Causal Discovery for Linear Mixed Data

Author names withheld

Editor: Under Review for CLeaR 2022

Abstract

Discovery of causal relationships from observational data, especially from mixed data that consist of both continuous and discrete variables, is a fundamental yet challenging problem. Traditional methods focus on polishing the data type processing policy, which may lose data information. Compared with such methods, the constraint-based and score-based methods for mixed data derive certain conditional independence tests or score functions from the data's characteristics. However, they may return the Markov equivalence class due to the lack of identifiability guarantees, which may limit their applicability or hinder their interpretability of causal graphs. Thus, in this paper, based on the structural causal models of continuous and discrete variables, we provide sufficient identifiability conditions in bivariate as well as multivariate cases. We show that if the data follow our proposed restricted Linear Mixed causal model (LiM), such a model is identifiable. In addition, we proposed a two-step hybrid method to discover the causal structure for mixed data. Experiments on both synthetic and real-world data empirically demonstrate the identifiability and efficacy of our proposed LiM model.

Keywords: causal discovery, structural causal models, mixed data, identifiability

1. Introduction

Identifying the causal structure from purely observational data, termed as causal discovery, has been rapidly developed for the past decades with growing interest and has been widely applied in many domains (Pearl, 2000; Spirtes et al., 1993; Shimizu, 2014; Zhang and Hyvärinen, 2016). The traditional approaches to causal discovery roughly fall into two categories, namely constraint-based methods, e.g., PC (Spirtes and Glymour, 1991) and Fast Causal Inference (FCI) (Spirtes et al., 1995), and score-based ones, e.g., Greedy Equivalence Search (GES) (Chickering, 2002). Since they may output Markov equivalence class, i.e., a set of causal structures entailing the same conditional independence, they do not offer complete causal information. To distinguish different causal structures in the Markov equivalence class, several scholars derive additional assumptions on the data distribution and propose causal methods based on Structural Causal Models (SCM). These methods including Linear Non-Gaussian Acyclic Models (LiNGAM) (Shimizu et al., 2006), Additive Nonlinear Models (ANM) (Hoyer et al., 2009), and Post Nonlinear (PNL) (Zhang and Hyvärinen, 2009), achieve the unique identifiability of the causal structure. Most existing causal discovery methods focus on cases when the involved variables are either continuous or discrete only.

However, in many real-world scenarios such as economics (Wei et al., 2018), bioinformatics (Sedgewick et al., 2019), etc., the collected data often are a mixture of both continuous and discrete variables. When encountering such mixed data, one may ignore the discrete variables and apply the methods for continuous variables to estimate the partial causal network; or utilize a discretization policy to discretize the continuous variables, so that they can use those methods for discrete causal networks (Monti and Cooper, 1998; Chen et al., 2017). Both methods attempt to

convert mixed data types into the same type, which are naive and are easy to lose data information, inducing non-negligible estimation errors. Apart from these methods, by and large, causal discovery algorithms for mixed data can be categorized into two classes: constraint-based and score-based ones. Constraint-based algorithms are those variants of PC (Pearl, 2000; Spirtes and Glymour, 1991), including Cui et al. (2016); Sedgewick et al. (2019); Tsagris et al. (2018), which can not guarantee identifiability and are sensitive to samples.

Unlike the PC variants, score-based algorithms for mixed data do not use (un)conditional independence tests, but instead, optimize a likelihood score derived from mixed data's characteristics with the commonly-used greedy equivalence search framework. Such efforts include Li and Shimizu (2018); Huang et al. (2018); Andrews et al. (2019); Wei et al. (2018), etc. In particular, the first three efforts employ different score functions, i.e., with LiNGAM and the logistic regression model, regression model in RKHS, and degenerate distributions, respectively, whereas they do not provide the model's identifiability results and may return Markov equivalence class. Wei et al. (2018) developed a mixed causal model and proved its identifiability in the bivariate cases. However, the bivariate identifiability is not qualified enough to handle multivariate cases whereas the multivariate data ordinarily exist in many applications.

Thus, in this paper, we propose a structural causal model that consists of both continuous and discrete variables following Li and Shimizu (2018), and give its sufficient identifiability conditions in bivariate as well as multivariate cases. Compared with the mixed model developed by Wei et al. (2018), we allow more non-Gaussian distributions to be followed by the noises of continuous variables in the proof of the identifiability. Further, we derive a two-step hybrid method to uniquely estimate the causal structure without discretization. In the first phase, we develop a log-likelihood score function to characterize the joint distribution for mixed data. It is optimized accompanied with the acyclicity and sparsity constraints in a continuous optimization manner. The output causal structure here may fall into the solutions up to the ground truth's skeleton. To mitigate this issue, in the second phase, we search structures over the skeleton spaces and find the graph with the best score. Experiments on synthetic and real-world data demonstrate our proposed method's efficacy, compared with other methods.

Our contributions mainly are detailed in two-fold:

(i) For the mixed causal models that contain both continuous and discrete data, we prove the identifiability conditions in bivariate as well as multivariate cases. with which we enrich the identifiability space for causal discovery with mixed data.

(ii) We propose a score-based optimization method to infer the causal structure between mixed data. It is robust to the sample sizes and the ratio by the number of discrete variables to that of continuous variables.

2. Model Definition

We consider linear mixed causal models. Speaking concretely, suppose we are given p observed random variables, including discrete and continuous ones, i.e., $X = \{x_1, ..., x_p\}$. Since a categorical variable with T class can be regarded as $(T-1)$ binary variables, we assume that each discrete variable is binary (Wei et al., 2018). And further, we make the following assumptions, which is the same as in the SCM's definition of Li and Shimizu (2018):

A1. Observed variables x_i $(i = 1, \ldots, p)$ form a Directed Acyclic Graph (DAG).

A2. The value assigned to each continuous variable x_i is a linear function of its parent variables denoted by $x_{\mathrm{pa}(i)}$ plus a non-Gaussian error term e_i, that is,

$$x_i = e_i + c_i + \sum_{j \in \mathrm{pa}(i)} b_{ij} x_j, \quad e_i \sim Non - Gaussian(\cdot), \tag{1}$$

where the error terms e_i are continuous random variables with non-Gaussian densities, and the error variables e_i are independent of each other. The coefficients b_{ij} and intercepts c_i are constants.

A3. For each discrete variable x_i, its value equals to 1 if the linear function of its parent variables $x_{\mathrm{pa}(i)}$ plus a Logistic error term e_i is larger than 0, otherwise, its value equals to 0. That is,

$$x_i = \begin{cases} 1, & e_i + c_i + \sum_{j \in \mathrm{pa}(i)} b_{ij} x_j > 0 \\ 0, & \text{otherwise} \end{cases}, \quad e_i \sim Logistic(0, 1), \tag{2}$$

where the error terms e_i are identical to those in Eq.(1), but follow the Logistic distribution.

Definition 1 (Linear Mixed causal model, LiM) *If a causal model for mixed data satisfies the assumptions A1-A3, then this SCM is called a Linear Mixed causal model, abbreviated as LiM.*

Let $\mathcal{F} = \{f_c, f_d | f_c(x, e) = bx + e + c, f_d(x, e) = \begin{cases} 1, & bx + e + c > 0 \\ 0, & \text{otherwise} \end{cases}\}$ be a set of two functions which work on continuous and discrete variables, respectively. $\mathcal{P} = \{P_c, P_d\}$ denotes the set of probabilistic distributions for continuous and discrete variables. Using these notations, our model can be rewritten as:

$$x_i = f_i(x_{\mathrm{pa}_i}, e_i), \quad e_i \sim P(e_i), \tag{3}$$

where $f_i \in \mathcal{F}$, and $P(e_i) \in \mathcal{P}$.

3. Identifiability Conditions of the LiM

Here we provide a sufficient identifiability condition for the LiM model, using similar ideas of Wei et al. (2018) and Peters et al. (2014).

3.1. Bivariate cases

The LiM model of Section 2 is equivalent to the model of Wei et al. (2018) if the intercepts c_i are taken to be zeros and the error terms e_i follow the Laplace distributions $L(0, \alpha_i)$. Laplace distributions are commonly used in non-Gaussian models including independent component analysis and are known to be robust against the misspecification of the distributions if the right distribution is super-Gaussian (Hyvärinen et al., 2001). Wei et al. (2018) provided a sufficient identifiability condition for two-variable cases of their model, i.e., the two variables do not have the same marginal distributions if they are binary, all the probabilities and densities are positive, and the error variables are of non-zero variance. We can show the identifiability of our model in a similar manner to Wei et al. (2018). The difference lies in the fact that we allow more non-Gaussian distributions to be followed by the continuous error terms, rather than only the Laplace distributions.

Now we characterize the condition about the non-Gaussian distributions.

Condition 1 *The limit of non-Gaussian density ratio λ, defined as $\lambda := \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)}$, satisfies:*

$$\lambda = C, \tag{4}$$

where C is a non-zero finite constant. In other words, it follows that λ is neither equal to zero nor infinity ($\lambda \neq 0, \infty$).

Corollary 1 *If λ satisfies the Condition 1, then one of the following statements must hold.*

S1. The density ratio equals to a non-zero finite constant C, i.e., $\frac{P_c(x)}{P_c(x-b)} = C$.

S2. The density ratio equals to the product of a function of x and a non-zero finite constant C_0, while the limit of such a function as x goes to infinity is another non-zero finite constant C_1. That is, $\frac{P_c(x)}{P_c(x-b)} = C_0 \cdot g(x)$ and $\lim_{x\to\pm\infty} g(x) = C_1$, where $C = C_0 \times C_1$, resulting in $\lambda = \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)} = C_0 \cdot \lim_{x\to\pm\infty} g(x) = C$.

Proof We employ the contradiction method, i.e., if statements S1 and S2 both violated, we prove that either $\lambda = 0$ or $\lambda = \infty$ holds. If statements S1 and S2 both violated, we derive the following four cases: i) $\frac{P_c(x)}{P_c(x-b)} = 0$; ii)$\frac{P_c(x)}{P_c(x-b)} = \infty$; iii) $\frac{P_c(x)}{P_c(x-b)} = g(x) \cdot C_0$ and $\lim_{x\to\pm\infty} g(x) = 0$; iv) $\frac{P_c(x)}{P_c(x-b)} = g(x) \cdot C_0$ and $\lim_{x\to\pm\infty} g(x) = \infty$. For the case i), we have $\lambda = \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)} = 0$. For the case ii), we have $\lambda = \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)} = \infty$. For the case iii), we get $\lambda = \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)} = \lim_{x\to\pm\infty} g(x) \cdot C_0 = 0$. For the case iv), we have $\lambda = \lim_{x\to\pm\infty} \frac{P_c(x)}{P_c(x-b)} = \lim_{x\to\pm\infty} g(x) \cdot C_0 = \infty$. To be concluded, the corollary is proved. ■

Besides, it's noteworthy that basic non-Gaussian distributions satisfy the Condition 1, as illustrated in the Corollary 1. For instance, the Laplace distribution, Uniform distribution and Exponential distribution follow the statement S1; while the Gamma distribution follows the statement S2 with $\lim_{x\to\pm\infty} g(x) = 1$. With Condition 1 in the LiM model, we obtain our identifiability result in the bivariate case.

Theorem 2 *Let the data $\boldsymbol{X} = \{x_i, x_j\}$ be generated by the LiM model in Eqs.(1)-(2) with Condition 1. Under the conditions that x_i and x_j do not share the same marginal distributions if they are both discrete, and all the probabilities are positive, the model is identifiable.*

Proof We prove the identifiability for the bivariate case from three aspects: i) both variables are continuous; ii) both variables are discrete; iii) one is continuous and the other is discrete.

i), if the two variables are continuous, the model of Section 2 is a LiNGAM model (Shimizu et al., 2006). Therefore, the model is identifiable.

ii), suppose that two variables $\{x_i, x_j\}$ are binary. Assume that all the probabilities are positive and their marginal distributions are different. Then, we compare the following two models. The condition probability $P(x_i \mid x_j)$ of the first model $x_j \to x_i$ is written as:

$$P(x_i = 1 \mid x_j) = \frac{1}{1 + e^{-(c_i + b_{ij} x_j)}}, \tag{5}$$

$$P(x_i = 0 \mid x_j) = 1 - P(x_i = 1 \mid x_j). \tag{6}$$

The condition probability $P(x_j \mid x_i)$ of the second model $x_i \to x_j$ is written as:

$$P(x_j = 1 \mid x_i) = \frac{1}{1 + e^{-(c_i + b_{ji} x_i)}}, \tag{7}$$

$$P(x_j = 0 \mid x_i) = 1 - P(x_j = 1 \mid x_i). \tag{8}$$

Then, as in the proof of Theorem 6 of Wei et al. (2018), assume that the two models give the same joint distribution of observed variables x_i and x_j. Denote $P(x_i = 1)$ by k_i and $P(x_j = 1)$ by k_j. Then,

$$k_i^{x_i}(1 - k_i)^{1 - x_i}(\frac{1}{1 + e^{-(c_i + b_{ij} x_j)}})^{x_j}(1 - \frac{1}{1 + e^{-(c_i + b_{ij} x_j)}})^{1 - x_j} \tag{9}$$

$$= k_j^{x_j}(1 - k_j)^{1 - x_j}(\frac{1}{1 + e^{-(c_j + b_{ji} x_i)}})^{x_i}(1 - \frac{1}{1 + e^{-(c_j + b_{ji} x_i)}})^{1 - x_i}. \tag{10}$$

This induces $k_i = k_j$ for cases with $x_i = 0$ and $x_j = 0$, which contradicts with the assumption that the marginal distributions of x_i and x_j are different.

iii), suppose that one is continuous and the other is binary. Without loss of generality, assume that x_i is continuous and x_j is binary. Then, we compare the following two models. The first model $x_i \to x_j$ is written as:

$$x_j = \begin{cases} 1, & e_j + c_j + b_{ji} x_i > 0 \\ 0, & \text{otherwise} \end{cases}, \quad e_j \sim Logistic(0, 1), \tag{11}$$

where $x_i = e_i \sim Non - Gaussian(\cdot)$. The second model $x_j \to x_i$ is written as:

$$x_i = e_i + c_i + b_{ij} x_j, e_i \sim Non - Gaussian(\cdot), \tag{12}$$

where $x_j = e_j \sim Logistic(0, 1)$. Then, as in the proof of Theorem 7 of Wei et al. (2018), assume that the two models give the same distribution of observed variables.

For the first model, the conditional probability of $x_j = 1$ given x_i is given by

$$P(x_j = 1 \mid x_i) = \frac{1}{1 + e^{-(c_i + b_{ji} x_i)}}. \tag{13}$$

Then,

$$\lim_{x_i \to \infty} P(x_j = 1 \mid x_i) = \lim_{x_i \to \infty} \frac{1}{1 + e^{-(c_i + b_{ji} x_i)}} \tag{14}$$

$$= \begin{cases} 1 & (b_{ji} > 0) \\ 0 & (b_{ji} < 0) \end{cases}, \tag{15}$$

$$\lim_{x_i \to -\infty} P(x_j = 1 \mid x_i) = \lim_{x_i \to -\infty} \frac{1}{1 + e^{-(c_i + b_{ji} x_i)}} \tag{16}$$

$$= \begin{cases} 0 & (b_{ij} > 0) \\ 1 & (b_{ij} < 0) \end{cases}. \tag{17}$$

For the second model, the conditional probability of $x_j = 1$ given x_i is given by

$$P(x_j = 1 \mid x_i) = \frac{P(x_j = 1, x_i)}{P(x_i)} \tag{18}$$

$$= \frac{P(x_i \mid x_j = 1)P(x_j = 1)}{P(x_i \mid x_j = 1)P(x_j = 1) + P(x_i \mid x_j = 0)P(x_j = 0)} \tag{19}$$

$$= \frac{P(x_j = 1)}{P(x_j = 1) + \frac{P(x_i|x_j=0)}{P(x_i|x_j=1)}P(x_j = 0)}. \tag{20}$$

Due to the Condition 1, we obtain the limit of the density ratio as

$$\lim_{x_i \to \pm\infty} \frac{P(x_i \mid x_j = 0)}{P(x_i \mid x_j = 1)} = \lim_{x_i \to \pm\infty} \frac{P(x_i - c_i)}{P(x_i - c_i - e_i)} \tag{21}$$

$$= \lambda, \tag{22}$$

$$\tag{23}$$

where $\lambda = C \neq 0$ and $\lambda \neq \infty$. Thus, the limits $\lim_{x_i \to \pm\infty} P(x_j = 1 \mid x_i)$ under the second model are greater than 0 and smaller than 1 due to the assumption $P(x_j = 1) > 0$. This means that the limits $\lim_{x_i \to \pm\infty} P(x_j = 1 \mid x_i)$ under the second model are different from those of the first model, which contradicts the assumption that the two models give the same distribution of observed variables.

Thus, the model is bivariate identifiable if the two variables do not have the same marginal distributions in case that they are binary, all the probabilities and densities are positive, and the error variables are of non-zero variance. ■

The same would apply when Eq.(1) is replaced by a nonlinear model like ANM (Hoyer et al., 2009).

$$x_i = e_i + f_i(x_j), \quad x_j \in \mathrm{pa}(i), \tag{24}$$

where $f_i(0)$ and $f_i(1)$ are assumed to be finite.

Definition 3 (Bivariate Identifiable Set) *Let* $\mathcal{F} = \{f_c, f_d | f_c(x, e) = bx + e + c, f_d(x, e) = \begin{cases} 1, & bx + e + c > 0 \\ 0, & \text{otherwise} \end{cases}\}$ *be a set of two functions which work on continuous and discrete variables, respectively.* $\mathcal{P} = \{P_c, P_d\}$ *denotes the set of probabilistic distributions for continuous and discrete variables. Consider a mixed causal model with two variables* x_i *and* x_j*, i.e.,* $x_j = e_j$ *and* $x_i = f_i(x_j, e_i)$ *with* $x_j \perp e_i$*. We call a set* $\mathcal{B} \subseteq \mathcal{F} \times \mathcal{P} \times \mathcal{P}$ *as a bivariate identifiability set if the triple* $(f_i, P(x_j), P(e_i))$ *where* $f \in \mathcal{F}$*, and* $P(x_j), P(e_i) \in \mathcal{P}$ *hold, follows our LiM model's assumptions.*

Using the definition of the bivariate identifiable set $\mathcal{B}$, if the triple $(f_i, P(x_j), P(e_i))$ follows our LiM model, we have $(f_i, P(x_j), P(e_i)) \in \mathcal{B}$, which means that the bivariate mixed causal model is identifiable.

3.2. From bivariate cases to multivariate cases

Hoyer et al. (2009) proposed a structural causal model named as additive noise model for causal discovery and considered the identifiability condition for two variables, i.e., its bivariate identifiability condition. Peters et al. (2014) showed the model also is identifiable for more than two variables. In the proof of Theorem 28 of Peters et al. (2014), under the assumption of causal minimality and that of positive densities, they showed that if two different causal graphs are assumed to give the same distributions of observed variables, it results in contradiction to the bivariate identifiability. Remark 30 of Peters et al. (2014) says that

> Whenever we have a restriction like Condition 19 that ensures identifiability in the bivariate case (Theorem 20), the multivariate version (Theorem 28) remains valid. The proof we provide in the appendix stays exactly the same.

In fact, most parts of their proof other than the bivariate identifiability condition use only the general properties of non-parametric structural causal models with no hidden common causes and no cycles, and do not depend on the assumptions of additive noise models. Therefore, our LiM model also is identifiable for more than two variable cases using the idea of Peters et al. (2014) since our model is of bivariate identifiability under some conditions as shown above in Subsection 3.1.

Here, we first define the restricted mixed causal models to constrain conditional distributions and thereafter give identifiability analysis, in a similar manner to Peters et al. (2014).

Definition 4 (Restricted LiM model) *Consider a LiM model with p variables. We call this SCM a restricted LiM model if for all $i \in \{1, ..., p\}$, $j \in \mathrm{pa}(i)$, and all sets $\mathbf{S} \subseteq \{1, ..., p\}$ with $\mathrm{pa}(i)\backslash\{j\} \subseteq \mathbf{S} \subseteq \mathrm{nd}(i)\backslash\{i, j\}$, there exists an $x^*_{\mathbf{S}}$ with $P_{\mathbf{S}}(x^*_{\mathbf{S}}) > 0$, s.t.*

$$(f_i(x_{pa(i)\backslash\{j\}}, \underbrace{\cdot}_{x_j}), P(x_j|x_{\mathbf{S}} = x^*_{\mathbf{S}}), P(e_i)) \tag{25}$$

satisfies the assumptions and the Condition 1 of the LiM's model in Section 2, i.e.,

$$(f_i(x_{pa(i)\backslash\{j\}}, \underbrace{\cdot}_{x_j}), P(x_j|x_{\mathbf{S}} = x^*_{\mathbf{S}}), P(e_i)) \in \mathcal{B}, \tag{26}$$

where the underbrace with x_j represents x_j's component with f_i, and $f_i \in \mathcal{F}$. Further, we require that the noise variables to have non-vanishing densities.

Theorem 5 *Let the data $\boldsymbol{X} = \{x_1, ..., x_p\}$ be generated by a restricted LiM model. Under the conditions that any two discrete variables do not share the same marginal distributions, all the probabilities are positive, and $P(\boldsymbol{X})$ satisfies the Markov and faithful conditions, the model is identifiable.*

Proof The theorem is proved by contradiction. We assume that our restricted mixed causal model is not identifiable, i.e., there exist two restricted mixed causal models $\mathcal{G}_1$ and $\mathcal{G}_2$, which induce the identical joint distribution $P(\mathbf{X})$. In such a case, we will show that $\mathcal{G}_1 = \mathcal{G}_2$ to induce the identifiability.

Consider two variables x_i and x_j in $\boldsymbol{X}$ where for the sets $\mathbf{Q} := \mathrm{pa}(i)^{\mathcal{G}_1}\backslash\{j\}$, $\mathbf{R} := \mathrm{pa}(j)^{\mathcal{G}_2}\backslash\{i\}$, and $\mathbf{S} := \mathbf{Q} \cup \mathbf{R}$, they satisfy i) $x_j \to x_i$ in $\mathcal{G}_1$ and $x_i \to x_j$ in $\mathcal{G}_2$; and ii) $\mathbf{S} \subseteq \mathrm{nd}(i)^{\mathcal{G}_1}\backslash\{j\}$

and $\mathbf{S} \subseteq \mathrm{nd}(j)^{\mathcal{G}_2} \backslash \{i\}$. Such two variables do exist (Peters et al., 2014). Firstly, due to ii) we get $e_i \perp (x_j, x_\mathbf{S})$ and $e_j \perp (x_i, x_\mathbf{S})$. Let $x^*_\mathbf{S} = \{x_\mathrm{q}, x_\mathrm{r}\}$. For the graph $\mathcal{G}_1$, we get $(f_i(x_\mathrm{q}, \cdot), P(x_j|x_\mathbf{S} = x^*_\mathbf{S}), P(e_i)) \in \mathcal{B}$, which satisfies the assumptions of our bivariate mixed causal model. It induces

$$x_i = f_i(x_\mathrm{q}, x^*_j), \quad x_\mathrm{q} \perp x^*_j, \tag{27}$$

where $x^*_j := x_j|x_\mathbf{S} = x_{\mathbf{S}^*}$ and $f_i \in \mathcal{F}$. For the graph $\mathcal{G}_2$, we get $(f_j(x_\mathrm{r}, \cdot), P(x_i|x_\mathbf{S} = x^*_\mathbf{S}), P(e_j)) \in \mathcal{B}$, which satisfies the assumptions of our bivariate LiM model. It induces

$$x_j = f_j(x_\mathrm{r}, x^*_i), \quad x_\mathrm{r} \perp x^*_i, \tag{28}$$

where $x^*_i := x_i|x_\mathbf{S} = x_{\mathbf{S}^*}$ and $f_j \in \mathcal{F}$. The above analysis contradicts the bivariate identifiability result in Theorem 2, hence we have $\mathcal{G}_1 = \mathcal{G}_2$. ■

4. Optimization Method

To uncover the causal structure for mixed data that consist of both continuous and discrete variables, we propose an integrated hybrid score-based learning method. The objective function is based on the negative log-likelihood of the data. By instantiating the negative log-likelihood with the joint probability distribution of the mixed data, we get

$$\begin{aligned}
\mathcal{L}(\mathbf{B}) &= -\log(P(\mathbf{X})) && (29)\\
&= -\log[\prod_t^n \prod_i^p P_d(x_{i,t} \mid x_{\mathrm{pa}(i),t})^{z_i} P_c(x_{i,t} \mid x_{\mathrm{pa}(i),t})^{1-z_i}] && (30)\\
&= -\sum_t^n \sum_i^p z_i \log[P_d(x_{i,t} \mid x_{\mathrm{pa}(i),t})] + (1 - z_i) \log[P_c(x_{i,t} \mid x_{\mathrm{pa}(i),t})], && (31)\\
&= -\sum_t^n \sum_i^p z_i \{x_{i,t} \log[\sigma(\mathbf{B})] + (1 - x_{i,t}) \log[1 - \sigma(\mathbf{B})]\} + && (32)\\
&\quad (1 - z_i) \log p_i(x_{i,t} - \sum_{k \in \mathrm{pa}(i)} b_{ik} x_{k,t}), && (33)
\end{aligned}$$

where $\mathbf{B}$ is the adjacency matrix, and $x_{i,t}$ is the t^{th} sample of the i^{th} variable x_i. n is the sample size. P_d and P_c denote the probability distribution of discrete and continuous variables, respectively. z_i is an indicator variable, where $z_i = 1$ if x_i is discrete while $z_i = 0$ otherwise. $\sigma(\mathbf{B}) = \frac{1}{1+e^{-(c_i + \sum_{l \in \mathrm{pa}(i)} b_{il} x_{l,t})}}$, while p_i is the density function of the non-Gaussian error terms e_i of continuous variables. In our method, we specified p_i to be the density function of Laplace distributions. But, we can use other density functions as well. Thereafter, we seek to solve the following continuous optimization problem:

$$\begin{aligned}
\min_{\mathbf{B}} \quad & \mathcal{L}(\mathbf{B}) + \Lambda ||\mathbf{B}||_1 \\
\text{subject to} \quad & h(\mathbf{B}) = 0,
\end{aligned} \tag{34}$$

where $h(\mathbf{B})$ is an acyclicity constraint which ensures that $\mathbf{B}$ is a DAG (Zheng et al., 2018), Λ is a regularization parameter, and $|| \cdot ||_1$ is an l_1 sparsity regularization. Following the optimization

Algorithm 1 LiM Algorithm

Require: Data $\boldsymbol{X}$; indicator vector Z; threshold ϵ; tolerance parameter ω.
Ensure: Connection strengths matrix $\mathbf{B}^*$.
Phase I: Global search
Optimize Eq.(34) to obtain $\hat{\mathbf{B}}$ using QPM with the tolerance parameter ω,
Rule out edges whose connection strengths are below ϵ: $\hat{\mathbf{B}} = \hat{\mathbf{B}} \circ \mathbf{0}(\hat{b}_{ij} < \epsilon)$.
Phase II: Local search
Initiate a temporary minimum log-likelihood as $\mathcal{L}_{tmp}(\mathbf{B}^*) = \mathcal{L}(\hat{\mathbf{B}})$.
while $\mathbf{B} \in \mathrm{Ske}(\hat{\mathbf{B}})$ **and** $h(\mathbf{B}) < \omega$ **do**
  Compute the negative log-likelihood $\mathcal{L}(\mathbf{B})$ for $\mathbf{B}$.
  if $\mathcal{L}(\mathbf{B}) < \mathcal{L}_{tmp}(\mathbf{B}^*)$ **then**
    $\mathcal{L}_{tmp}(\mathbf{B}^*) = \mathcal{L}(\mathbf{B})$.
  end if
end while
return $\mathbf{B}^*$ with the minimal log-likelihood $\mathcal{L}_{tmp}(\mathbf{B}^*)$.

procedure in (Zeng et al., 2021), we leverage the Quadratic Penalty Method (QPM) to estimate $\mathbf{B}$, converting Eq.(34) into an unconstrained function

$$\min_{\mathbf{B}} \mathcal{S}(\mathbf{B}), \tag{35}$$

where $\mathcal{S} = \mathcal{L}(\mathbf{B}) + \Lambda||\mathbf{B}||_1 + \frac{\rho}{2}h(\mathbf{B})^2$ is the quadratic penalty function, and ρ is a regularization parameter. Then we utilize the L-BFGS-B (Byrd et al., 1995) to solve Eq.(35). Due to the machine precision, it is well-known that the estimated b_{ij} is hard to receive absolute zeros if such pairs have no edges (Zheng et al., 2018). Hence we give a small fixed threshold ϵ to rule out those whose estimated effects are lower than ϵ.

However, such an optimization method may easily fall into its skeleton solutions due to the sub-optimality. To mitigate this issue, after obtaining the estimated adjacency matrix $\hat{\mathbf{B}}$, we tackle further the following combinatorial optimization problem:

$$\mathbf{B}^* = \arg \min_{\mathbf{B} \in \mathrm{Ske}(\hat{\mathbf{B}}), h(\mathbf{B}) < \omega} \mathcal{L}(\mathbf{B}), \tag{36}$$

where $\mathrm{Ske}(\hat{\mathbf{B}})$ represents a set where the containing DAGs entail the same skeleton as $\hat{\mathbf{B}}$ and ω is a tolerance parameter. Compared with the traditional approaches that search over the discrete space of DAGs with the full graph, we perform our structure search over the narrowed space of DAGs within the estimated skeleton, which possesses an advantage in computational efficiency. The full algorithm is outlined in Algorithm 1.

As demonstrated in Algorithm 1, our LiM approach firstly performs global updates, estimating the connection strengths matrix $\hat{\mathbf{B}}$ in one step with continuous optimization techniques. Then, it performs a local update to search over the skeleton space, estimating one candidate DAG with one changing edge at each iteration in a combinatorial optimization manner. To conclude, the LiM approach is a two-step hybrid method, which takes advantages of both global and local search to avoid falling into sub-optimal solutions and to be more computationally efficient.

5. Experiments

In this section, we performed simulation experiments and employed our method to real-world application data to learn the causal graph with mixed data, evaluating the efficacy of our proposed method.

5.1. Synthetic data

To generate the data in simulations, we firstly established a randomly unweighted Directed Acyclic Graph (DAG) according to the ER models, where the number of edges was randomly selected. Given the DAG, we assigned uniformly the edge weights from $[-2,-1]\bigcup[1,2]$ to get an adjacency matrix $\mathbf{B}$. Without loss of generality, the number of discrete variables was selected randomly from $[1,(p-1)]$, and thereafter we randomly assigned the discrete and continuous variables. Finally, the data were generated according to our LiM model in Eqs.(1)-(2).

We compared our method with a constraint-based method, a variant of PC algorithm (PC) (Spirtes and Glymour, 1991) as a representative. It discretized all continuous variables into discrete ones, following Li and Shimizu (2018). Besides, since PC may return a DAG pattern (PDAG) instead of a unique DAG, we took its instance for evaluation. We compared with score-based methods, including the scores of Notears with the Logistic (logistic) or Laplace (laplace) distributions. We also compared with a commonly-used functional-based method, the LiNGAM method (LiNGAM). To emphasize the necessity of our local search phase, we took our LiM method without the second phase as a comparison as well (mixed).

In these experiments, we evaluated the performance of all methods in terms of precision (Pre.), recall (Re.) and F1 score (F1), where the F1 is defined as $\mathrm{F1}=\frac{2\times\mathrm{Pre.}\times\mathrm{Re.}}{\mathrm{Pre.}+\mathrm{Re.}}$. For those continuous optimization methods, we chose the threshold $\epsilon=0.1$, the tolerance $\omega=1e-8$, and the regularization parameter $\Lambda=0.1$, while for those which exploit conditional independence tests, we fixed the significance level to be 0.01. For other parameters, we adopted their default settings.

We performed the simulations using i) different sample sizes, i.e., $n=$ 50, 100, 500, 1000, 2000, 5000, with bivariate and 5 mixed variables in turn. In addition, we generated the data with ii) different numbers of discrete variables ranging from 1 to $(p-1)$ where each graph has 5000 samples, to test the robustness of our proposed LiM method. For each setting, we experimented with 50 realizations and reported the average results.

Sensitivity to different sample sizes. Figure 1 gives the Re., Pre. and F1 of the recovered causal graph with 2 or 5 mixed variables, compared with PC, logistic, laplace, LiNGAM, and mixed methods. The x-axis shows the sample sizes, while the y-axis is the Re., Pre, or F1. Overall, our LiM method gives the best accuracy in both settings, which verified the identifiability results, especially in bivariate cases. More specifically, the LiM, and PC methods' accuracies increase remarkably along with the sample sizes. Though our LiM method is sensitive to sample sizes for multi-variate causal networks, LiM does perform better than other comparisons in the scheme of small sample size $n=50$. On the contrary, PC's unsatisfactory performance is basically due to the usage of conditional independence tests. Score-based methods are more robust to the sample sizes compared with the constraint-based ones. However, since their score functions may not fit the mixed data or they may be trapped in sub-optimal problems, their performances are not comparable despite stability.

Sensitivity to different discrete variables. Figure 2 reports the Re., Pre. and F1 of the recovered causal graph with different numbers of discrete variables $p_d=$1, 2, 3, where there are a total of

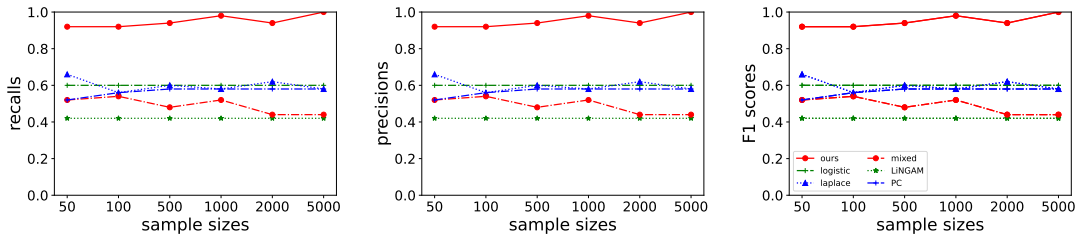

(a) Re. with 2 vars. (b) Pre. with 2 vars. (c) F1 with 2 vars.

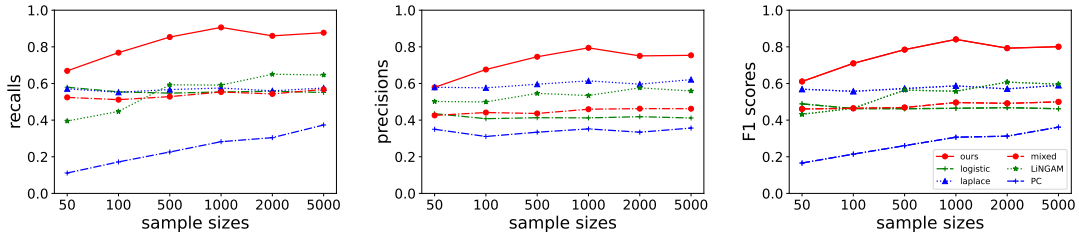

(d) Re. with 5 vars. (e) Pre. with 5 vars. (f) F1 with 5 vars.

Figure 1: Recalls (Re.), Precisions (Pre.), and F1 scores (F1.) of recovered causal graphs for bivariate and 5 mixed variables with different sample sizes. Higher F1, Re., and Pre. mean higher accuracies.

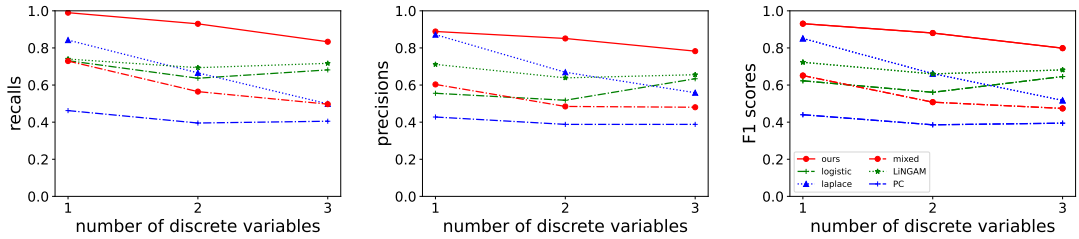

(a) Re. with 4 vars. (b) Pre. with 4 vars. (c) F1 with 4 vars.

Figure 2: Recalls (Re.), Precisions (Pre.), and F1 scores (F1.) of recovered causal graphs with different numbers of discrete variables, where the sample size is 5000.

4 observed variables. The x-axis shows the number of discrete variables, while the y-axis is the Re., Pre, or F1. As shown, we can see that overall, our method performed better than other methods, indicating the capability of handling mixed data.

5.2. Real-world data

Boston housing data set. We then applied our LiM method to a real-world Boston housing data set, which was collected at the UCI Repository (Dua and Graff, 2017). Such a data set contains 506 data points and we chose 11 variables for the experiments, where the chosen continuous variables

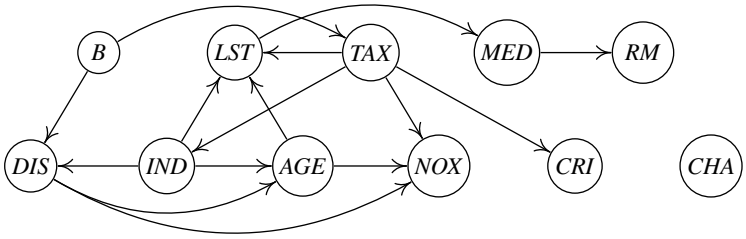

Figure 3: Result of the LiM method applied to the Boston housing data set.

are identical to Zhang et al. (2011) and the only binary variable is included. They are *CRI* (continuous, per capita crime rate by town), *INDUS* (proportion of non-retail business acres per town), *CHA* (binary, if tract bounds river or not), *NOX* (continuous, nitric oxides concentration), *RM* (continuous, average number of rooms per dwelling), *AGE* (continuous, proportion of owner-occupied units built prior to 1940), *DIS* (continuous, weighted distances to five Boston employment centres), *TAX* (continuous, full-value property-tax rate per 10,000 dollars), *B* (continuous, the proportion of blacks by town), *LST* (continuous, lower status of the population), and *MED* (continuous, median value of owner-occupied homes). We used the same settings as in the simulation experiments, i.e., we set the ruled-out threshold $\epsilon = 0.1$ and tolerance parameter $\omega = 1e - 8$. To provide reliable performance, due to the different scales of the large number of variables, we standardized the continuous variables before employing our method. The resulting causal graph is demonstrated in Figure 3. Though it is arguable that *RM* may not an effect variable, we still found some interesting conclusions which were accordance with our common understandings. For example, *MED* is influenced by *LST*, which are determined by some house-related indicators, i.e., *IND*, *AGE* and *TAX*; there is no direct link between *NOX* and *MED* but they are dependent through intermediate causal relationships (Margaritis, 2005); it is reasonable that *TAX*, which reflects the government's housing policy, influences *IND*, *LST*, and *CRI*. The results illustrated the effectiveness of our proposed LiM method in inferring causal graphs from mixed data.

6. Conclusions

In this paper, we provided complete identifiability conditions for causal discovery with linear mixed data that consist of continuous and discrete variables, both in bivariate and multivariate cases. Further, we proposed a two-step hybrid approach to uniquely identify the causal structure. Experiments on synthetic as well as real-world data demonstrate that our LiM method outperformed the comparisons. A line of our future research is to generalize the identifiability to cases where there are nonlinear relationships or confounders in the underlying causal graphs.

References

Bryan Andrews, Joseph Ramsey, and Gregory F. Cooper. Learning high-dimensional directed acyclic graphs with mixed data-types. In *The 2019 ACM SIGKDD Workshop on Causal Discovery*, pages 4–21. PMLR, 2019.

Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on scientific computing*, 16(5):1190–1208, 1995.

Yi-Chun Chen, Tim A. Wheeler, and Mykel J. Kochenderfer. Learning discrete bayesian networks from continuous data. *Journal of Artificial Intelligence Research*, 59:103–132, 2017.

David Maxwell Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 2002.

Ruifei Cui, Perry Groot, and Tom Heskes. Copula PC algorithm for causal discovery from mixed data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 377–392. Springer, 2016.

Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL http://archive.ics.uci.edu/ml.

Patrik O. Hoyer, Dominik Janzing, Joris Mooij, Jonas Peters, and Bernhard Schölkopf. Nonlinear causal discovery with additive noise models. In *Advances in Neural Information Processing Systems 21*, pages 689–696. 2009.

Biwei Huang, Kun Zhang, Yizhu Lin, Bernhard Schölkopf, and Clark Glymour. Generalized score functions for causal discovery. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1551–1560, 2018.

Aapo Hyvärinen, J Karhunen, and E Oja. *Independent Component Analysis*. Wiley, New York, 2001.

Chao Li and Shohei Shimizu. Combining linear non-gaussian acyclic model with logistic regression model for estimating causal structure from mixed continuous and discrete data. *arXiv preprint arXiv:1802.05889*, 2018.

Dimitris Margaritis. Distribution-free learning of Bayesian network structure in continuous domains. In *AAAI*, volume 5, pages 825–830, 2005.

Stefano Monti and Gregory F. Cooper. A multivariate discretization method for learning Bayesian networks from mixed data. In *Proc. 14th Conference on Uncertainty in Artificial Intelligence (UAI1998)*, pages 404–413, 1998.

Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

Jonas Peters, Joris M. Mooij, Dominik Janzing, and Bernhard Schölkopf. Causal discovery with continuous additive noise models. *Journal of Machine Learning Research*, 15:2009–2053, 2014.

Andrew J Sedgewick, Kristina Buschur, Ivy Shi, Joseph D Ramsey, Vineet K Raghu, Dimitris V Manatakis, Yingze Zhang, Jessica Bon, Divay Chandra, Chad Karoleski, et al. Mixed graphical models for integrative causal analysis with application to chronic lung disease diagnosis and prognosis. *Bioinformatics*, 35(7):1204–1212, 2019.

Shohei Shimizu. LiNGAM: Non-Gaussian methods for estimating causal structures. *Behaviormetrika*, 41(1):65–98, 2014.

Shohei Shimizu, Patrik O. Hoyer, Aapo Hyvärinen, and Antti Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 2006.

Peter Spirtes and Clark Glymour. An algorithm for fast recovery of sparse causal graphs. *Social Science Computer Review*, 9:67–72, 1991.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer Verlag, 1993. (2nd ed. MIT Press 2000).

Peter Spirtes, Christopher Meek, and Thomas S Richardson. Causal inference in the presence of latent variables and selection bias. In *Proc. 11th Annual Conference on Uncertainty in Artificial Intelligence (UAI1995)*, pages 491–506, 1995.

Michail Tsagris, Giorgos Borboudakis, Vincenzo Lagani, and Ioannis Tsamardinos. Constraint-based causal discovery with mixed data. *International journal of data science and analytics*, 6 (1):19–30, 2018.

Wenjuan Wei, Feng Lu, and Chunchen Liu. Mixed causal structure discovery with application to prescriptive pricing. In *Proc. 27rd International Joint Conference on Artificial Intelligence (IJCAI2018)*, pages 5126–5134, 2018.

Yan Zeng, Shohei Shimizu, Ruichu Cai, Feng Xie, Michio Yamamoto, and Zhifeng Hao. Causal discovery with multi-domain lingam for latent factors. In *Proc. 30th International Joint Conference on Artificial Intelligence (IJCAI2021)*, 2021.

Kun Zhang and Aapo Hyvärinen. On the identifiability of the post-nonlinear causal model. In *Proc. 25th Conference on Uncertainty in Artificial Intelligence (UAI2009)*, pages 647–655, 2009.

Kun Zhang and Aapo Hyvärinen. Nonlinear functional causal models for distinguishing causes form effect. In W. Wiedermann and A. von Eye, editors, *Statistics and Causality: Methods for Applied Empirical Research*. Wiley & Sons, 2016.

Kun Zhang, Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. Kernel-based conditional independence test and application in causal discovery. In *27th Conference on Uncertainty in Artificial Intelligence (UAI 2011)*, pages 804–813. AUAI Press, 2011.

Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. Dags with no tears: Continuous optimization for structure learning. *Advances in Neural Information Processing Systems*, 31, 2018.