervention variable  $I_1$  may indicate whether a certain gene has been knocked out. As shown in the causal graph in Figure 1(a), we assume that  $Y$  is affected by  $A$  and affects  $Z$ , while  $I_1$  affects both  $A$  and  $Z$ . Suppose further that the relation between  $A$  and  $Y$  is about equally strong as the relation between  $Y$  and  $Z$ , but considerably more noisy. Then a feature selection method aiming to select the best subset of features to use for prediction of  $Y$  will prefer both  $\{Z\}$  and  $\{A, Z\}$  over  $\{A\}$  (because predicting  $Y$  from  $A$  leads to larger variance than predicting  $Y$  from  $Z$ , and to a larger bias than predicting  $Y$  from both  $A$  and  $Z$ ). However, under the intervention ( $I_1 = 1$ ),  $\mathbb{P}(Y | Z)$  and  $\mathbb{P}(Y | A, Z)$  both change,<sup>1</sup> so that the predictions of  $Y$  in this setting may be extremely biased. Because the conditional distribution of  $Y$  given  $A$  does not change under  $I_1 = 1$ , i.e.,  $\mathbb{P}(Y | A, I_1 = 0) = \mathbb{P}(Y | A, I_1 = 1)$ , predictions of  $Y$  based only on  $A$  remain as accurate in the interventional setting as in the observational setting.

Over the last years, various methods have been proposed to exploit the causal structure of the data generating process in order to address transfer learning tasks, relying on different assumptions. For example, Bareinboim and Pearl [2016] provide theory for identifiability under transfer (“transportability”) assuming that the causal graph is known, that interventions are perfect, and that the intervention targets are known. Hyttinen et al. [2015] also assume perfect interventions with known targets but do not rely on complete knowledge of the causal graph, and rather infer the relevant aspects of it from the data. Rojas-Carulla et al. [2016] make the assumption that if the conditional distribution of the target given some subset of covariates is invariant across different training settings, then this conditional distribution must also be the same in the test setting. The methods proposed in [Schölkopf et al., 2012, Zhang et al., 2013, 2015, Gong et al., 2016] all address challenging settings in which conditional independences that follow from the usual Markov and faithfulness assumptions alone do

<sup>1</sup>More precisely, we should say that  $\mathbb{P}(Y | Z, I_1 = 1)$  may differ from  $\mathbb{P}(Y | Z, I_1 = 0)$ , and similarly when conditioning on  $\{A, Z\}$ .

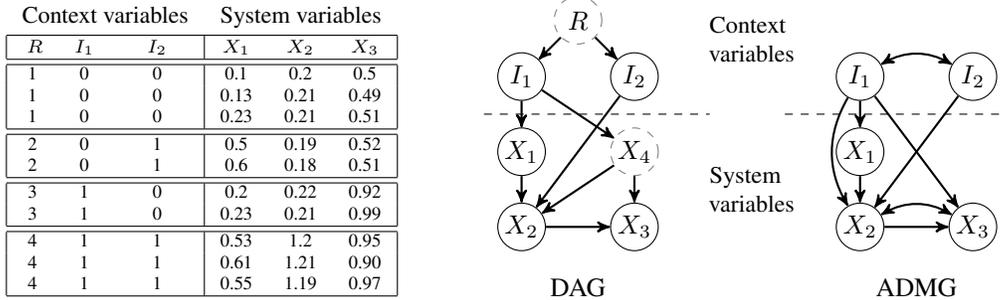


Figure 2: Example of JCI setting with four data sets. In this example,  $R = 1$  corresponds with the observational setting, whereas  $R = 2, 3, 4$  correspond with different interventional settings. Intervention variables  $I_1$  and  $I_2$  may indicate two different gene knockouts; system variables  $X_1, X_2, X_3$  could measure different blood cell phenotypes (e.g., red blood cell count, mean cell volume, hematocrit). Left: table with pooled data. Middle: causal DAG modeling the data generating process, with latent system variable  $X_4$ . Right: causal ADMG representation of the data generating process on observed variables, treating the regime variable as latent to avoid faithfulness violations induced by deterministic relations between context variables (e.g., in this particular experimental design,  $I_1$  and  $I_2$  are both deterministic functions of  $R$ ).

not suffice to solve the problem, but additional assumptions on the data generating process have to be made.

In this work, we will make no such additional assumptions, and address the challenging setting in which both the causal graph *and* the intervention types and targets may be (partially) unknown. This is a realistic setting in many practical applications. For example, in biology, many interventions that can be performed on organisms are known to result in measurable downstream effects, but the exact mechanism and direct intervention targets are unknown, and therefore it is not clear whether the knowledge gained may be transferred to other species. In pharmaceutical research, it is desirable to target the root causes of illness directly and minimize side-effects; however, as the causal mechanisms are often poorly understood, it is unclear what exactly a drug is doing and whether the results of a particular study on a subpopulation of patients (say, middle-aged males in the US) will generalize to other subpopulations (e.g., elderly women with dementia). In policy decisions, changing tax rules may have different repercussions for different socio-economic classes, but the exact workings of an economy can only be modeled to a certain extent. Machine learning may help to make such predictions more data-driven, but should then correctly take into account the transfer of distributions that result from interventions and context changes.

Our main contributions are the following. We propose a set of relatively weak assumptions in this general setting that make the problem well-posed. We then propose a method that can solve this class of causal transfer learning problems even when there are latent confounders and when types and targets of interventions are not known. The main idea is to select the subset of features  $\mathbf{A}$  that leads to the best predictions of  $Y$  on the training data, while satisfying *invariance* (i.e.,  $\mathbb{P}(Y | \mathbf{A})$  is the same in the training and the test distribution). To test whether the invariance condition is satisfied, we build on recent advances in causal discovery [Hyttinen et al., 2014, Magliacane et al., 2017] that exploit the information provided by multiple training sets from different interventional settings. We show that our method outperforms standard feature selection on synthetic data and a real-world example.

## 2 Non-deterministic Joint Causal Inference

The class of causal transfer learning problems that are the subject of study of this work will be defined in the next section. Before doing so, we first discuss Joint Causal Inference (JCI), a framework for constraint-based causal discovery from a combination of observational and experimental data that was recently proposed by Magliacane et al. [2017], and that we build upon in the next section. JCI can be seen as a generalization of the idea behind the LCD algorithm [Cooper, 1997] to multiple variables. For clarity of exposition, we give here a simplified treatment, allowing us to ignore the technical complications of deterministic relations between variables.

JCI can be applied to one or more data sets from different experimental conditions, corresponding for example with a baseline of purely observational data of the “natural” state of the system and different perturbations of the system caused by external interventions. JCI distinguishes *system variables*  $\{X_j\}_{j \in \mathcal{X}}$  describing the system of interest, and *context variables* describing the experimental setting. The context variables consist of the *regime variable*  $R$  that simply labels the data sets, and *intervention variables*  $\{I_i\}_{i \in \mathcal{I}}$  that indicate whether or not certain interventions have been performed on the system (or, more generally, specify *how* the interventions were performed, e.g., they could be continuous variables that encode the precise dosage of a prescribed drug). These interventions are not limited to the perfect (“surgical”) interventions modeled by the do-operator of Pearl [2009] (or alternatively, by “force variables” Pearl [1993]), but can also model more general types of interventions such as mechanism changes. An example of the JCI setting is given in Figure 2.

The main assumption of JCI is that the data generating process (on both system and context variables) can be represented as an acyclic Structural Causal Model (SCM) (see e.g., [Pearl, 2009]), with structural equations:

$$\mathcal{M} : \begin{cases} R = E_R \\ I_i = g_i(R, E_i), & i \in \mathcal{I} \\ X_j = f_j(X_{\text{pa}(X_j) \cap \mathcal{X}}, I_{\text{pa}(X_j) \cap \mathcal{I}}, E_j) & j \in \mathcal{X} \end{cases} \quad (1)$$

where the exogenous (noise) variables  $\{E_k\}_{k \in \mathcal{X} \cup \mathcal{I} \cup \{R\}}$  are jointly independent, and the intervention variables  $I_i$  are (possibly noisy) functions of the regime variable  $R$ .<sup>2</sup> The regime variable captures dependencies between intervention variables in the pooled data, which is a mixture of the data sets generated in the different regimes (e.g., in the pooled data in Figure 2,  $I_1$  is dependent of  $I_2$ ). In this setting, not all system variables are necessarily observed. Let  $\mathcal{X} = \mathcal{O} \cup \mathcal{L}$  be the disjoint union of observed system variables  $\{X_j\}_{j \in \mathcal{O}}$  and latent system variables  $\{X_j\}_{j \in \mathcal{L}}$ . For example, in Figure 2,  $\mathcal{O} = \{1, 2, 3\}$  and  $\mathcal{L} = \{4\}$ .

The SCM  $\mathcal{M}$  can be represented graphically by a causal Directed Acyclic Graph (DAG). Marginalizing onto the subset of variables  $\mathbf{V} := \{X_j\}_{j \in \mathcal{O}} \cup \{I_i\}_{i \in \mathcal{I}}$  (i.e., treating the regime variable  $R$  and the latent system variables  $\{X_j\}_{j \in \mathcal{L}}$  as latent), we obtain an Acyclic Directed Mixed Graph (ADMG)  $\mathcal{G}$ , also known as Semi-Markov Causal Model (see e.g., [Pearl, 2009]). In an ADMG, directed edges represent direct causal relationships, and bidirected edges represent hidden confounders (both relative to the set of variables in the ADMG). The (causal) *Markov assumption* holds [Richardson, 2003], i.e., any d-separation  $\mathbf{A} \perp \mathbf{B} \mid \mathbf{C} [\mathcal{G}]$  between sets of random variables  $\mathbf{A}, \mathbf{B}, \mathbf{C} \subseteq \mathbf{V}$  in the ADMG  $\mathcal{G}$  implies a conditional independence  $\mathbf{A} \perp \mathbf{B} \mid \mathbf{C} [\mathbb{P}(\mathbf{V})]$  in the marginal distribution  $\mathbb{P}(\mathbf{V})$  induced by the SCM  $\mathcal{M}$ . We assume that the joint distribution  $\mathbb{P}(\mathbf{V})$  is *faithful* with respect to the ADMG  $\mathcal{G}$  (i.e., that there are no other conditional independences in the joint distribution than those implied by d-separation). In particular, this implies that we also assume that there are no deterministic relations between the variables that would lead to faithfulness violations.<sup>3</sup> We will refer to this set of assumptions as the *non-deterministic JCI setting*.<sup>4,5</sup>

The non-deterministic JCI setting is very general as it allows to treat various types of interventions in a unified way: it can deal with perfect interventions [Pearl, 2009], mechanism changes

<sup>2</sup>We also allow for stochastic interventions, extending [Magliacane et al., 2017].

<sup>3</sup>This assumption differs from the more general setting described by Magliacane et al. [2017] that allows for certain deterministic relations between context variables. Here, for simplicity we use standard d-separation and the standard faithfulness assumption rather than D-separation [Spirtes et al., 2000] and the D-faithfulness assumption used in [Magliacane et al., 2017].

<sup>4</sup>Our graphical representation bears strong similarities with influence diagrams [Dawid, 2002], but a formal difference is that we consider the intervention variables to be random variables that reflect the empirical distribution of the experimental design, whereas in influence diagrams they are interpreted as decision variables rather than random variables. The advantage of treating intervention variables as random variables is that this allows one to apply standard causal discovery techniques (designed for random variables) *jointly* on system and intervention variables.

<sup>5</sup>Our graphical representation also bears some similarities with selection diagrams [Bareinboim and Pearl, 2013], but one crucial difference is that we are modeling the *joint* distribution on the intervention and system variables, whereas a selection diagram represents the *conditional* distribution of the system variables given the intervention (“selection”) variables. Because we are modeling the joint distribution and not only the conditional one, we can apply standard causal discovery techniques directly on pooled data, something that would not be as trivial when using selection diagrams instead.

Table 1: Example of causal transfer learning problem. The task is to predict the values of  $Y = X_2$  given the observations of  $X_1, X_3, I_2$  and  $I_3$  in the test task ( $I_1 = 1$ ), given various training data sets ( $I_1 = 0$ ). The distributions of the various training sets and the test set may differ due to the interventions  $I_1, I_2$  and  $I_3$  that have been performed. All variables have been observed in all tasks, except that the target  $Y = X_2$  is unobserved in the test task ( $I_1 = 1$ ).

Task	$I_1$	$I_2$	$I_3$	$X_1$	$X_2$	$X_3$
training	0	0	0	observed	observed	observed
training	0	0	1	observed	observed	observed
training	0	1	1	observed	observed	observed
test	1	0	0	observed	<b>unobserved</b>	observed

[Tian and Pearl, 2001], soft interventions [Markowitz et al., 2005], fat-hand interventions [Eaton and Murphy, 2007], activity interventions [Mooij and Heskes, 2013], and stochastic versions of all these. Knowledge of the intervention *targets* is not necessary (but is certainly helpful), as these can be learnt from the data to some extent, similarly to how the effects of system variables can be learnt. On the other hand, for certain types of interventions (e.g., perfect interventions on known targets), JCI does not take advantage of all the available information that other methods, e.g., Hyttinen et al. [2014], Triantafillou and Tsamardinos [2015], would exploit.

Any causal discovery method that does not assume causal sufficiency can be used in the non-deterministic JCI setting. Identifiability greatly benefits from taking into account the following background knowledge:

**Assumption 1.** *The ADMG  $\mathcal{G}$  with variables  $\mathbf{V}$  (consisting of system variables  $\{X_j\}_{j \in \mathcal{O}}$  and intervention variables  $\{I_i\}_{i \in \mathcal{I}}$ ) satisfies the following constraints:*

- *no variable directly causes any intervention variable  $I_i$  with respect to  $\mathbf{V}$*   
( $\forall A \in \mathbf{V}, \forall i \in \mathcal{I} : A \rightarrow I_i \notin \mathcal{G}$ );
- *any pair of intervention variables<sup>6</sup> is confounded with respect to  $\mathbf{V}$*   
( $\forall i, j \in \mathcal{I} : I_i \leftrightarrow I_j \in \mathcal{G}$ );
- *no system variable is confounded with an intervention variable with respect to  $\mathbf{V}$*   
( $\forall j \in \mathcal{O}, \forall i \in \mathcal{I} : X_j \leftrightarrow I_i \notin \mathcal{G}$ ).

Logic-based causal discovery methods, such as [Hyttinen et al., 2014, Triantafillou and Tsamardinos, 2015, Magliacane et al., 2016], are ideally suited to exploit the background knowledge. For other methods, e.g., FCI [Spirtes et al., 2000, Zhang, 2008], incorporating all background knowledge is less straightforward and as far as we know cannot be done with off-the-shelf implementations.

### 3 Causal transfer learning

In this section, we first define the class of causal transfer learning problems of interest. Then we discuss an approach to obtain predictions under transfer based on invariance of the conditional distribution of the target variable given certain “separating” subsets of features that guarantee an upper bound on the generalization error on the test data, which can be estimated from the training data. Finally, we discuss how such separating sets of features can be identified from data even when the graph is unknown.

#### 3.1 Problem setting

An example of the causal transfer learning problems that we study here is provided in Table 1. We assume that the non-deterministic JCI setting (as described in Section 2) applies. For simplicity, we assume that we have one or more training tasks, in each of which  $I_1 = 0$ , and all variables have been observed. In addition, there is one test task, in which  $I_1 = 1$ . In the test task, all variables in  $\mathbf{V}$ , except for some target variable  $Y = X_j$  (where  $j \in \mathcal{O}$ ), have been observed. The goal is to predict the values of the target variable  $Y$  given the observations in the test task, i.e., from observations of  $\mathbf{V} \setminus \{Y\}$  in the context  $I_1 = 1$ . This is a transfer learning problem, because all training sets and the

<sup>6</sup>For simplicity, we assume here that there are no (conditional) independences between intervention variables in the experimental design.

test set may have different distributions due to the different interventions that have been performed. In order to make the problem well-posed, we strengthen the assumptions of the non-deterministic JCI setting:

**Assumption 2** (Causal transfer learning assumptions). *Let  $\mathcal{G}$  be an ADMG with variables  $\mathbf{V}$  (consisting of system variables  $\{X_j\}_{j \in \mathcal{O}}$  and intervention variables  $\{I_i\}_{i \in \mathcal{I}}$ ), and  $\mathbb{P}(\mathbf{V})$  be the distribution on  $\mathbf{V}$ . Assume that:*

- (i)  $\mathcal{G}$  satisfies Assumption 1;
- (ii) the mixture of all (training and test) distributions  $\mathbb{P}(\mathbf{V})$  is Markov and faithful w.r.t.  $\mathcal{G}$ ;
- (iii) the mixture of all training distributions,  $\mathbb{P}(\mathbf{V} \setminus I_1 | I_1 = 0)$ , is faithful w.r.t. the induced sub-ADMG  $\mathcal{G}_{\mathbf{V} \setminus \{I_1\}}$ ;<sup>7</sup>
- (iv)  $I_1$  has no direct effect on  $Y$  w.r.t.  $\mathbf{V}$ , i.e.,  $I_1 \rightarrow Y \notin \mathcal{G}$ .

These assumptions are not testable (from the data), but we believe that they are as weak as possible to make the problem well-posed and without having to introduce other (untestable) assumptions. In particular, Assumption 2(iv) gets weaker the more relevant system variables are observed.<sup>8</sup> Intuitively, Assumption 2(iii) implies that the mixture of training distributions is faithful to the same ADMG (without  $I_1$ ) as the test distribution, i.e., there are no (conditional) dependences that are present in the test setting but absent in any of the training settings. A sufficient condition for this is if the test setting does not introduce new causal relations or confounders between variables (except for  $I_1$ ).

An important consequence of that additional faithfulness assumption is that it enables us to transfer any conditional independence from the training distribution to the test distribution, via the graph  $\mathcal{G}$  (proof provided in the Supplementary Material):

**Lemma 1.** *Under Assumption 2,*

$$\mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} [I_1 = 0] \iff \mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} \cup \{I_1\} [\mathcal{G}] \iff \mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} \cup \{I_1\}$$

for subsets  $\mathbf{A}, \mathbf{B}, \mathbf{C} \subseteq \mathbf{V}$  not containing  $I_1$ .<sup>9</sup>

Note that additional *independences* may hold in the test distribution, e.g., when  $I_1$  models a perfect intervention. According to the assumptions,  $I_1$  can be any type of intervention that does not introduce new conditional dependencies between variables that are not already present in the mixture of the training distributions. This Lemma will turn out to be useful later, as we will be making heavy use of conditional independences to draw conclusions about the causal structure, but not all conditional independences can be directly tested in the data (because the values of  $Y$  are missing in context [ $I_1 = 1$ ]).

### 3.2 Causal feature selection

Our approach to addressing these causal transfer learning problems is based on finding a *separating set*  $\mathbf{A} \subseteq \mathbf{V} \setminus \{I_1, Y\}$  of (intervention and system) variables which satisfies  $I_1 \perp\!\!\!\perp Y | \mathbf{A} [\mathcal{G}]$ . If such a separating set  $\mathbf{A}$  can be found, then the distribution of  $Y$  conditional on  $\mathbf{A}$  is *invariant* under transferring from the training tasks to the test task, i.e.,  $\mathbb{P}(Y | \mathbf{A}, I_1 = 0) = \mathbb{P}(Y | \mathbf{A}, I_1 = 1)$ . As the former can be estimated from the training data, we directly obtain a prediction for the latter, which then enables us to predict the values of  $Y$  from the observed values of  $\mathbf{A}$  in the test set.<sup>10</sup>

For simplicity of the exposition, we use the squared loss function and we ignore finite-sample issues. When predicting  $Y$  from a subset of features  $\mathbf{A} \subseteq \mathbf{V} \setminus \{Y, I_1\}$ , the optimal predictor is defined as the function  $\hat{Y}$  mapping the domain of  $\mathbf{A}$  to the domain of  $Y$  that minimizes the expected *test risk*  $\mathbb{E}((Y - \hat{Y})^2 | I_1 = 1)$ , and is given by the conditional expectation (regression function)

<sup>7</sup>For an ADMG  $\mathcal{G}$  defined on variables  $\mathbf{V}$ , and any subset  $\mathbf{W} \subseteq \mathbf{V}$ , the induced sub-ADMG  $\mathcal{G}_{\mathbf{W}}$  is the ADMG with  $\mathbf{W}$  as nodes and all (directed and bidirected) edges that exist in  $\mathcal{G}$  between pairs of nodes in  $\mathbf{W}$ . Note that this can be different from the marginalized ADMG on  $\mathbf{W}$ .

<sup>8</sup>For some proposals on what to do when this assumption is violated, see e.g., [Schölkopf et al., 2012, Zhang et al., 2013, 2015, Gong et al., 2016].

<sup>9</sup>Here, with  $\mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} [I_1 = 0]$  we mean  $\mathbf{A} \perp\!\!\!\perp \mathbf{B} | \mathbf{C} [\mathbb{P}(\mathbf{V} | I_1 = 0)]$ , i.e., the conditional independence of  $\mathbf{A}$  from  $\mathbf{B}$  given  $\mathbf{C}$  in the mixture of the training distributions  $\mathbb{P}(\mathbf{V} | I_1 = 0)$ .

<sup>10</sup>This trivial observation is not novel; see e.g. [Ch. 7, p. 164, Spirtes et al., 2000]. It also follows as a special case of [Theorem 2, Pearl and Bareinboim, 2011]. The main novelty of this work is the proposed strategy to identify such separating sets.

$\hat{Y}_{\mathbf{A}}^1(\mathbf{a}) := \mathbb{E}(Y | \mathbf{A} = \mathbf{a}, I_1 = 1)$ . Since  $Y$  is not observed in the test data, we cannot directly estimate this regression function from the data.

One approach that is often used in practice is to ignore the difference in distribution between training and test data, and use instead the predictor  $\hat{Y}_{\mathbf{A}}^0(\mathbf{a}) := \mathbb{E}(Y | \mathbf{A} = \mathbf{a}, I_1 = 0)$  which minimizes the expected *training risk*  $\mathbb{E}((Y - \hat{Y})^2 | I_1 = 0)$ . This approximation introduces a bias  $\hat{Y}_{\mathbf{A}}^1 - \hat{Y}_{\mathbf{A}}^0$  that we will refer to as the *transfer bias* (when predicting  $Y$  from  $\mathbf{A}$ ). When ignoring that training and test set come from different distributions, any standard machine learning method can be used to predict  $Y$  from  $\mathbf{A}$ . As the transfer bias can become arbitrarily large (as we have seen in Example 1), the prediction accuracy by this solution strategy may be arbitrarily bad (even in the infinite-sample limit).

Instead, we propose to only predict  $Y$  from  $\mathbf{A}$  when the set  $\mathbf{A}$  of features satisfies the following *separating set* property:

$$I_1 \perp\!\!\!\perp Y | \mathbf{A} [\mathcal{G}], \quad (2)$$

i.e., it d-separates  $I_1$  from  $Y$  in  $\mathcal{G}$ . By the causal Markov property, this implies  $I_1 \perp\!\!\!\perp Y | \mathbf{A} [\mathbb{P}(\mathbf{V})]$ . In other words, for separating sets, the distribution of  $Y$  conditional on  $\mathbf{A}$  is *invariant* under transferring from the training tasks to the test task, i.e.,  $\mathbb{P}(Y | \mathbf{A}, I_1 = 0) = \mathbb{P}(Y | \mathbf{A}, I_1 = 1)$ . By virtue of the invariance, regression functions are identical for the training and test distributions, i.e.,  $\hat{Y}_{\mathbf{A}}^0 = \hat{Y}_{\mathbf{A}}^1$ , and hence also the expected training and test risks when using  $\hat{Y}_{\mathbf{A}}^0$  are identical:

$$I_1 \perp\!\!\!\perp Y | \mathbf{A} [\mathcal{G}] \implies \mathbb{E}((Y - \hat{Y}_{\mathbf{A}}^0)^2 | I_1 = 1) = \mathbb{E}((Y - \hat{Y}_{\mathbf{A}}^0)^2 | I_1 = 0). \quad (3)$$

The r.h.s. can be estimated from the training data, and the l.h.s. equals the generalization error on the test data when using the predictor  $\hat{Y}_{\mathbf{A}}^0$  trained on the training data (which equals the predictor  $\hat{Y}_{\mathbf{A}}^1$  that one could obtain if all test data were observed).<sup>11</sup> Although this approach leads to zero transfer bias, it introduces another bias: by using only a subset of the features  $\mathbf{A}$ , rather than *all available* features  $\mathbf{V} \setminus \{I_1, Y\}$ , we may miss relevant information to predict  $Y$ . We refer to this bias as the *incomplete information bias*,  $\hat{Y}_{\mathbf{V} \setminus \{Y, I_1\}}^1 - \hat{Y}_{\mathbf{A}}^1$ .

The total bias when using  $\hat{Y}_{\mathbf{A}}^0$  to predict  $Y$  is the sum of the transfer bias and the incomplete information bias:

$$\underbrace{\hat{Y}_{\mathbf{V} \setminus \{Y, I_1\}}^1 - \hat{Y}_{\mathbf{A}}^0}_{\text{total bias}} = \underbrace{(\hat{Y}_{\mathbf{A}}^1 - \hat{Y}_{\mathbf{A}}^0)}_{\text{transfer bias}} + \underbrace{(\hat{Y}_{\mathbf{V} \setminus \{Y, I_1\}}^1 - \hat{Y}_{\mathbf{A}}^1)}_{\text{incomplete information bias}}.$$

For some problems, one may be better off to simply ignore the transfer bias and minimize the incomplete information bias, while for other problems, it is crucial to take the transfer into account to obtain small prediction errors. In that situation, we could use any subset  $\mathbf{A}$  for prediction that satisfies the separating set property (2), implying zero transfer bias; obviously, the best predictions are then obtained by selecting a separating subset that also minimizes the expected training risk (i.e., minimizes the incomplete information bias). We term this solution strategy *causal feature selection*. We conclude that this strategy of selecting a subset  $\mathbf{A}$  to predict  $Y$  may yield an asymptotic guarantee on the prediction error by (3), whereas simply ignoring the shift in distribution may lead to unbounded prediction error since the transfer bias could be arbitrarily large.

### 3.3 Identifiability of separating sets

For the causal feature selection strategy discussed in Section 3.2, we need to find one or more sets  $\mathbf{A}$  that satisfy (2).<sup>12</sup> When the causal graph  $\mathcal{G}$  is known, it is easy to read off (2) directly using d-separation. Here we address the more challenging setting in which the ADMG is (partially) unknown, and even the targets of the interventions may be (partially) unknown.<sup>13</sup>

<sup>11</sup>Note that this equation only holds asymptotically; for finite samples, in addition to the transfer from training to test distributions, we have to deal with the generalization from empirical to population distributions and from the covariate shift if  $\mathbb{P}(\mathbf{A} | I_1 = 1) \neq \mathbb{P}(\mathbf{A} | I_1 = 0)$  [see e.g. Mansour et al., 2009].

<sup>12</sup>Any set  $\mathbf{A}$  that makes  $I_1$  independent of  $Y$  would already suffice (even in case of faithfulness violations), but the problem is that we cannot directly test in the data whether  $Y \perp\!\!\!\perp I_1 | \mathbf{A}$ , because the values of  $Y$  are missing for  $I_1 = 1$ .

<sup>13</sup>Another option, proposed by Rojas-Carulla et al. [2016], would be to *assume* that if  $p(Y | \mathbf{A})$  is invariant across all training data (i.e.,  $p(Y | \mathbf{A}, I_1 = 0, I_{\setminus 1} = i') = p(Y | \mathbf{A}, I_1 = 1, I_{\setminus 1} = i)$  for all  $i, i'$ ), then the same holds including the test data (i.e.,  $p(Y | \mathbf{A}, I_1 = 1, I_{\setminus 1} = i) = p(Y | \mathbf{A}, I_1 = 0, I_{\setminus 1} = i)$ ). This is a stronger assumption than the ones we are making here, and Example 2 is a simple case in which it is violated.

The existence of a separating set  $\mathbf{A}$  is not guaranteed in general. The following example (proof provided in the Supplementary Material) illustrates a case in which a separating set  $\mathbf{A} = \{A\}$  is actually identifiable.

**Example 2.** Assume that Assumption 2 holds for two intervention variables  $I_1, I_2$  and three system variables  $A, Y, Z$ . If the following conditional (in)dependencies all hold in the training data:

$$I_2 \perp\!\!\!\perp Y \mid A [I_1 = 0], \quad I_2 \not\perp\!\!\!\perp Y \mid \emptyset [I_1 = 0], \quad I_2 \perp\!\!\!\perp Z \mid Y [I_1 = 0],$$

then  $I_1 \perp\!\!\!\perp Y \mid A [\mathcal{G}]$ , i.e.,  $\{A\}$  is a separating set. One possible ADMG leading to those (in)dependencies is provided in Figure 1d. For that ADMG, and given enough data, feature selection applied to the training data will generically select  $\{A, Z\}$  as the optimal set of features for predicting  $Y$ , which can lead to an arbitrarily large prediction error. On the other hand, using the separating set  $\{A\}$  to predict  $Y$  is valid in the sense that it leads to zero transfer bias, and therefore provides a guarantee on the expected test risk (i.e., it provides an upper bound on the optimal expected test risk, which can actually be estimated from the training data).

Rather than characterising by hand all possible situations in which a separating set can be identified (like in Example 2), in this work we delegate the causal reasoning to an automatic theorem prover. Intuitively, the idea is to provide the automatic theorem prover with the conditional (in)dependencies that hold in the data, in combination with an encoding of the causal transfer learning assumptions into logical rules, and ask the theorem prover whether it can prove that  $I_1 \perp\!\!\!\perp Y \mid \mathbf{A}$  holds for a candidate set  $\mathbf{A}$  from the assumptions and provided conditional (in)dependencies. There are three possibilities: either it can prove the query (and then we can proceed to predict  $Y$  from  $\mathbf{A}$  and get an estimate of the expected test risk), or it can disprove the query (and then we know  $\mathbf{A}$  will generically give predictions that suffer from an arbitrarily large transfer bias), or it can do neither (in which case hopefully another subset  $\mathbf{A}$  can be found that does provably satisfy (2)).

### 3.4 Implementation details

Possibly the simplest (brute-force) causal feature selection scheme that one could think of works as follows: by using a standard feature selection method, produce a ranked list of subsets  $\mathbf{A} \subseteq \mathbf{V} \setminus \{Y, I_1\}$ , ordered ascendingly with respect to the empirical training risk. Going through this list of subsets (starting with the one with the smallest empirical training risk), test whether the separating set property can be inferred from the data by querying the automated theorem prover. If (2) can be shown to hold, use that subset  $\mathbf{A}$  for prediction of  $Y$  and stop; if not, continue with the next candidate subset  $\mathbf{A}$  in the list. If no subset satisfies (2), abstain from making a prediction.<sup>14</sup>

To test the separating set condition (2), we use the approach proposed by Hyttinen et al. [2014], where we simply add the background knowledge (Assumption 1) concerning the non-deterministic JCI setting as constraints on the optimization problem, in addition to the transfer-learning specific assumption that  $I_1 \rightarrow Y \notin \mathcal{G}$  (Assumption 2(iv)). As inputs we use all directly testable conditional independence test p-values  $p_{\mathbf{A} \perp\!\!\!\perp \mathbf{B} \mid \mathbf{C}}$  in the pooled data (when  $Y \notin \mathbf{A} \cup \mathbf{B} \cup \mathbf{C}$ ) and all those resulting from Lemma 1 from the training data only (if  $Y \in \mathbf{A} \cup \mathbf{B} \cup \mathbf{C}$ ). If background knowledge on intervention targets or the causal graph is available, it can easily be added as well. We use the method proposed in Magliacane et al. [2016] to query for the confidence of whether some statement (e.g.,  $Y \perp\!\!\!\perp I_1 \mid \mathbf{A}$ ) is true or false. The theory in Magliacane et al. [2016] shows that this approach is sound under oracle inputs, and asymptotically consistent whenever the statistical conditional independence tests used are asymptotically consistent. In other words, in this way the probability of wrongly deciding whether a subset  $\mathbf{A}$  is a separating set converges to zero as the sample size increases. We chose this approach because it is easy to implement on top of existing open source code.<sup>15</sup> The computational cost quickly increases with the number of variables, limiting the number of variables that can be considered simultaneously.

One important issue is how to predict  $Y$  when an optimal set  $\mathbf{A}$  has been found. As the distribution of  $\mathbf{A}$  may shift when transferring from training to test task, this means that there is a *covariate shift*

<sup>14</sup>Abstaining from predictions can be advantageous when trading off recall and accuracy. If a prediction *has* to be made, we can fall back on some other method or simply accept the risk that the transfer bias may be large.

<sup>15</sup>We build on the source code provided by Magliacane et al. [2016] which in turn extends the source code provided by Hyttinen et al. [2014]. The full source code of our implementation and the experiments will be made available under an open source license on publication.

to be taken into account when predicting  $Y$ . Any method (e.g., least-squares regression) could in principle be used to predict  $Y$  from a given set of covariates, but it is advisable to use a prediction method that works well under covariate shift, e.g., [Sugiyama et al., 2008].

## 4 Evaluation

We perform an evaluation on both synthetic data and a real-world dataset based on a causal inference challenge.<sup>16</sup> The latter dataset consists of hematology-related measurements from the International Mouse Phenotyping Consortium (IMPC), which collects measurements of phenotypes of mice with different single-gene knockouts.

In both evaluations we compare a standard feature selection method (which uses Random Forests) with our method that builds on top of it and selects from its output the best separating set. First, we score all possible subsets of features by their out-of-bag score using the implementation of Random Forest Regressor from `scikit-learn` [Pedregosa et al., 2011] with default parameters. For the baseline we then select the best performing subset and predict  $Y$ . Instead, for our causal feature selection method we try to find a subset of features  $A$  that is also a separating set, starting from the subsets with the best scores. To test whether  $A$  is a separating set, we use the method described in Section 3.4, using the ASP solver `clingo 4.5.4` [Gebser et al., 2014]. We provide as inputs the independence test results from a partial correlation test with significance level  $\alpha = 0.05$  and combine it with the weighting scheme from [Magliacane et al., 2016]. We then use the first subset  $A$  in the ranked list of predictive sets of features found by the Random Forest method for which the confidence that  $I_1 \perp Y \mid A$  holds is positive. If there is no set  $A$  that satisfies this criterium, then we abstain from making a prediction.

For the synthetic data, we generate randomly 1000 linear acyclic models with latent variables and Gaussian noise, each with three system variables, and sample 1000 data points each for the observational and two experimental cases, where we assume soft interventions on randomly selected targets. We randomly select which intervention variable will be  $I_1$  and which system variable will be  $Y$ . We disallow direct effects of  $I_1$  on  $Y$ , and enforce that no intervention can directly affect all variables simultaneously. Figure 3a shows a boxplot of the  $L_2$  loss of the predicted  $Y$  values with respect to the true values for both the baseline and our causal feature selection method. The latter performs substantially better than the baseline, and abstains from making a prediction in 430 cases out of 1000.

For the real-world dataset, we select a subset of the variables considered in the CRM Causal Inference Challenge. Specifically, for simplicity we focus on 16 phenotypes that are not deterministically related to each other. The dataset contains measurements for 441 “wild type” mice and for about 10 “mutant” mice for each of 13 different single gene knockouts. We then generate 1000 datasets by randomly selecting subsets of 3 variables and 2 gene knockouts interventions, and always include also “wild type” mice. For each dataset we randomly choose  $Y$  and  $I_1$ , and remove the values of  $Y$  for  $I_1 = 1$ . Figure 3b shows a boxplot of the  $L_2$  loss of the predicted  $Y$  values with respect to the real values for the baseline and our method. Given the small size of the datasets, this is a very difficult problem. Even in this case, our method performs better than the baseline, and abstains from making a prediction for 203 cases out of 1000.

## 5 Discussion and conclusion

We have defined a general class of causal transfer learning problems and proposed a method that can identify sets of features that lead to transferrable predictions in practice. Our assumptions are very general and do not require the causal graph or the intervention targets to be known. The method gives promising results on simulated and real-world data. More work remains to be done on the implementation side, for example, scaling up to more variables. We hope that this work will also inspire further research on the interplay between bias, variance and causality from a statistical learning theory perspective.

---

<sup>16</sup>Part of the CRM workshop on Statistical Causal Inference and Applications to Genetics, Montreal, Canada (2016). See also [http://www.crm.umontreal.ca/2016/Genetics16/competition\\_e.php](http://www.crm.umontreal.ca/2016/Genetics16/competition_e.php)

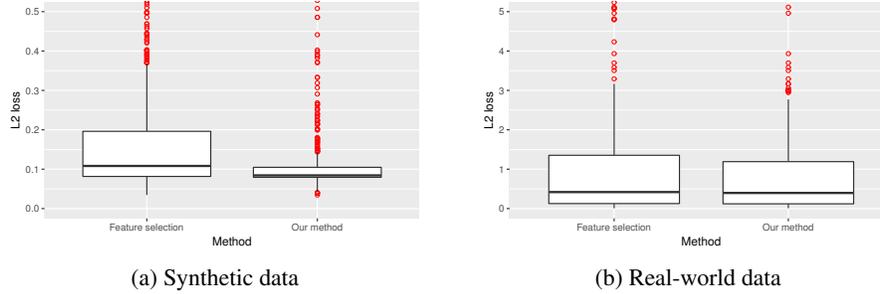


Figure 3: Evaluation results (see main text for details).

## Acknowledgments

We thank Patrick Forré for proofreading a draft of this work. We thank Renée van Amerongen and Lucas van Eijk for sharing their domain knowledge about the hematology-related measurements from the International Mouse Phenotyping Consortium (IMPC). SM, TC, SB, and PV were supported by NWO, the Netherlands Organization for Scientific Research (VIDI grant 639.072.410). SM was also supported by the Dutch programme COMMIT/ under the Data2Semantics project. TC was also supported by NWO grant 612.001.202 (MoCoCaDi), and EU-FP7 grant agreement n.603016 (MATRICS). TvO and JMM were supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement 639466).

## References

- E. Bareinboim and J. Pearl. A general algorithm for deciding transportability of experimental results. *Journal of Causal Inference*, 1:107–134, 2013.
- E. Bareinboim and J. Pearl. Causal inference and the data-fusion problem. *Proceedings of the National Academy of Sciences*, 113(27):7345–7352, 2016.
- G. F. Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Mining and Knowledge Discovery*, 1(2):203–224, 1997.
- P. Dawid. Influence diagrams for causal modelling and inference. *International Statistical Review*, 70(2): 161–189, 2002.
- D. Eaton and K. Murphy. Exact Bayesian structure learning from uncertain interventions. In *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics, (AISTATS-07)*, volume 2 of *Proceedings of Machine Learning Research*, pages 107–114, 2007.
- M. Gebser, R. Kaminski, B. Kaufmann, and T. Schaub. *Clingo = ASP + control: Extended report*. Technical report, University of Potsdam, 2014. URL <http://www.cs.uni-potsdam.de/wv/pdfformat/gekakasc14a.pdf>.
- M. Gong, K. Zhang, T. Liu, D. Tao, C. Glymour, and B. Schölkopf. Domain adaptation with conditional transferable components. In *Proceedings of the 33rd International Conference on Machine Learning (ICML 2016)*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 2839–2848, 2016.
- A. Hyttinen, F. Eberhardt, and M. Järvisalo. Constraint-based causal discovery: Conflict resolution with answer set programming. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, (UAI-14)*, pages 340–349, 2014.
- A. Hyttinen, F. Eberhardt, and M. Järvisalo. Do-calculus when the true graph is unknown. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence (UAI 2015)*, pages 395–404, 2015.
- S. Magliacane, T. Claassen, and J. M. Mooij. Ancestral causal inference. In *In Proceedings of Advances in Neural Information Processing Systems, (NIPS-16)*, pages 4466–4474, 2016.
- S. Magliacane, T. Claassen, and J. M. Mooij. Joint causal inference from observational and interventional datasets. *arXiv.org preprint*, arXiv:1611.10351v2 [cs.LG], Mar. 2017. URL <http://arxiv.org/abs/1611.10351>.
- Y. Mansour, M. Mohri, and A. Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv.org preprint*, arXiv:0902.3430v2 [cs.LG], 2009. URL <https://arxiv.org/pdf/0902.3430>.

- F. Markowetz, S. Grossmann, and R. Spang. Probabilistic soft interventions in conditional Gaussian networks. In *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics, (AISTATS-05)*, pages 214–221, 2005.
- J. M. Mooij and T. Heskes. Cyclic causal discovery from continuous equilibrium data. In *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 431–439, 2013.
- S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, Oct. 2010.
- J. Pearl. Comment: Graphical models, causality, and intervention. *Statistical Science*, 8:266–269, 1993.
- J. Pearl. *Causality: models, reasoning and inference*. Cambridge University Press, 2009.
- J. Pearl and E. Bareinboim. Transportability of causal and statistical relations: A formal approach. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 247–254, 2011.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- J. Quiñero-Candela, M. Suyiyama, A. Schwaighofer, and N. D. Lawrence, editors. *Dataset Shift in Machine Learning*. MIT Press, 2009.
- T. Richardson. Markov properties for acyclic directed mixed graphs. *Scandinavian Journal of Statistics*, 30: 145–157, 2003.
- M. Rojas-Carulla, B. Schölkopf, R. Turner, and J. Peters. Causal transfer in machine learning. *arXiv.org preprint*, arXiv:1507.05333v3 [stat.ML], Aug. 2016. URL <http://arxiv.org/abs/1507.05333v3>.
- B. Schölkopf, D. Janzing, J. Peters, E. Sgouritsa, K. Zhang, and J. M. Mooij. On causal and anticausal learning. In *Proceedings of the 29th International Conference on Machine Learning (ICML 2012)*, pages 1255–1262, 2012.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT press, 2nd edition, 2000.
- A. Storkey. When training and test sets are different: characterizing learning transfer. In *Dataset Shift in Machine Learning*, chapter 1, pages 3–28. MIT Press, 2009.
- M. Sugiyama, S. Nakajima, H. Kashima, P. V. Buenau, and M. Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *In Proceedings of Advances in Neural Information Processing Systems (NIPS-08)*, pages 1433–1440, 2008.
- J. Tian and J. Pearl. Causal discovery from changes. In *Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence, (UAI-01)*, 2001.
- S. Triantafillou and I. Tsamardinos. Constraint-based causal discovery from multiple interventions over overlapping variable sets. *Journal of Machine Learning Research*, 16:2147–2205, 2015.
- J. Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17):1873–1896, 2008.
- 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*, volume 28 of *Proceedings of Machine Learning Research*, pages 819–827, 2013.
- K. Zhang, M. Gong, and B. Schölkopf. Multi-source domain adaptation: A causal view. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 3150–3157, 2015.

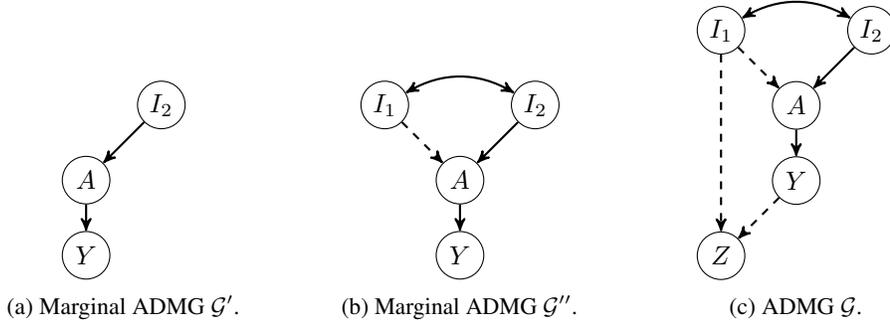


Figure 4: ADMGs for proof of Example 2. Each dashed edge can either be present or absent.

## Supplementary material

### 5.1 Proofs

**Proof of Lemma 1.** If the conditional independence  $A \perp\!\!\!\perp B \mid C [I_1 = 0]$  holds in the mixture of the training distributions, then by the additional faithfulness assumption,  $A \perp\!\!\!\perp B \mid C [\mathcal{G}_{\mathbf{V} \setminus \{I_1\}}]$ . Since by the JCI assumptions,  $I_1$  must be a noncollider on any path in  $\mathcal{G}$  that contains it, this implies  $A \perp\!\!\!\perp B \mid C \cup \{I_1\} [\mathcal{G}]$ , and therefore by the Markov property,  $A \perp\!\!\!\perp B \mid C \cup \{I_1\} [\mathbb{P}(\mathbf{V})]$ . On the other hand, if  $A \not\perp\!\!\!\perp B \mid C [I_1 = 0]$ , then by definition of independence,  $A \not\perp\!\!\!\perp B \mid C \cup \{I_1\}$ .  $\square$

**Proof of Example 2.** In the JCI setting, we assume that in the full ADMG  $\mathcal{G}$  over variables  $\{I_1, I_2, A, Y, Z\}$ ,  $I_1$  and  $I_2$  are confounded (by the latent regime variable  $R$ ) and not caused by system variables  $A, Y, Z$ . Furthermore, no pair of system variable and intervention variable is confounded.

In the context  $[I_1 = 0]$ , if the conditional independences  $I_2 \perp\!\!\!\perp Y \mid A [I_1 = 0]$  and  $I_2 \not\perp\!\!\!\perp Y \mid \emptyset [I_1 = 0]$  hold, then we can also derive that  $I_2 \not\perp\!\!\!\perp A \mid \emptyset [I_1 = 0]$ , for example using Rule (9) from Magliacane et al. [2016]. Moreover, we know that  $I_2$  is not caused by  $A$  and  $Y$ , or in other words  $A \not\rightarrow I_2$  and  $Y \not\rightarrow I_2$ . Thus we conclude that  $(I_2, A, Y)$  is an LCD triple Cooper [1997] in the context  $I_1 = 0$ . Since in addition, in this case  $I_2$  and  $A$  are unconfounded, the marginal ADMG  $\mathcal{G}'$  on  $\{I_2, A, Y\}$  (in the context  $I_1 = 0$ , and hence by Lemma 1 in all contexts) must be given by Figure 4a.

Therefore, the extended marginal ADMG  $\mathcal{G}''$  on variables  $\{I_1, I_2, A, Y\}$  must also have a directed path from  $I_2$  to  $A$  and from  $A$  to  $Y$ .  $I_1$  cannot be on these paths, as none of the variables causes  $I_1$ , and therefore  $\mathcal{G}''$  also contains the directed edges  $I_2 \rightarrow A$  and  $A \rightarrow Y$ . Moreover,  $\mathcal{G}''$  cannot contain any edge between  $I_2$  and  $Y$ , nor a bidirected edge between  $A$  and  $Y$ , because that would violate the conditional independence. By construction, in the JCI setting there is a bidirected edge between  $I_1$  and  $I_2$ , and that is the only bidirected edge connecting to  $I_1$  or  $I_2$ . As we assumed there is no direct effect of  $I_1$  on target  $Y$ , there is no edge between  $I_1$  and  $Y$  in  $\mathcal{G}''$ . There is also no directed edge  $A \rightarrow I_1$  in  $\mathcal{G}''$ , as the JCI assumption implies none of the other variables causes  $I_1$ . Therefore, the marginal ADMG  $\mathcal{G}''$  is given by Figure 4b. either with the directed edge  $I_1 \rightarrow A$  present, or without that edge.

If it additionally holds that  $I_2 \perp\!\!\!\perp Z \mid Y [I_1 = 0]$ , we have two possibilities:

1. if  $I_2 \perp\!\!\!\perp Z \mid \emptyset [I_1 = 0]$  holds, then  $Z$  is not caused by  $I_2$ . This means it cannot be on any directed path from  $I_2$  to  $A$ , from  $A$  to  $Y$ , or be a descendant of  $Y$ . Therefore the full ADMG  $\mathcal{G}$  also necessarily contains the directed edges  $I_2 \rightarrow A$  and  $A \rightarrow Y$ .
2. if  $I_2 \not\perp\!\!\!\perp Z \mid \emptyset [I_1 = 0]$  holds, then in conjunction with  $I_2 \perp\!\!\!\perp Z \mid Y [I_1 = 0]$  we can derive  $Y \rightarrow Z$ , for example using Rule (5) from [Magliacane et al., 2016]. This means  $Z$  must be a descendant of  $Y$  in the full ADMG  $\mathcal{G}$ , which implies it cannot be on the directed path from  $I_2$  to  $A$ , or on the one from  $A$  to  $Y$ . Therefore the full ADMG  $\mathcal{G}$  also necessarily contains the directed edges  $I_2 \rightarrow A$  and  $A \rightarrow Y$ .

Because of the independence statements and JCI assumptions, there cannot be a bidirected edge between  $Z$  and  $A, Y, I_1$  or  $I_2$ . Similarly, there cannot be directed edges from  $Z$  to one of those nodes. The edges  $A \rightarrow Z$  and  $I_2 \rightarrow Z$  must also be absent.

In both cases, there can be a directed edge from  $I_1$  to  $Z$ . Therefore, the full ADMG  $\mathcal{G}$  is given by Figure 4c. In all cases we see that  $I_1 \perp Y \mid A [\mathcal{G}]$ , and we conclude that  $\{A\}$  is a valid separating set for causal feature selection.

If the ADMG is as in Figure 1d, then a standard feature selection method would asymptotically prefer the subset  $\{A, Z\}$  to predict  $Y$  over the subset  $\{A\}$  (note that the Markov blanket of  $Y$  in context  $[I_1 = 0]$  is  $\{A, Z\}$ ). As a result, any prediction method trained on all available features in context  $[I_1 = 0]$  may incur a possibly unbounded prediction error when used to predict  $Y$  in the test context  $[I_1 = 1]$  (for example, if  $Z$  is an almost deterministic copy of  $Y$  if  $I_1 = 0$ , but has a drastically different distribution if  $I_1 = 1$ ).  $\square$