
Causal Discovery in the Presence of Missing Data

Ruibo Tu^{*1}, Cheng Zhang^{*2}, Paul Ackermann³, Karthika Mohan⁴,
Hedvig Kjellström¹, and Kun Zhang^{*5}

¹KTH Royal Institute of Technology, ²Microsoft Research, Cambridge, ³Karolinska Institute,

⁴University of California, Berkeley, ⁵Carnegie Mellon University

Abstract

Missing data are ubiquitous in many domains such as healthcare. When these data entries are not missing completely at random, the (conditional) independence relations in the observed data may be different from those in the complete data generated by the underlying causal process. Consequently, simply applying existing causal discovery methods to the observed data may lead to wrong conclusions. In this paper, we aim at developing a causal discovery method to recover the underlying causal structure from observed data that are missing under different mechanisms, including missing completely at random (MCAR), missing at random (MAR), and missing not at random (MNAR). With missingness mechanisms represented by missingness graphs (m-graphs), we analyze conditions under which additional correction is needed to derive conditional independence/dependence relations in the complete data. Based on our analysis, we propose Missing Value PC (MVPC), which extends the PC algorithm to incorporate additional corrections. Our proposed MVPC is shown in theory to give asymptotically correct results even on data that are MAR or MNAR. Experimental results on both synthetic data and real healthcare applications illustrate that the proposed algorithm is able to find correct causal relations even in the general case of MNAR.

1 Introduction

Determining causal relations plays a pivotal role in many disciplines of science, especially in healthcare. In particular,

* Authors contributed equally.

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

understanding causality in healthcare can facilitate effective treatments to improve quality of life. Traditional approaches (Domej-Arverud et al., 2016) to identify causal relations are usually based on randomized controlled trials, which are expensive or even impossible in certain domains. In contrast, owing to the availability of purely observational data and recent technological developments in computational and statistical analysis, causal discovery from observational data is potentially widely applicable (Spirtes et al., 2001; Peters et al., 2017). In recent years, causal discovery from observational data has become popular in medical research (Sokolova et al., 2017; Klasson et al., 2017).

Most existing algorithms for causal discovery are designed for complete data (Pearl, 2000; Peters et al., 2017), such as the widely used PC algorithm (Spirtes et al., 2001). Unfortunately, missing data entries are common in many domains. For example, in healthcare, missing entries may come from imperfect data collection, compensatory medical instruments, and fitness of the patients etc. (Robins, 1986).

All missing data problems fall into one of the following three categories (Rubin, 1976): Missing Completely At Random (MCAR), Missing At Random (MAR), and Missing Not At Random (MNAR). Data are MCAR if the cause of missingness is purely random, e.g., some entries are deleted due to a random computer error. Data are MAR when the direct cause of missingness is fully observed. For example, a dataset consists of two variables: gender and income, where gender is always observed and income has missing entries. MAR missingness would occur when men are more reluctant than women to disclose their income (i.e., gender causes missingness). Data that are neither MAR nor MCAR fall under the MNAR category. In the example above, MNAR would occur when gender also has missing entries. These missingness mechanisms can be represented by causal graphs as introduced in Section 2. While it might be tempting to remove samples corrupted by missingness and perform analysis solely with complete cases, it will reduce sample size and, more importantly, bias the outcome especially when data are MAR or MNAR (Rubin, 2004; Mohan et al., 2013; Shpitser, 2016).

This paper is concerned with how to find the underlying

causal structure over observed data even in the situation of MAR or MNAR. For simplicity, we assume causal sufficiency in the paper, as assumed by many causal discovery methods including PC (Spirites et al., 2001). Recoverability of the data distribution under missing data has been discussed in a number of contributions; see, e.g., (Mohan et al., 2013; Mohan and Pearl, 2014a). A straightforward solution is to recover all relevant distributions that are needed for Conditional Independence (CI) tests involved in the CI-based causal search procedure, such as PC. But compared to the CI test on independent and identically distributed observations, the CI test on corrected distributions is generally harder because it involves simulating new data or importance reweighting with density ratios.

Therefore, instead of correcting all CI tests of the PC algorithm, we aim to find under which condition, CI tests in the observed data produce erroneous edges, and then we correct only such edges by further applying CI tests on corrected distributions. Our main contributions are:

- We provide a theoretical analysis of the error that different missingness mechanisms introduce in the results given by traditional causal discovery methods, such as the PC algorithm (Section 3). We will show that naive deletion-based method may lead to incorrect results due to the bias caused by missing data. One immediate way to extend constraint-based methods to handle the missing data issue is correcting all the involved CI tests. This approach is neither data-efficient nor computation-efficient. Therefore, we identify possible errors that different missingness mechanisms lead to in the results given by deletion-based PC. We show that one needs to correct only a small number of CI tests in order to recover the true causal structure.
- We propose a novel, correction-based extension of the PC algorithm, *Missing Value PC (MVPC)*, that handles all three types of missingness mechanisms: MCAR, MAR, and MNAR (Section 4). Based on the result from Section 3, we identify where corrections are required and propose efficient correction methods for all three types of the missingness mechanisms.
- MVPC demonstrates superior performance in different settings, including two real-life healthcare scenarios (Section 5). We first evaluate the proposed MVPC on synthetic datasets under different settings. MVPC shows clear improvement over multiple baselines. We further apply MVPC to two real-world datasets: the US Cognition study and Achilles Tendon Rupture study. The results are consistent with medical domain knowledge and demonstrate the efficacy of our method.

2 Related work

We discuss closely related works, including traditional causal discovery algorithms and approaches that deal with

missing data from a causal perspective.

Causal discovery. Causal discovery from observational data has been of great interest in various domains in the past decades (Pearl, 2000; Spirites et al., 2001). In general, causal discovery consists of constraint-based methods, score-based methods, and methods based on functional causal models. Typical constraint-based methods include the PC algorithm and Fast Causal Inference (FCI). They assume that all CI relations are entailed from the causal Markov condition, according to the faithfulness assumption, and use CI constraints in the data to recover causal structure. The PC algorithm assumes no confounders (hidden direct common causes of two variables) and outputs a Completed Partially Directed Acyclic Graph (CPDAG), which is easy to interpret and often used in biomedical applications (Neto et al., 2008; Le et al., 2016). FCI allows confounders and selection bias, and outputs a Partial Ancestral Graph (PAG). For simplicity, we use the PC algorithm in this paper, but it is straightforward to transfer our framework to other constraint-based methods. Score-based methods (e.g., Greedy Equivalence Search (Chickering, 2002)) find the best Markov equivalence class (which contains DAGs that have the same CI relations) under certain score-based criterion, such as the Bayesian Information Criterion (BIC). Causal discovery based on functional causal models benefits from the additional assumptions on the data distribution and/or the functional classes to further determine the causal direction between variables. Typical functional causal models include the linear non-Gaussian acyclic model (LiNGAM) (Shimizu et al., 2006), the post-nonlinear (PNL) causal model (Zhang and Hyvärinen, 2009), and the nonlinear additive noise model (ANM) (Peters et al., 2017).

Dealing with data with missing values from a causal perspective. Recent years have witnessed a growing interest in analysing the problem of missing data from a causal perspective. In particular, the notions of recoverability and testability have been studied by modeling the missingness process using causal graphs (called missingness graphs or m-graphs) (Mohan et al., 2013). Given a m-graph, a query (such as conditional or joint distribution and causal effects) is deemed recoverable if it can be consistently estimated (Mohan and Pearl, 2014a). Testability, on the other hand, deals with finding testable implications, i.e., claims refutable by the (missing) data distribution (Mohan and Pearl, 2014b). As for causal discovery, it aims to find the structure of variables of interest rather than the missingness. Relations of variables of interest can be testable under appropriate assumptions, although relations between variables of interest and their missingness are untestable.

In causal discovery, there are few works for the MNAR case. FCI by test-wise deletion regards the missingness procedure as a particular type of selection bias to handle the MNAR missingness (Strobl et al., 2017). It shows that

FCI combined with test-wise deletion is still sound when one aims to estimate the PAG for the variables including the effect of missingness. Data missingness is usually different from selection bias, because in the selection bias case we only have the distribution of the selected samples but no clue about the population. However, in the missing data case, we may be able to check the (conditional) independence relation between two variables given others by making use of the available data for the involved variables. In the case where the missingness mechanisms to be known, this problem is closely related to the recoverability of models with missing data. Gain and Shpitser (2018) utilize Inverse Probability Weight (IPW) for each CI test, assuming the missing data model is known, which may not be realistic in many real-life applications. When the missing data model is unknown, they choose the sparsest resulting graph considering all possible missingness structures, which is usually computationally expensive.

3 Deletion-based PC: A first proposal and its behavior

We assume that there is no confounder or selection bias relative to the set of observed variables. When the available dataset has missing values, one may apply the PC algorithm for causal discovery by performing CI tests on those records which do not have missing values for the variables involved in the tests. We term this first proposal *deletion-based PC*. In this section, we discuss the influence of missing data on the result of deletion-based PC.

Primarily, we investigate the situations where errors occur to the output of deletion-based PC due to the missingness. Firstly, we utilize m-graphs and summarize the assumptions that we need for properly dealing with missingness. We then present the aforementioned deletion-based PC algorithm. Our analysis focuses on properties of the results given by this naive extension, and provides the conditions under which the deletion-based PC produces erroneous edges.

Missingness graph. We utilize the notation of the *m-graph* (Mohan et al., 2013). A m-graph is a causal DAG $G(\mathbb{V}, \mathbf{E})$ where $\mathbb{V} = \mathbf{V} \cup \mathbf{U} \cup \mathbf{V}^* \cup \mathbf{R}$. \mathbf{U} is the set of unobservable nodes; in this paper, we assume causal sufficiency, so \mathbf{U} is an empty set. \mathbf{V} is the set of *substantive* nodes (observable nodes) containing \mathbf{V}_o and \mathbf{V}_m . $\mathbf{V}_o \subseteq \mathbf{V}$ is the set of fully observed variables, denoted by white nodes in our graphical representation. $\mathbf{V}_m \subseteq \mathbf{V}$ is the set of partially observed variables that are missing in at least one record, which is shadowed in gray. \mathbf{R} is the set of *missingness indicators* that represent the status of missingness and are responsible for the values of proxy variables \mathbf{V}^* . For example, the proxy variable $Y^* \in \mathbf{V}^*$ is introduced as an auxiliary variable for the convenience of derivation. $R_y = 1$ means that the corresponding record value of Y is missing and Y^*

corresponds to a missing entry; $R_y = 0$ indicates that the corresponding record value of Y is observed and Y^* takes the value of Y .

In this work we adopt the CI-based definitions of missingness categories as stated in (Mohan et al., 2013). We denote an independent relation in a dataset by " $\perp\!\!\!\perp$ " and d-separation in a m-graph by " $\perp\!\!\!\perp_d$ ". As shown in Figure 1, data are MCAR if $\{\mathbf{V}_m, \mathbf{V}_o\} \perp\!\!\!\perp_d \mathbf{R}$ holds in the m-graph, MAR if $\mathbf{V}_m \perp\!\!\!\perp_d \mathbf{R} \mid \mathbf{V}_o$ holds, and MNAR otherwise.

Assumptions for dealing with missingness. Apart from the assumptions for the asymptotic correctness of the PC algorithm (including the causal Markov condition, faithfulness, and no confounding or selection bias), we introduce some additional assumptions that we make use of to deal with missingness.

Assumption 1 (Missingness indicators are not causes). *No missingness indicator can be the cause of any substantive (observed) variable.*

This assumption is employed in most related work using m-graphs (Mohan et al., 2013; Mohan and Pearl, 2014a). Consequently, under this assumption, if variables of interest X and Y are not d-separated by a variable set $\mathbf{Z} \subseteq \mathbf{V} \setminus \{X, Y\}$, they are not d-separated by \mathbf{Z} together with their missingness indicators. Under the faithfulness assumption, this means that if they are conditionally independent given \mathbf{Z} together with their missingness indicators, they are conditionally independent given only \mathbf{Z} . Now the problem is that generally speaking, we cannot directly verify whether they are conditionally independent given \mathbf{Z} and their missingness variables because we do not have the records for the considered variables when their missingness indicators take value one. We then need the following assumptions.

Assumption 2 (Faithful observability). *Any conditional independence relation in the observed data also holds in the unobserved data; formally, $X \perp\!\!\!\perp Y \mid \{\mathbf{Z}, \mathbf{R}_K = \mathbf{0}\} \iff X \perp\!\!\!\perp Y \mid \{\mathbf{Z}, \mathbf{R}_K = \mathbf{1}\}$. Here \mathbf{R}_K is the missingness indicator set $\{R_x, R_y, R_z\}$. $\mathbf{R}_K = \mathbf{0}$ means all the missingness indicators in \mathbf{R}_K taking the value zero; $\mathbf{R}_K = \mathbf{1}$ means at least one missingness indicator in \mathbf{R}_K taking the value one.*

This implies $X \perp\!\!\!\perp Y \mid \{\mathbf{Z}, \mathbf{R}_K = \mathbf{0}\} \iff X \perp\!\!\!\perp Y \mid \{\mathbf{Z}, \mathbf{R}_K\}$, which means that conditional independence relations in the observed data also hold in the complete data, i.e., there is no accidental conditional independence relation caused by missingness.

Assumption 3 (No causal interactions between missingness indicators). *No missingness indicator can be a deterministic function of any other missingness indicators.*

Assumption 4 (No self-masking missingness). *Self-masking missingness refers to missingness in a variable that is caused by itself. In the m-graph this is depicted by an edge from X to R_x , for $X \in \mathbf{V}_m$ (as shown in Figure 1d). We assume that there is no such edges in the m-graph.*




Figure 1: Exemplar missingness graphs in MCAR, MAR, MNAR, and self-masking missingness. X , Y , Z , and W are random variables. In m-graphs, gray nodes are partially observed variables, and white nodes are fully observed variables. R_x , R_y , and R_w are the missingness indicators of X , Y , and W .

Assumption 3 and 4 guarantee the recoverability of a joint distribution of substantive variables, as shown in (Mohan et al., 2013). As discussed in Appendix A.1, in the linear Gaussian case, the "self-masking" only affects causal discovery results when R_x has direct causes other than X .

In the end, we assume linear Gaussian causal models in this work. Thus, one can check CI relations with the partial correlation test, a simple CI test method. Note that our proposed algorithm also works well for general situations. In the non-linear case, we can use a suitable non-linear or non-parametric one (Zhang et al., 2011).

Effect of missing data on the deletion-based PC. In the presence of missing data, the list-wise deletion PC algorithm deletes all records that have any missing value and then applies the PC algorithm to the remaining data. In contrast, the test-wise deletion PC algorithm only deletes records with missing values for variables involved in the current CI test when performing the PC algorithm (which can be seen as the PC algorithm realization of (Strobl et al., 2017)). Test-wise deletion is more data-efficient than list-wise deletion. In this paper, we focus on the Test-wise Deletion PC algorithm (TD-PC).

TD-PC gives asymptotically correct results when data are MCAR since $\{\mathbf{V}_m, \mathbf{V}_o\} \perp\!\!\!\perp_d \mathbf{R}$ is satisfied. Consider Figure 1a as an example. $R_y \perp\!\!\!\perp_d \{X, Y, Z\}$ holds; thus, we have $X \perp\!\!\!\perp_d Y \mid Z \iff X \perp\!\!\!\perp_d Y \mid \{Z, R_y\}$. With the faithfulness assumption on m-graphs, $X \perp\!\!\!\perp Y \mid Z \iff X \perp\!\!\!\perp Y \mid \{Z, R_y\}$. Furthermore, with the faithful observability assumption, we conclude $X \perp\!\!\!\perp Y \mid Z \iff X \perp\!\!\!\perp Y^* \mid \{Z, R_y = 0\}$. When applying the CI test to the test-wise deleted data of concerned variables X , Y , and Z , we test whether $X \perp\!\!\!\perp Y^* \mid \{Z, R_y = 0\}$ holds. Therefore, CI results imply d-separation/d-connection relations of concerned variables in m-graphs when data are MCAR, which guarantees the asymptotic correctness of TD-PC.

In cases of MAR and MNAR, TD-PC may produce erroneous edges because $\{\mathbf{V}_m, \mathbf{V}_o\} \perp\!\!\!\perp_d \mathbf{R}$ does not hold. Therefore, in what follows in this section, we mainly address the problems of TD-PC in cases of MAR and MNAR.

Erroneous edges produced by TD-PC. Since TD-PC may produce erroneous edges when data are MAR and MNAR, in the following propositions (proofs are given in Appendix A.2), we first show that the causal skeleton (undirected graph) given by TD-PC has no missing edges, but may contain extraneous edges. We then determine the conditions under which extraneous edges occur in the output of TD-PC.

Proposition 1. *Under Assumptions 1-4 the CI relation in test-wise deleted data, $X \perp\!\!\!\perp Y \mid \{Z, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}\}$, implies the CI relation in complete data, $X \perp\!\!\!\perp Y \mid Z$, where X and Y are random variables and $Z \subseteq \mathbf{V} \setminus \{X, Y\}$.*

Proposition 1 shows that CI relations in test-wise deleted data implies the true corresponding d-separation relations in a m-graph. However, dependence relations in test-wise deleted data may imply the wrong corresponding relations in the m-graph because $X \not\perp\!\!\!\perp Y \mid \{Z, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}\} \not\Rightarrow X \not\perp\!\!\!\perp Y \mid Z$. In other words, TD-PC may wrongly treat some d-separation relations of concerned variables as to be not d-separated in a m-graph. Thus, TD-PC produces extraneous edges in the causal skeleton result rather than missing edges. For example, in Figure 1b, we have $X \not\perp\!\!\!\perp Y^* \mid \{Z, R_y = 0\}$ in the test-wise deleted data, but the true d-separation relation is $X \perp\!\!\!\perp_d Y \mid Z$ instead of $X \not\perp\!\!\!\perp_d Y \mid Z$. Thus, TD-PC produces an extraneous edge between X and Y . Fortunately, such extraneous edges appear only under special circumstances, as shown in the following proposition.

Proposition 2. *Suppose that X and Y are not adjacent in the true causal graph and that for any variable set $Z \subseteq \mathbf{V} \setminus \{X, Y\}$ such that $X \perp\!\!\!\perp Y \mid Z$, it is always the case that $X \not\perp\!\!\!\perp Y \mid \{Z, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}\}$. Then under Assumptions 1-4 for at least one variable in $\{X\} \cup \{Y\} \cup Z$, its missingness indicator is either the direct common effect or a descendant of the direct common effect of X and Y .*

Proposition 2 indicates that extraneous edges can be identified from the output of TD-PC. For example, in Figure 1b and Figure 1c, W is the direct common effect of X and Y and the missingness indicator R_y is a descendant of W . Thus, the extraneous edge occurs between X and Y in the causal skeleton produced by TD-PC.

4 Proposed method: Missing-value PC

In this section, we present our proposed approach, Missing-Value PC (MVPC), for causal discovery in the presence of missing data based on PC. We introduce the general MVPC framework in Section 4.1 and present our correction methods for removing extraneous edges in Section 4.2.

4.1 Overview of MVPC

Algorithm 1 summarizes the framework of MVPC. We perform TD-PC on \mathbf{V} (Step 1), and then involve \mathbf{R} (Step 2). This is equivalent to performing TD-PC on $\mathbf{V} \cup \mathbf{R}$ under Assumption 1, 3, and 4. More details of Step 2 are introduced in Appendix A.3. We then identify potential extraneous edges (Step 3). These are the edges between variables of which direct common effects are missingness indicators or ancestors of missingness indicators. Since we do not have orientation information at this stage, we cannot directly locate such extra edges; however, we can find potentially incorrect edges, as a superset of the incorrect edges. Next, we perform correction for these candidate edges (Step 4). Finally, we orient edges of the recovered causal skeleton with the same procedure as the PC algorithm.

Algorithm 1 Missing-value PC

- 1: *Skeleton search with deletion-based PC:*
 - a *Graph initialization:* Build a complete undirected graph G on the node set \mathbf{V} .
 - b *Causal skeleton discovery:* Remove edges in G with the same procedure as the PC algorithm (Spirtes et al., 2001) with the test-wise deleted data.
 - 2: *Detecting direct causes of missingness indicators:*
For each variable $V_i \in \mathbf{V}$ containing missing values and for each j that $j \neq i$, test the CI relation of R_i and V_j . If they are independent given a subset of $\mathbf{V} \setminus \{V_i, V_j\}$, V_j is not a direct cause of R_i .
 - 3: *Detecting potential extraneous edges:*
For each $i \neq j$, if V_i and V_j are adjacent and have at least one common adjacent variable or missingness indicator, the edge between V_i and V_j is potentially extraneous.
 - 4: *Recovering the true causal skeleton:*
Perform correction methods for removing the extraneous edges in G as shown in Section 4.2.
 - 5: *Determining the orientation:*
Orient edges in G with the same orientation procedure as the PC algorithm.
-

4.2 Recovery of the true causal skeleton

As shown in Section 3, TD-PC produces extraneous edges in the causal skeleton result in the situations of Proposition 2. In this section, we introduce our correction methods to

move the extraneous edges. We first introduce Permutation-based Correction (PermC) with an example. We then show that PermC handles most of the missingness cases. Next, we propose an alternative solution, named Density Ratio Weighted correction (DRW), for the cases which PermC does not cover.

Permutation-based correction. We use an example to demonstrate how to remove the extraneous edges with PermC. For example, suppose that we have a dataset with missing values of which the underlying m-graph is shown in Figure 1b. As discussed in Section 3, when applying TD-PC to this dataset, we produce an extraneous edge between X and Y in the output of TD-PC. The problem is that data samples from joint distribution $P(X, Y, Z)$ are not available in the observed dataset. In this case, we test the CI relations in the test-wise deleted data from $P(X, Y^*, Z | R_y = 0)$, which leads to producing the extraneous edge.

PermC solves this problem by testing the CI relations in the reconstructed virtual dataset utilizing the observed data concerning:

$$\begin{aligned} P(X, Y, Z) &= \int_W P(X, Y, Z | W)P(W)dW \\ &= \int_W P(X, Y^*, Z | W, R_y = 0)P(W)dW, \quad (1) \end{aligned}$$

such that reconstructed data follow the joint distribution $P(X, Y, Z)$. As shown in the first step of Equation 1, we introduce a random variable W which is the direct cause of R_y in Figure 1b to reconstruct the dataset and then marginalize it out. With W , the joint distribution $P(X, Y, Z)$ is estimated by 1) learning the model for $P(X, Y, Z | W)$ from test-wise deleted data, 2) plugging in the values of W in the dataset, as data samples from $P(W)$, and 3) disregarding the input W and keeping the generated virtual data for $\{X, Y, Z\}$ to marginalize W out. Given virtual data of X , Y , and Z that follow the joint distribution $P(X, Y, Z)$, one can test CI relations in the complete data.

Now the issue is that the data samples from $P(X, Y, Z | W)$ are not directly available. Nevertheless, we learn a model for $P(X, Y^*, Z | W, R_y = 0)$ to generate virtual data of X , Y , and Z from W , as shown in the second step of Equation 1. Under Assumptions 1~4 we have $P(X, Y, Z | W) = P(X, Y^*, Z | W, R_y = 0)$ because $R_y \perp\!\!\!\perp \{X, Y, Z\} | W$; moreover, data samples from $P(X, Y^*, Z | W, R_y = 0)$ can be constructed by test-wise deletion. For simplicity, under the linear Gaussian assumption we apply linear regression to learning the model for $P(X, Y^*, Z | W, R_y = 0)$ as :

$$X = \alpha_1 W + \varepsilon_1, \quad Y = \alpha_2 W + \varepsilon_2, \quad Z = \alpha_3 W + \varepsilon_3, \quad (2)$$

where α_i is the parameter of linear regression models and ε_i is the residual.

Next, we sample the input values from the probability distribution $P(W)$. Estimating $P(W)$ for sampling input values is

unnecessary in this case because we have the complete data of W which follow $P(W)$. However, to generate virtual data with linear regression models, we cannot directly input the test-wise deleted data of W and add the residuals from the linear regression models in Equation 2. In this way, the input values follow the conditional distribution $P(W|R_y=0)$ instead of $P(W)$. Thus, we shuffle the values of W in the observed dataset such that $P(W^S|R_y=0) = P(W^S)$ where W^S denotes the shuffled W . We then feed test-wise deleted values of W^S into the linear regression models as :

$$\widehat{X} := \alpha_1 W^S + \varepsilon_1, \quad \widehat{Y} := \alpha_2 W^S + \varepsilon_2, \quad \widehat{Z} := \alpha_3 W^S + \varepsilon_3, \quad (3)$$

where we denote the random variables with generated virtual values by \widehat{X} , \widehat{Y} , and \widehat{Z} . Finally, we test the CI relations among \widehat{X} , \widehat{Y} , and \widehat{Z} . PermC for this example is summarized in Algorithm 2.

Algorithm 2 Permutation-based correction

Input: data of concerned variables, such as X , Y , and Z in Figure 1b, and the direct causes of their corresponding missingness indicators, such as the direct cause W of R_y in Figure 1b

Output: The CI relations among concerned variables, such as the CI relations among X , Y , and Z .

- 1: Delete records containing any missing value. We denote the deleted dataset by D_d , and denote the original dataset by D_o .
 - 2: Regress X , Y , and Z on W with D_d as Equation 2
 - 3: Shuffle data of W in D_o , denoted by W^S , and delete records containing any missing value in D_o (included W^S).
 - 4: Generate virtual data of \widehat{X} , \widehat{Y} , and \widehat{Z} , with W^S and the residuals according to Equation 3.
 - 5: Test the CI relations among \widehat{X} , \widehat{Y} , and \widehat{Z} in the generated virtual data.
 - 6: **return** The CI relations among X , Y , and Z .
-

Without loss of generality, we summarize the conditions under which PermC correctly removes extraneous edges. Suppose that we need to test the CI relation of X and Y given $\mathbf{Z} \subseteq \mathbf{V} \setminus \{X, Y\}$ in the generated virtual data. We denote the direct causes of missingness indicators by $Pa(\mathbf{R})$. The conditions for the validity of PermC are as follows.

- (i) $\{R_x, R_y, \mathbf{R}_z, \mathbf{R}_w\} \perp\!\!\!\perp \{X, Y, \mathbf{Z}\} \mid \mathbf{W}$, where the variable set \mathbf{W} is the set of direct causes of missingness indicators R_x , R_y , and \mathbf{R}_z ; if variables in \mathbf{W} also have missing values, the direct causes of their missingness indicators \mathbf{R}_w are also included in \mathbf{W} ; formally, $\mathbf{W} = Pa(R_x, R_y, \mathbf{R}_z, \mathbf{R}_w)$;
- (ii) In the m-graph, the missingness indicators of \mathbf{W} follow the condition that $X \perp\!\!\!\perp_d Y \mid \mathbf{Z} \iff X \perp\!\!\!\perp_d Y \mid \{\mathbf{Z}, \mathbf{R}_w\}$.

Under Conditions (i) and (ii), we have

$$\begin{aligned} & P(X, Y, \mathbf{Z} \mid \mathbf{R}_w = \mathbf{0}) \\ &= \int_{\mathbf{W}^*} P(X^*, Y^*, \mathbf{Z}^* \mid \mathbf{W}^*, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}, \mathbf{R}_w = \mathbf{0}) \times \\ & \quad P(\mathbf{W}^* \mid \mathbf{R}_w = \mathbf{0}) d\mathbf{W}^*. \quad (4) \end{aligned}$$

To test the CI relation of X and Y given \mathbf{Z} in data samples from $P(X, Y, \mathbf{Z})$, it is valid to test the CI relation in the generated data samples from $P(X, Y, \mathbf{Z} \mid \mathbf{R}_w = \mathbf{0})$. Under Condition (ii) the conditional independence/dependence relations in $P(X, Y, \mathbf{Z})$ also hold in $P(X, Y, \mathbf{Z} \mid \mathbf{R}_w = \mathbf{0})$. Moreover, linear regression models in PermC are valid. Under Condition (i), we have $P(X, Y, \mathbf{Z} \mid \mathbf{W}, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}, \mathbf{R}_w = \mathbf{0}) = P(X, Y, \mathbf{Z} \mid \mathbf{W})$, in which X , Y , and \mathbf{Z} are conditionally Gaussian distributed given \mathbf{W} . Thus, we use linear regression to estimate $P(X^*, Y^*, \mathbf{Z}^* \mid \mathbf{W}^*, R_x = 0, R_y = 0, \mathbf{R}_z = \mathbf{0}, \mathbf{R}_w = \mathbf{0})$ and use them in the correction.

Density ratio weighted correction. DRW removes extraneous edges in situations where Condition (i) and Condition (ii) are not satisfied (e.g., Figure 1c). In these cases, we consistently estimate the joint distribution $P(\mathbf{V}_a)$ of concerned variables X , Y , and Z in a CI test and the direct causes $\mathbf{W} = Pa(R_x, R_y, R_z, \mathbf{R}_w)$, based on Theorem 2 of (Mohan et al., 2013), as shown in the first line of Equation 5. Here, \mathbf{R} represents the missingness indicators of \mathbf{V}_a . Equation 5 provides a way to reconstruct the observed dataset:

$$\begin{aligned} P(\mathbf{V}_a) &= \frac{P(\mathbf{R} = \mathbf{0}, \mathbf{V}_a)}{\prod_i P(R_i = 0 \mid Pa(R_i), R_{Pa(R_i)} = 0)} \\ &= P(\mathbf{V}_a \mid \mathbf{R} = \mathbf{0}) \times c \times \prod_i \beta_{Pa(R_i)}, \quad (5) \end{aligned}$$

where $c = \frac{P(\mathbf{R} = \mathbf{0})}{\prod_i P(R_i = 0 \mid R_{Pa(R_i)} = 0)}$ and $\beta_{Pa(R_i)} = \frac{P(Pa(R_i) \mid R_{Pa(R_i)} = 0)}{P(Pa(R_i) \mid R_i = 0, R_{Pa(R_i)} = 0)}$. In the second line of Equation 5, every (conditional) probability distribution can be consistently estimated. We first apply test-wise deletion to the observed data of \mathbf{V}_a . Then, we reweight such data with the density ratios $\prod_i \beta_{Pa(R_i)}$ and the normalizing constant c . We estimate density ratios $\prod_i \beta_{Pa(R_i)}$ with the kernel density estimation (Sheather and Jones, 1991) and compute the normalizing constant c . Finally, we test CI the relations of the concerned variables in the reweighted data samples from their corresponding joint distribution.

5 Experiments

We evaluate our method, MVPC, on both synthetic and real-world datasets. We first show experimental results on synthetic data (Section 5.1) and the behavior of our method in a situation with ground truth. After that, we apply our method to two healthcare datasets where data entries are significantly missing. The first is from the Cognition and

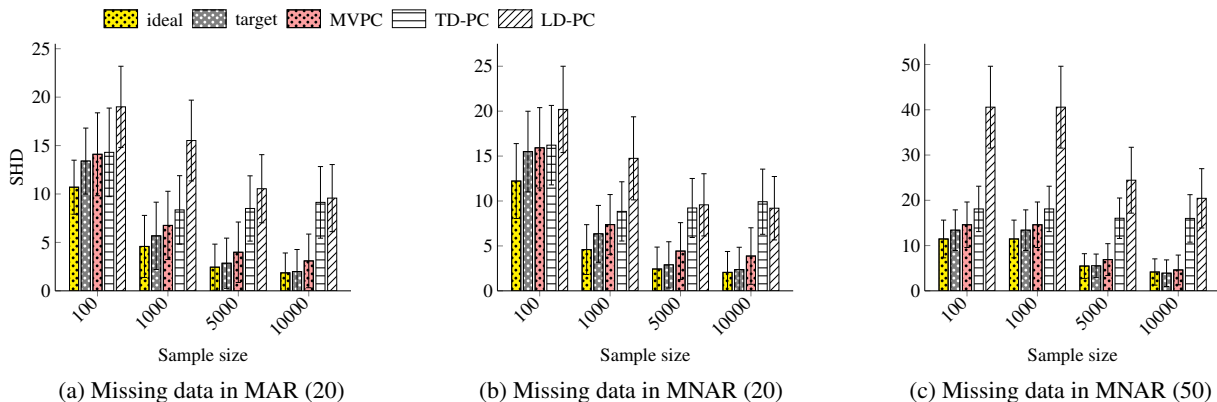


Figure 2: Performance comparison using structural Hamming distance. Lower value is better. Panel (a) shows the performance for MAR with 20 variables. Panel (b) and (c) show the performance for MNAR with 20 and 50 variables.




Figure 3: Precision and recall for adjacencies and orientation comparison (Higher is better). All experiments above use 20 nodes with 10000 data samples.

Aging USA (CogUSA) study (McArdle et al., 2015) (Section 5.2), and the second is about Achilles Tendon Rupture (ATR) rehabilitation research study (Praxitelous et al., 2017; Domeij-Arverud et al., 2016). MVPC demonstrates superior performance compared to multiple baseline methods.

5.1 Synthetic data evaluation

To best demonstrate the behavior of different causal discovery methods, we first perform the evaluation on synthetic data where the ground truth of causal graphs is known.

Baselines. Our baseline methods include deletion-based PC algorithms (as mentioned in Section 3): TD-PC and List-wise Deletion PC (LD-PC). Additionally, we apply the PC algorithm to the oracle data (without missing data), denoted by "ideal". Finally, to decouple the effect of sample size, we construct virtual datasets in MCAR with the same sample size as in each CI test of TD-PC. PC with such virtual MCAR data as a reference is denoted by "target".

Data Generation. We follow the procedures in (Colombo et al., 2012; Strobl et al., 2017) to randomly generate Gaussian DAG and sample data based on the given DAG. Additionally, we include at least two collider structures in the random Gaussian DAG, in order for deletion-based PC to have erroneous edges, as implied by Proposition 2. We generate two groups of synthetic data to show the scala-

bility of our methods: One group has 20 variables (with 6-10 partially observed variables), and the other is with 50 variables (with 10-14 partially observed variables) for MAR and MNAR. Note that in MNAR case, we assume that the direct causes of missingness indicators are partially observed. This is different from (Strobl et al., 2017), which assumes that the cause is a hidden variable. For each group of the experiments, we generate 400 DAGs with sample size of 100, 1000, 5000, and 10000, respectively.

Result. In all different experimental settings, we compare the results of different algorithms with structural Hamming distance from the ground truth, shown in Figure 2, and with the precision and recall of their adjacency and orientation, given in Figure 3. Across both metrics, as seen from Figure 2 and Figure 3, our proposed algorithm consistently has superior performance compared to both TD-PC and LD-PC, and is very close to the "target" performance. Similar to (Strobl et al., 2017), TD-PC also performs better than LD-PC in the context of PC. Additionally, our proposed method benefits from large volume of data samples as shown in Figure 2, in contrast to (Strobl et al., 2017).

5.2 The Cognition and aging USA (CogUSA) study

In this experiment, we aim to discovery causal relations in the CogUSA study as in (Strobl et al., 2017). This is a typical survey based healthcare dataset with a large amount

variables with missing values. In this scenario, the missingness mechanism is unknown and we could expect MCAR, MAR, and MNAR occur.

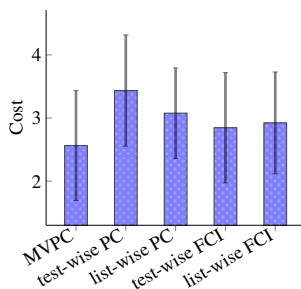


Figure 4: Performance of different methods on CogUSA study. Lower cost is better. The cost is the count of errors comparing with known causal constrains from experts.

We use the same 16 variables of interest in the CogUSA study as in (Strobl et al., 2017). Since the missingness indicators of the 16 variables can be caused by other variables, we utilize the rest variables when applying MVPC to the dataset. We use the BIC score for CI test (likelihood ratio test with the BIC penalty as the threshold). Figure 4 shows the performance evaluated using the known causal constraints: 1) Variables are in two groups with no inter-group causal relation; 2) there are causal relations between two pairs of variables given by the domain expertise. Each violation of these known causal relations adds 1 in the cost shown in Figure 4. Our proposed method obtains the best performance (lowest cost) comparing with deletion-based PC and deletion-based FCI (Strobl et al., 2017). This demonstrates the capabilities of our method in real life applications.

5.3 Achilles Tendon Rupture study

In the end, we perform causal discovery on a Achilles Tendon Rupture (ATR) study dataset (Praxitelous et al., 2017; Hamesse et al., 2018), collected in multiple hospitals¹. ATR is a type of soft tissue injury involving a long rehabilitation process. Understanding causal relations among various factors and healing outcomes is essential for practitioners. The list-wise deletion method is not applicable for this case because about 70% of the data entries are missing, which means that very rare patients have complete data. Thus, we apply our method and TD-PC to this dataset. We ran experiments on the full dataset with more than 100 variables. Figure 5a shows part of the causal graph.

We find that age, gender, BMI (body mass index), and LSI (Limb Symmetry Index) in the causal graph given by MVPC do not affect the healing outcome measured by Foot Ankle Outcome Score (FAOS). This result is consistent with (Praxitelous et al., 2017; Domeij-Arverud et al., 2016). To test the

¹In the ATR study experiment, only Paul Ackermann and Ruibo Tu get access to the ATR dataset.

effectiveness of MNAR, we further introduce an auxiliary variable S which is generated from two variables: Operation time (OP_{time}) and FAOS. This variable further causes the missingness indicator of FAOS. Figure 5b and 5c show the results of these variables using TD-PC and our proposed method. Our proposed MVPC is able to correctly remove the extraneous edge between Operation time and FAOS.

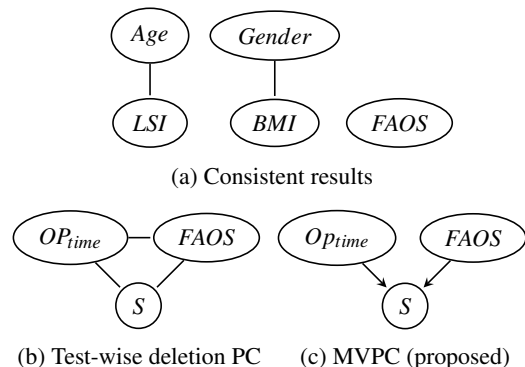


Figure 5: Causal discovery results in the ATR study. Experiments were run over all variables. We show only a part of the whole causal graph. Panel (a) shows the relations among five variables given by MVPC. The relations are consistent with medical studies. Panel (b) and (c) show an example where MVPC is able to correct the error of TD-PC.

6 Discussion

In this work, we address the problem of causal discovery in the presence of missing data. We first provide theoretical analysis to identify possible errors in the results given by a simple extension of PC. We then show that erroneous causal edges occur only in particular graph structures. Based on our analysis, we propose a novel algorithm MVPC, which corrects erroneous edges under mild assumptions. We demonstrate the asymptotic correctness and the effectiveness of our method on both synthetic data and real-world applications. As future work, we will explore the possibility of further relaxing the assumptions in MVPC, as well as work jointly with practitioners on causal analysis of large-scale healthcare applications in the presence of missing data.

Acknowledgements Kun would like to acknowledge the supported by the United States Air Force under Contract No. FA8650-17-C-7715, by National Institutes of Health under Contract No. NIH-1R01EB022858-01, FAINR01EB022858, NIH-1R01LM012087, NIH-5U54HG008540-02, and FAIN-U54HG008540, and by National Science Foundation EAGER Grant No. IIS-1829681. The National Institutes of Health, the U.S. Air Force, and the National Science Foundation are not responsible for the views reported in this article.

References

- D. M. Chickering. Optimal structure identification with greedy search. *Journal of machine learning research*, 3 (Nov):507–554, 2002.
- D. Colombo, M. H. Maathuis, M. Kalisch, and T. S. Richardson. Learning high-dimensional directed acyclic graphs with latent and selection variables. *The Annals of Statistics*, pages 294–321, 2012.
- E. Domeij-Arverud, P. Anundsson, E. Hardell, G. Barreng, G. Edman, A. Latifi, F. Labruto, and P. Ackermann. Ageing, deep vein thrombosis and male gender predict poor outcome after acute achilles tendon rupture. *Bone Joint J*, 98(12):1635–1641, 2016.
- A. Gain and I. Shpitser. Structure learning under missing data. In *International Conference on Probabilistic Graphical Models*, pages 121–132, 2018.
- C. Hamesse, P. Ackermann, H. Kjellstrom, and C. Zhang. Simultaneous measurement imputation and outcome prediction for achilles tendon rupture rehabilitation. In *ICML/IJCAI Joint Workshop on Artificial Intelligence in Health*, 2018.
- M. Klasson, K. Zhang, B. C. Bertilson, C. Zhang, and H. Kjellström. Causality refined diagnostic prediction. *arXiv preprint arXiv:1711.10915*, 2017.
- T. Le, T. Hoang, J. Li, L. Liu, H. Liu, and S. Hu. A fast pc algorithm for high dimensional causal discovery with multi-core pcs. *IEEE/ACM transactions on computational biology and bioinformatics*, 2016.
- J. McArdle, W. Rodgers, and R. Willis. Cognition and aging in the usa (cogusa) 2007-2009. *Assessment*, 2015.
- K. Mohan and J. Pearl. Graphical models for recovering probabilistic and causal queries from missing data. In *Advances in Neural Information Processing Systems*, pages 1520–1528, 2014a.
- K. Mohan and J. Pearl. On the testability of models with missing data. In *Artificial Intelligence and Statistics*, pages 643–650, 2014b.
- K. Mohan, J. Pearl, and J. Tian. Graphical models for inference with missing data. In *Advances in neural information processing systems*, pages 1277–1285, 2013.
- E. C. Neto, C. T. Ferrara, A. D. Attie, and B. S. Yandell. Inferring causal phenotype networks from segregating populations. *Genetics*, 2008.
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
- J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. MIT press, 2017.
- P. Praxitelous, G. Edman, and P. W. Ackermann. Microcirculation after achilles tendon rupture correlates with functional and patient-reported outcome. *Scandinavian journal of medicine & science in sports*, 2017.
- J. Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986.
- D. B. Rubin. Inference and missing data. *Biometrika*, 63(3): 581–592, 1976.
- D. B. Rubin. *Multiple imputation for nonresponse in surveys*, volume 81. John Wiley & Sons, 2004.
- S. J. Sheather and M. C. Jones. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 683–690, 1991.
- S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7(Oct):2003–2030, 2006.
- I. Shpitser. Consistent estimation of functions of data missing non-monotonically and not at random. In *Advances in Neural Information Processing Systems*, pages 3144–3152, 2016.
- E. Sokolova, D. von Rhein, J. Naaijen, P. Groot, T. Claassen, J. Buitelaar, and T. Heskes. Handling hybrid and missing data in constraint-based causal discovery to study the etiology of adhd. *International journal of data science and analytics*, 3(2):105–119, 2017.
- P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. MIT Press, Cambridge, MA, 2nd edition, 2001.
- E. V. Strobl, S. Visweswaran, and P. L. Spirtes. Fast causal inference with non-random missingness by test-wise deletion. *International Journal of Data Science and Analytics*, pages 1–16, 2017.
- K. Zhang and A. Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proceedings of the twenty-fifth conference on uncertainty in artificial intelligence*, pages 647–655. AUAI Press, 2009.
- K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*, Barcelona, Spain, 2011.