

A novel method for Causal Structure Discovery from EHR data: A demonstration on type-2 diabetes mellitus

Xinpeng Shen

University of Minnesota

Sisi Ma

University of Minnesota

Prashanthi Vemuri

Mayo Clinic

M. Regina Castro

Mayo Clinic

Pedro J. Caraballo

Mayo Clinic

Gyorgy J. Simon (✉ simo0342@umn.edu)

University of Minnesota

Research Article

Keywords: Causality, Causal Structure Discovery, Electronic Health Records, Type 2 Diabetes Mellitus, Machine Learning

Posted Date: May 14th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-486814/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

A novel method for Causal Structure Discovery from EHR data:
A demonstration on type-2 diabetes mellitus

Xinpeng Shen^a, Sisi Ma^{a,b}, Prashanthi Vemuri^c, M. Regina Castro^d, Pedro J. Caraballo^e, Gyorgy J. Simon^{a,b*}

^aInstitute for Health Informatics, University of Minnesota, Minneapolis, MN, USA

^bDepartment of Medicine, University of Minnesota, Minneapolis, MN, USA

^cDepartment of Radiology, Mayo Clinic, Rochester, MN, USA

^dDivision of Endocrinology, Mayo Clinic, Rochester, MN, USA

^eDepartment of Internal Medicine, and Department of Health Sciences Research, Mayo Clinic, Rochester, MN, USA

Keywords: Causality, Causal Structure Discovery, Electronic Health Records, Type 2 Diabetes Mellitus, Machine Learning

Abstract

Modern AI-based clinical decision support models owe their success in part to the very large number of predictors they use. Safe and robust decision support, especially for intervention planning, requires causal, not associative, relationships. Traditional methods of causal discovery, clinical trials and extracting biochemical pathways, are resource intensive and may not scale up to the number and complexity of relationships sufficient for precision treatment planning. Computational causal structure discovery (CSD) from electronic health records (EHR) data can represent a solution, however, current CSD methods fall short on EHR data. This paper presents a CSD method tailored to the EHR data. The application of the proposed methodology was demonstrated on type-2 diabetes mellitus. A large EHR dataset from Mayo Clinic was used as development cohort, and another large dataset from an independent health system, M Health Fairview, as external validation cohort. The proposed method achieved very high recall (.95) and substantially higher precision than the general-purpose methods (.84 versus .29, and .55). The causal relationships extracted from the development and external validation cohorts had a high (81%) overlap. Due to the adaptations to EHR data, the proposed method is more suitable for use in clinical decision support than the general-purpose methods.

Introduction

Diagnostic tools based on artificial intelligence (AI) have recently demonstrated human-like performance[1-4], owing their high performance to their ability to synthesize information from many features. Consistent with this observation, national initiatives such as the Precision Medicine Initiative[5] and the Learning Health Systems[6] encourage the inclusion of a wide-range of information about the patient into the decision making process. Increasingly, clinical decision support systems start to include treatment planning and selection tools[7]. Such tools require causal knowledge, not merely the associations (correlations). Intervening on correlates rather than causal factors of the disease leads to lack of efficacy, under- or overtreatment, and in worst case, to iatrogenic harm[8].

The gold standard for discovering causal relationships is conducting a randomized clinical trial or elucidating the underlying biochemical pathways. In many cases, clinical trials are impractical, unethical, if not outright impossible. Computational causal structure discovery (CSD) methods to discover causal relationships have demonstrated great success in many domains[9-11] and their application to EHR data could offer a solution for causal discovery from observational real world medical data. However, to unlock their full potential, these general-purpose algorithms need to be adapted to address study design and data quality challenges specific to the EHR data.

We propose an algorithm with three adaptations. First, we incorporate *study design considerations*. EHR data as it exists in the system does not follow any study design. Billing codes in particular are recorded for reimbursement purposes and do not distinguish between new incidences and pre-existing conditions. Understanding this difference is critical for study design. Second, *time stamps can be unreliable*. The time stamp of a diagnosis often does **not** coincide with the onset time of the disease, but rather reflects the documentation time. In some cases, the temporal ordering of diseases may be reversed. Partly for this reason, general purpose CSD algorithms applied to the EHR data occasionally report “causal” relationships that are in the opposite direction of the natural disease progression. Third, *general-purpose CSD methods sometimes fail to orient edges*. Even when a clear causal direction exists and is not masked by data artifacts, CSD algorithms can have difficulty distinguishing the cause from the effect due to statistical equivalence[12]. Leveraging the longitudinal nature of EHR data and incorporating time information as part of the causal discovery process can enhance the identification of edge orientation.

In this paper, (1) we propose a data transformation procedure that distinguishes new incidences from pre-existing conditions, which allows us to determine the temporal order of the disease-related events despite the inaccurate (or rather noisy) timestamps in the EHR data. (2) We then present modifications to an existing CSD method, (Fast) Greedy Equivalence Search (GES)[13 14], to infer the direction of causal relationships more robustly using longitudinal information and takes the above study design considerations into account.

We demonstrate this methodology through the clinical example of type-2 diabetes mellitus (T2D), its risk factors and complications. T2D is an exceptionally well-studied disease with numerous clinical trials having produced a vast knowledge base, making the evaluation of the methodology possible. The goal of this work is not to uncover new causal relationships in diabetes, but to present a methodology for discovering causal relationships from EHR data that are sufficiently robust to support model development for clinical decision support tools. While we use T2D as our use case, we expect our methods to generalize to other diseases, typically chronic diseases, that exhibit similar characteristics and suffer from the same EHR shortcomings.

Methods

Study Source and Population

This retrospective cohort study utilized EHR data sets from two independent health systems, Mayo Clinic (MC) in Rochester, Minnesota and M Health Fairview (FV) in Minneapolis, Minnesota. Two 2-year time windows 2003-2004 and 2006-2007 for MC; and 2008-2009, and 2011-2012 for FV were defined. Dates for the time windows differed between MC and FV due to data availability. We extracted diagnoses, prescriptions, laboratory results, and vital signs from the two EHR data sets with the same inclusion and exclusion criteria: patients must have at least two blood pressure measurements, one before the first time window and one after the second time window; aged 18+ at the end of the first time window; and sex and age must be known. Figure 1A shows an overview of the study design of MC EHR (the study design for FV is similar). We used the MC EHR as the development cohort.

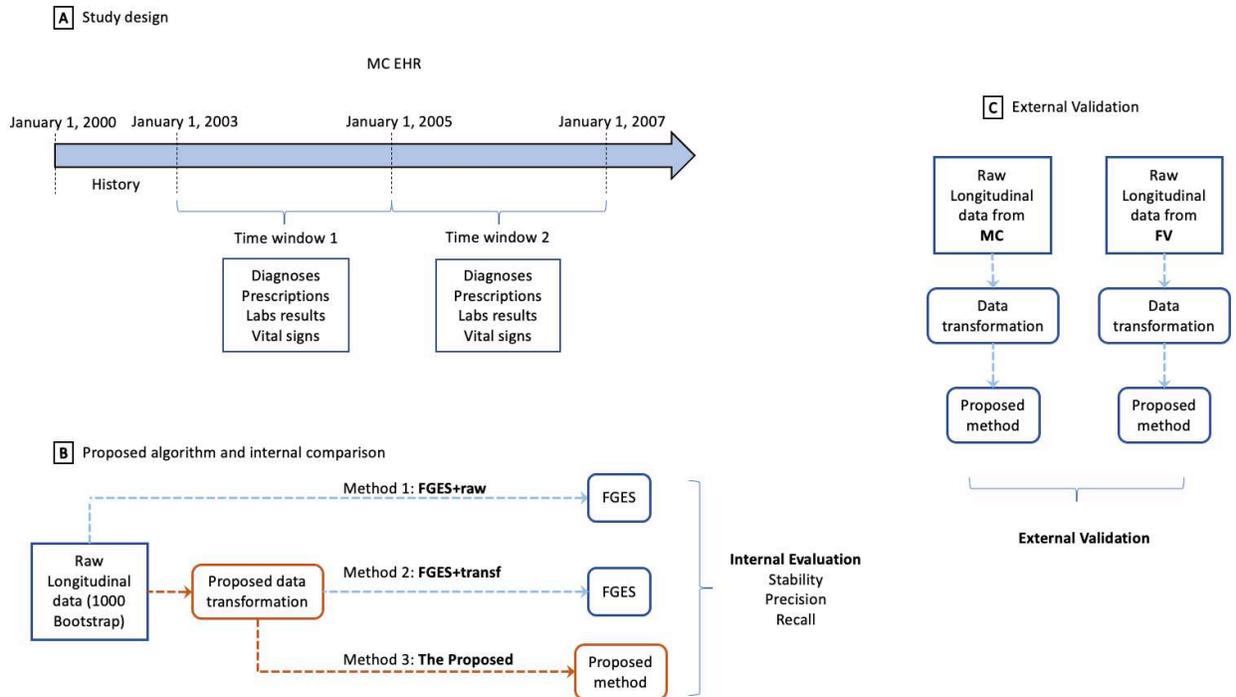


Figure 1. Study design and evaluations

A. Overview of the study design for Mayo Clinic (MC) EHR. **B.** The workflow of the internal evaluation. Three methods FGES+raw, FGES+transf, and the proposed algorithm were compared using stability, precision, and recall. Orange color highlights the proposed method (Method 3). **C.** The workflow of external comparison. The proposed method was applied to two datasets, MC and M Health Fairview (FV), and the resulting graphs were compared.

Variables

Diagnosis codes are aggregated into the disease categories of obesity, hyperlipidemia, pre-diabetes, type 2 diabetes mellitus, coronary artery disease, myocardial infarction, heart failure, chronic renal failure, cerebrovascular disease, and stroke based on ICD-9 and codes following our previous work[15]. Medications indicated for the above conditions were rolled up into NDF-RT therapeutic subclasses. Relevant laboratory results and vital signs were categorized based on cutoffs from the American Diabetes Association guidelines[16].

Causal Structure Discovery

A relationship between two events is **causal** if manipulating the earlier event causes the other (later) event to change. For example, prescribing a medication reduces the probability of downstream events (complications). Causation differs from association. For example, blood sugar is associated with risk of stroke: diabetic patients with higher blood sugar have a higher risk of stroke; however, this relationship is likely not causal in diabetic patients since attempts to reduce the risk of stroke by reducing blood sugar consistently failed in clinical trials[17 18]. If two events share a common cause (a **confounder**) and are not otherwise causally related, then manipulating one event will not affect the other variable as long as the common cause remains unchanged. The confounder can be observed or latent. The term **causal structure** refers to the set of all existing causal relationships among all events and can be visualized as a graph. The causal graph consists of nodes, which corresponds to events, and the nodes are connected

by edges that denote causal relationships. General-purpose CSD methods are designed to work with observational data to derive a causal structure that are consistent with the joint probability of the data.

Several general-purpose CSD algorithms have been proposed and the interested reader is referred to the Supplements II where we present an overview of the major methods. In this work, we focus on (Fast) Greedy Equivalence Search (FGES) as the comparison method, because we previously found it to outperform other CSD methods[19]. Briefly, FGES finds the optimal causal graph by a greedy search guided by a goodness-of-fit score (e.g. BIC or BDeu) over all possible graphs. Particularly, it starts with an empty graph, and iteratively adds individual edges that maximize the score given the current graph, until adding edges no longer improves the score. Then, FGES iteratively removes individual edges that maximizes the score, until edge removal ceases to improve the score. The output of FGES is a pattern, which can contain undirected edges, where the causal effect direction could not be determined due to statistical equivalency. FGES has good mathematical properties and been shown to be consistent under a set of assumptions[14 20].

Proposed Methods

The workflow of the proposed methods is described in Figure 1B, method 3 (colored in orange). We propose two methods, a data transformation and a causal search method. The former method transforms the longitudinal EHR data into **disease-related events**, so that we can determine the temporal ordering of events (diseases) despite inaccuracies in the EHR data and extracts all pairs of diseases where a clear precedence ordering exists. The search method constructs the causal graph using the transformed data and the set of precedence pairs.

Data Transformation Method

A disease-related **event** is defined as a diagnosis, a prescription, an abnormal lab result, or abnormal vital sign. An event is **incident** if it occurs in the second time window but is not present in the first time window although the patient is observed in the first time window. Conversely, a disease event is **pre-existing** if the patient presented with it in or before the first time window. An event A **precedes** another event B if among patients who have both A and B in the second time window, B is significantly more likely to be incident than A . Note that precedence implies neither causation nor association; however, if a causal effect exists, it must follow the precedence direction. Formal mathematical definitions of these concepts can be found in the Supplement I. The output from this step is (i) an event-based data set consisting of the incident and pre-existing conditions for each patient in each of the two time windows, (ii) a set \mathcal{C} of precedence relationships of all pairs (v_i, v_j) of events for which event v_i clearly precedes v_j .

The proposed CSD Search Algorithm

Given \mathcal{C} , we construct the causal graph \mathcal{G} by iteratively adding edge (v_i, v_j) from \mathcal{C} that maximizes the goodness of fit of \mathcal{G} . The orientation of this edge must be consistent with the precedence relationship, namely from v_i to v_j . The goodness of fit is defined by the BIC criteria. Let $X^{(1)}, X^{(2)}$ denote the data sets collected in the two distinct time windows, where $X^{(2)}$ follows $X^{(1)}$. The likelihood of the \mathcal{G} is

$$\begin{aligned} \mathcal{L}(\mathcal{G}|X^{(1)}, X^{(2)}) &= P(X^{(2)}, X^{(1)} | \mathcal{G}) \\ &= P(X^{(2)} | X^{(1)}, \mathcal{G}) P(X^{(1)} | \mathcal{G}) \\ &= \prod_s \prod_{v \in \mathcal{V}} P(v_s^{(2)} | x_s^{(1)}, \mathcal{G}) P(x_s^{(1)} | \mathcal{G}) \end{aligned}$$

$$= \prod_s \prod_{v \in \mathcal{V}} P(v_s^{(2)} | pa(v, \mathcal{G})_s^{(1)}) P(x_s^{(1)} | \mathcal{G}), \quad (4)$$

where $x_s^{(t)}$ is the observation vector for subject s at the cross-section t ; $v_s^{(t)}$ is the observation of variable (event) v for subject s at the cross-section t ; and $pa(v, \mathcal{G})_s^{(1)}$ is the observation vector for the parents of v in the causal structure \mathcal{G} , at cross section 1 for subject s .

The algorithm estimates $P(v_s^{(2)} | pa(v, \mathcal{G})_s^{(1)})$ using logistic regression on the subjects that do not have v at the first cross section and are under observations for both cross sections. For subjects who have v at the first cross section, the probability of having v at the second cross section is 1. Since \mathcal{G} represents the transition graph, the term $P(x_s^{(1)} | \mathcal{G})$ is a constant.

Finally, the BIC score is

$$\text{BIC}(\mathcal{G}) = -2\ln\mathcal{L}(\mathcal{G}|X^{(1)}, X^{(2)}) + \ln(n)|\mathcal{G}|, \quad (5)$$

where n is the number of observations that are common in the two cross sections, and $|\mathcal{G}|$ is the number of edges in the causal structure \mathcal{G} .

Algorithm 1: The proposed causal search algorithm.

Input: the precedence set \mathcal{C} ; the data set X
Output: the discovered DAG, \mathcal{G}

```

1 Initialize  $\mathcal{G} = \emptyset$ 
2 while the candidate set  $\mathcal{C}$  is not empty and  $\Delta\text{BIC} > 0$  do
3   foreach candidate pair  $(V_i < V_j) \in \mathcal{C}$  do
4      $\Delta\text{BIC} = \text{BIC}(X, \mathcal{G}) - \text{BIC}(X, \mathcal{G} \cup (V_i < V_j))$ 
5     if  $\Delta\text{BIC} > 0$  then
6        $\mathcal{G} = \mathcal{G} \cup (V_i < V_j)$  // add the candidate pair on the current graph
7        $\mathcal{C} = \mathcal{C} \setminus (V_i < V_j)$  // Remove the candidate pair from  $\mathcal{C}$  else stop
8     end
9   end
10 end
11 Return  $\mathcal{G}$ 

```

Algorithm 1. The proposed causal search algorithm.

Algorithm 1 describes the proposed algorithm for constructing the causal graph \mathcal{G} . \mathcal{G} is a directed acyclic graph (DAG), with nodes representing variables and edges representing causal effects between a pre-existing and an incident variable.

Evaluation

Clinical evidence

The standard way to evaluate CSD methods is to compare the resulting graph to a gold standard graph. However, such a gold standard graph does not exist and possibly many relationships are unknown.

However, there exists (i) Associative Evidence: a large body of observational studies documenting risk factors and outcomes for diabetes. Results from these studies have already been distilled into summaries.[21] (ii) Clinical trials can support both the existence (positive) and also the lack (negative) of hypothesized causal relationships. We compiled a list of causal relationships from clinical trials considering 175 clinical trials with a primary endpoint of any of the conditions we studied, including composite end points. We excluded trials with inclusion criteria that are too strict (trial results would not generalize to our population) and the interventions that are out of the scope of our study. 14 trials remained yielding 19 positive and 18 negative causal relationships. These trials and the evidence they produced are listed in Supplement III, Table S1. These relationships are used as causal evidence to compute recall.

Internal Evaluation

We evaluated the method and the resulting graphs from the following four perspectives.

Stability. We run 1000 bootstrap replicas on the development cohort. An edge has ambiguous orientation if it is present in at least half of the 1000 graphs (edge is not noise) and both orientations appear in at least 30% of the graphs that contain this edge (it does not have a dominant direction). We report the percentage of ambiguous edges.

Precision. Based on the causal graph derived from the training cohort, an edge is **incorrect** if there is no associative evidence of a relationship between the two events; or if causal evidence specifically indicates the lack of a causal relationship. We define **precision** as one minus the proportion of incorrect edges among the discovered edges.

Causal recall. Causal recall is computed on a single graph discovered from the training cohort, quantifying the percentage of the known causal relationships discovered. A known causal relationship from A to B is discovered if there is a node in the graph that maps to A , another node that maps to B and (a) a direct causal relationship $A \rightarrow B$ in the graph exists or (b) a causal path $A \rightarrow X \rightarrow B$ exists and no causal evidence states that in patients with X , A does not cause B . For example, if the evidence states that blood pressure (without specifying whether it is systolic or diastolic) increases the risk of stroke, then the path `sbp→cevd→stroke` would satisfy this relationship.

Associative recall. Associative recall is also computed on a single graph discovered from the training cohort and it quantifies the percentage of known associative relationships that can be explained by the discovered causal graph. An associative relationship between A and B is explained by the graph if there is a node in the graph that maps to A , another node that maps to B , and a path between A and B exists in the graph.

External Validation

We performed 1000 bootstrap replications on both data sets independently using the proposed method. On each data set, all edges from the 1000 graphs were pooled, resulting in two sets of pooled edges. We compared these two sets and pointed out the edges that were discordant between the MC and FV data, as shown in Figure 1C.

Method Comparison

Figure 1B depicts an overview of the method comparison. Three methods are compared, (i) **FGES+raw**: FGES is applied directly to the raw data; (ii) **FGES+transf**: data is transformed using the proposed transformation method and FGES is applied to the transformed data; and (iii) **Proposed**: the proposed

search algorithm is applied to the transformed data. Comparing FGES+raw and FGES+transf isolates the effect of the proposed transformation method, and comparing FGES+transf and Proposed highlights the effect of the proposed search algorithm.

Results

Baseline characteristics

Table 1 presents descriptive statistics for the MC and FV data sets at the end of the first time window and incidence rates for the diseases in the second window. Differences between datasets are tested through the t-test (for age) and the chi-square test (all other variables).

Table 1. Characteristics of the MC and FV data sets. For age, mean (sd) is indicated; for the disease-related events, percentage (%) of positive is indicated. New events rate at the second time windows is reported.

	MC (N = 57332)		FV (N = 79486)		P-val
	events in window 1	new events in window 2	events in window 1	new events in window 2	
Demographics					
age	48.1 (18.2)		50.4 (14.6)		0.000
male	0.43		0.34		0.000
Vitals & labs					
BMI ≥ 25 & ≤ 30	27.1	2.9	27.5	3.4	0.097
BMI ≥ 30	32.6	3.6	43.1	4.9	0.000
SBP ≥ 140	10.3	3.4	4.5	2.9	0.000
DBP ≥ 90	2.3	1.0	1.6	1.2	0.000
LDL ≥ 130	18.4	3.6	15.4	4.3	0.000
HDL abnormal	20.2	1.7	24.6	3.0	0.000
Triglyceride ≥ 150	22.6	3.7	17.6	4.3	0.000
FPG 100-125	24.4	7.2			
FPG > 125	11.9	3.7			
A1C > 5.7 and A1C < 6.5			6.8	0.6	
A1C ≥ 6.5			7.0	0.9	
Diagnoses					
Hypertension	28.4	5.6	30.6	8.4	0.000
Obesity	11.5	1.2	11.3	1.3	0.320

Hyperlipidemia	31.9	8.3	36.4	9.4	0.000
Pre-diabetes mellitus	0.9	3.5	2.4	2.4	0.000
Diabetes mellitus	7.9	5.1	9.5	4.3	0.000
Chronic renal failure	1.2	0.2	0.2	0.3	0.000
Coronary heart failure	2.4	1.7	1.2	1.4	0.000
Coronary artery disease	9.4	3.5	5.6	3.4	0.000
Myocardial infarction	2.4	1.2	0.9	1.6	0.000
Cerebrovascular disease	3.6	2.3	1.8	1.4	0.000
Stroke	1.2	1.1	0.6	1.0	0.000
Treatments					
Hypertension	20.6	8.3	31.5	13.9	0.000
Hyperlipidemia	15.7	8.0	24.6	9.1	0.000
Diabetes mellitus	4.4	2.4	7.2	4.3	0.000

Table 2. Directional stability. The table shows the number of distinct edges that appeared in half of the 1000 bootstrap replications, and the percentage of ambiguously oriented edges.

Method	Number of distinct edges	Ambiguously oriented
FGES+raw	125	45%
FGES+transf	75	24%
Proposed	64	0%

Directional stability

The proposed data transformation reduced the percentage of ambiguously oriented edges from 45% to 24%, and finally, the proposed search method eliminated ambiguously oriented edges (Table 2).

Table 3. Metrics from clinical evidence

	Precision	Associative Recall	Causal Recall
1. FGES+raw	0.294	1.000	1.000
2. FGES+transf	0.549	0.985	1.000

3. Proposed

0.838

1.000

0.947

Correctness and completeness

Table 3 shows the precision, associative recall and causal recall of the graphs discovered by the three methods. All three methods achieved almost perfect recall; FGES+raw achieved the lowest precision of 0.294: less than third of the events reported as causally related are even associated. By using the proposed transformation, the precision increased to 0.55, but almost half of the reported causal relationships are still incorrect. Finally, the proposed method achieved a precision of 0.838. We present the causal graph discovered by the proposed methods in the Supplement IV, Figure S1.

External validation

We compared the graphs discovered from the MC and FV data sets. There are 74 distinct edges that were discovered from at least one of the data sets. Sixty (81%) edges coincided across the two datasets, while 14 (19%) differed. Table 4 shows the discordant edges, the percentage of bootstrap iterations in which the edge was present and the main reason for the discordance.

Table 4. External Validation. The table describes the edges that were discordant between the Mayo Clinic (MC) and M Health Fairview (FV) data sets. It shows the percentage of the bootstrap iterations in which the edge was discovered at MC and FV and a brief reason for the discrepancy.

Edge	Discovery %		Reason
	MC	FV	
hdl→trigl	0	91.7	There is no clear precedence relationship, the two events often coincide.
htn.dx→crf	88.5	0.1	
trigl→dm.dx	100	0	
trigl→hl.tx	100	0	
ldl→hl.dx	72.1	0	
fasting.125→dm.dx	100	0	FV uses A1C, not FPG.
trigl→fasting.125	99.5	0.2	
dbp→htn.tx	91.5	0	The criteria for diagnosis and treatment are institution specific.
sbp→hl.tx	99.3	1.7	
sbp→htn.tx	100	29.1	
trigl→htn.tx	83.7	0	
chf→mi	0	67.6	SBP is a common cause for CHF and MI, but at FV, this effect was too weak in 68% of the bootstrap iterations
hl.dx→trigl	0	87.6	While the main driver of TG is BMI, at FV, the diagnosis of HL helps explain the variation in TG.
hl.tx→cad	0	74.3	LDL drives both HL treatment and CAD.

There are three broad reasons for differences in edges. The main reason, affecting half of the edges was that of policy differences. These include preferred lab results (A1C vs FPG) and decisions regarding therapeutic interventions. The second reason, affecting four edges, is a lack of clear precedence in the relationships among the events. For example, the abnormal TG→HL treatment edge was not discovered at FV because the first abnormal TG precedes or follows the HL treatment in statistically equal proportions. The final reason, affecting the remaining three edges, is differential degree of confounding

between the two sites. For example, SBP is a confounder of CHF and MI. When the algorithm fails to detect the SBP→MI edge, the effect of SBP on MI was shown through CHF (which depends on SBP more than MI). For the HL diagnosis→TG edge, the common cause is BMI, and for the HL treatment→CAD edge, it is LDL. The reason for differential confounding was likely a combination of population and institutional differences as well as data artifacts.

Discussion

We proposed a new data transformation method and a new search algorithm specifically designed for EHR data, and showed how the resulting graph achieved close to 90% precision (90% of the edges were correct), almost 100% recall (the graph could explain all known associations and almost all known causal relationships), and the graph was remarkably stable in face of data perturbation (no edge disappeared or changed direction). Due to its built-in facility to handle inaccuracies in the EHR data and study design considerations, our method outperformed general purpose methods by a large margin.

While the two graphs from the two independent health systems are reassuringly similar, small differences exist. None of these differences implies an incorrect physiological or pathophysiological effect. Among the 14 edges that differed, seven captured differences between the population and the institutions, such as institution-specific triggers for prescriptions and the use of different laboratory tests for the same purpose (fasting plasma glucose versus A1c). Depending on the goal of the modeling, it may be desirable to include such differences. We believe that the discovered causal graphs offer adequate information about causal (including confounding) factors to support the development of clinical decision support models and can also support clinical research efforts.

Generalizability beyond diabetes. The proposed method was demonstrated on type 2 diabetes, but it can generalize to other applications as long as the target application benefits from some of the improvements: reducing the impact of inaccuracies in the EHR data, accounting for the temporal ordering of events and distinguishing pre-existing and incident conditions.

Future work: The algorithm requires longitudinal data with at least two time-points. Different disease and their symptoms might manifest at different rates, incorporating this knowledge into the discovery process may enhance the performance of the algorithms. Further, variable semantics (such as SBP and DBP being measures related to hypertension) is an essential component of the proposed algorithm, but it is not always available in a computable form. Further, both datasets in this study are from the Midwest with a predominantly white patient population. The generalizability of the discovered causal relations can be further tested by examining a broader patient population.

Conclusions

We have demonstrated that the graph produced by the proposed transformation and search algorithm is more stable across bootstrap iterations and as complete as other methods yet it contained substantially fewer errors (had higher precision) than graphs produced by general-purpose methods. The resulting graph was successfully validated using longitudinal EHR data from an independent health system. We conclude that the proposed method is more suitable for use in clinical studies using EHR data.

Funding acknowledgement

This work is supported by the National Institutes of Health (NIH) grants AG056366, TR002494, and LM011972. Contents of this document are the sole responsibility of the authors and do not necessarily represent official views of the NIH.

Author contributions

G.S. and S.M. designed the study. X.S. conducted the analyses and drafted the paper. G.S. and S.M. designed and oversaw the statistical analysis. P.V., R.C., and P.C. supervised and contributed to the data interpretation. All authors reviewed the manuscript and provided suggestions.

Conflict of interest statement

The authors have no competing interests to declare.

Data availability

The data that support the findings of this study are not publicly available since they contain patient health information. Authorization to access patient data can be requested from the Mayo Clinic and University of Minnesota Institutional Review Board.

Statement of human rights and informed consent

The study was approved by both Mayo Clinic and University of Minnesota Institutional Review Board (IRB). Informed consent was obtained from all patients. All relevant guidelines and regulations were followed.

Reference

1. Loh E. Medicine and the rise of the robots: a qualitative review of recent advances of artificial intelligence in health. *BMJ Leader* 2018;leader-2018-000071 doi: 10.1136/leader-2018-000071.
2. Semigran HL, Levine DM, Nundy S, Mehrotra A. Comparison of Physician and Computer Diagnostic Accuracy. *JAMA Internal Medicine* 2016;**176**(12):1860-61 doi: 10.1001/jamainternmed.2016.6001.
3. Beam AL, Kohane IS. Big Data and Machine Learning in Health Care. *JAMA* 2018;**319**(13):1317-18 doi: 10.1001/jama.2017.18391.
4. Trister AD, Buist DSM, Lee CI. Will Machine Learning Tip the Balance in Breast Cancer Screening? Will Machine Learning Tip the Balance in Breast Cancer Screening? Will Machine Learning Tip the Balance in Breast Cancer Screening? *JAMA Oncology* 2017;**3**(11):1463-64 doi: 10.1001/jamaoncol.2017.0473.
5. Ashley EA. The Precision Medicine Initiative: A New National Effort. *JAMA* 2015;**313**(21):2119-20 doi: 10.1001/jama.2015.3595.
6. Friedman CP, Wong AK, Blumenthal D. Achieving a Nationwide Learning Health System. *Science Translational Medicine* 2010;**2**(57):57cm29 doi: 10.1126/scitranslmed.3001456.
7. Mukherjee S. A.I. Versus M.D. What happens when diagnosis is automated? . *The New Yorker*, 2017.
8. Segura-Egea JJ, Cabanillas-Balsera D, Jiménez-Sánchez MC, Martín-González J. Endodontics and diabetes: association versus causation. *International Endodontic Journal* 2019;**52**(6):790-802 doi: 10.1111/iej.13079.
9. Li Y, Torralba A, Anandkumar A, Fox D, Garg A. Causal Discovery in Physical Systems from Videos. *NeurIPS*, 2020.

10. Anker JJ, Kummerfeld E, Rix A, Burwell SJ, Kushner MG. Causal Network Modeling of the Determinants of Drinking Behavior in Comorbid Alcohol Use and Anxiety Disorder. *Alcoholism: Clinical and Experimental Research* 2019;**43**(1):91-97 doi: <https://doi.org/10.1111/acer.13914>.
11. Ebert-Uphoff I, Deng Y. Causal Discovery for Climate Research Using Graphical Models. *Journal of Climate* 2012;**25**(17):5648-65 doi: 10.1175/jcli-d-11-00387.1.
12. Pearl J. *Causality*. Cambridge: Cambridge University Press, 2009.
13. Meek C. Causal Inference and Causal Explanation with Background Knowledge. *Proceedings of the Eleventh conference on Uncertainty in artificial intelligence* 1995:403-10
14. Ramsey JD. Scaling up Greedy Equivalence Search for Continuous Variables. *CoRR* 2015;**abs/1507.07749**.
15. Ngufor C, Caraballo PJ, O'Byrne TJ, et al. Development and Validation of a Risk Stratification Model Using Disease Severity Hierarchy for Mortality or Major Cardiovascular Event. *JAMA network open* 2020;**3**(7):e208270-e70 doi: 10.1001/jamanetworkopen.2020.8270.
16. Standards of Medical Care in Diabetes—2019 Abridged for Primary Care Providers. *Clinical Diabetes* 2019;**37**(1):11 doi: 10.2337/cd18-0105.
17. Effects of Intensive Glucose Lowering in Type 2 Diabetes. *New England Journal of Medicine* 2008;**358**(24):2545-59 doi: 10.1056/NEJMoa0802743.
18. Intensive blood-glucose control with sulphonylureas or insulin compared with conventional treatment and risk of complications in patients with type 2 diabetes (UKPDS 33). *The Lancet* 1998;**352**(9131):837-53 doi: 10.1016/S0140-6736(98)07019-6.
19. Shen X, Ma S, Vemuri P, et al. Challenges and Opportunities with Causal Discovery Algorithms: Application to Alzheimer's Pathophysiology. *Scientific Reports* 2020;**10**(1):2975 doi: 10.1038/s41598-020-59669-x.
20. Chickering DM. Optimal Structure Identification With Greedy Search. *Journal of Machine Learning Research* 2002;**3**:507-54.
21. Mayo Clinic Patient Care and Health Information. <https://www.mayoclinic.org/patient-care-and-health-information>. Accessed Nov 21, 2020.
22. Spirtes P, Clark G, Scheines R. *Causation, Prediction, and Search*: The MIT Press, 2000.
23. Peter S, Christopher M, Thomas R. Causal inference in the presence of latent variables and selection bias. *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence* 1995:499–506
24. Colombo D, Maathuis MH, Kalisch M, Richardson TS. LEARNING HIGH-DIMENSIONAL DIRECTED ACYCLIC GRAPHS WITH LATENT AND SELECTION VARIABLES. *The Annals of Statistics* 2012;**40**(1):294-321
25. Tsamardinos I, Brown LE, Aliferis CF. The max-min hill-climbing Bayesian network structure learning algorithm. *Machine Learning* 2006;**65**(1):31-78 doi: 10.1007/s10994-006-6889-7.
26. Ogarrio JM, Spirtes P, Ramsey J. A Hybrid Causal Search Algorithm for Latent Variable Models. In: Alessandro A, Giorgio C, Cassio Polpo C, eds. *Proceedings of the Eighth International Conference on Probabilistic Graphical Models*. *Proceedings of Machine Learning Research: PMLR*, 2016:368--79.
27. Hoyer PO, Janzing D, M. MJ, Jonas P, Bernhard SI. Nonlinear causal discovery with additive noise models. 2009:689--96
28. Sanchez-Romero R, Ramsey JD, Zhang K, Glymour MRK, Huang B, Glymour C. Estimating feedforward and feedback effective connections from fMRI time series: Assessments of statistical methods. *Network neuroscience (Cambridge, Mass.)* 2019;**3**(2):274-306 doi: 10.1162/netn_a_00061.
29. Jabbari F, Visweswaran S, Cooper GF. Instance-Specific Bayesian Network Structure Learning. In: Václav K, Milan S, eds. *Proceedings of the Ninth International Conference on Probabilistic Graphical Models*. *Proceedings of Machine Learning Research: PMLR*, 2018:169--80.

30. Jabbari F, Cooper G. An Instance-Specific Algorithm for Learning the Structure of Causal Bayesian Networks Containing Latent Variables, 2020:433-41.
31. Zheng X, Aragam B, Ravikumar PK, P. XE. DAGs with NO TEARS: Continuous Optimization for Structure Learning. 2018:9472--83
32. Yu Y, Chen J, Gao T, Yu M. DAG-GNN: DAG Structure Learning with Graph Neural Networks. In: Kamalika C, Ruslan S, eds. Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research: PMLR, 2019:7154--63.
33. Zheng X, Dan C, Aragam B, Ravikumar P, Xing E. Learning Sparse Nonparametric DAGs. In: Silvia C, Roberto C, eds. Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research: PMLR, 2020:3414--25.
34. Lachapelle S, Brouillard P, Deleu T, Lacoste-Julien S. *Gradient-Based Neural DAG Learning*, 2019.
35. Young J, Andrews B, Cooper G, Lu X. *Learning Latent Causal Structures with a Redundant Input Neural Network*, 2020.
36. Bellot AavdSM. Conditional Independence Testing using Generative Adversarial Networks. 2019:2202--11
37. Goudet O, Kalainathan D, Caillou P, Guyon I, Lopez-Paz D, Sebag M. Learning Functional Causal Models with Generative Neural Networks. In: Escalante HJ, Escalera S, Guyon I, et al., eds. Explainable and Interpretable Models in Computer Vision and Machine Learning. Cham: Springer International Publishing, 2018:39-80.
38. Wang Y, Menkovski V, Wang H, Du X, Pechenizkiy M. *Causal Discovery from Incomplete Data: A Deep Learning Approach*, 2020.
39. Zhu S, Ng I, Chen Z. Causal Discovery with Reinforcement Learning. International Conference on Learning Representations 2020
40. Huang X, Zhu F, Holloway L, Haidar A. *Causal Discovery from Incomplete Data using An Encoder and Reinforcement Learning*, 2020.
41. Designing Optimal Dynamic Treatment Regimes: A Causal Reinforcement Learning Approach. ICML 2020; 2020.
42. Effects of an Angiotensin-Converting–Enzyme Inhibitor, Ramipril, on Cardiovascular Events in High-Risk Patients. New England Journal of Medicine 2000;**342**(18):1376-76 doi: 10.1056/NEJM200005043421819.
43. Rubins HB, Robins SJ, Collins D, et al. Gemfibrozil for the Secondary Prevention of Coronary Heart Disease in Men with Low Levels of High-Density Lipoprotein Cholesterol. New England Journal of Medicine 1999;**341**(6):410-18 doi: 10.1056/nejm199908053410604.
44. Effects of Intensive Blood-Pressure Control in Type 2 Diabetes Mellitus. New England Journal of Medicine 2010;**362**(17):1575-85 doi: 10.1056/NEJMoa1001286.
45. Effects of Intensive Glucose Lowering in Type 2 Diabetes. New England Journal of Medicine 2008;**358**(24):2545-59 doi: 10.1056/NEJMoa0802743.
46. Intensive blood-glucose control with sulphonylureas or insulin compared with conventional treatment and risk of complications in patients with type 2 diabetes (UKPDS 33). The Lancet 1998;**352**(9131):837-53 doi: 10.1016/S0140-6736(98)07019-6.
47. Effects of Combination Lipid Therapy in Type 2 Diabetes Mellitus. New England Journal of Medicine 2010;**362**(17):1563-74 doi: 10.1056/NEJMoa1001282.
48. Cardiovascular Effects of Intensive Lifestyle Intervention in Type 2 Diabetes. New England Journal of Medicine 2013;**369**(2):145-54 doi: 10.1056/NEJMoa1212914.
49. Collins R, Armitage J, Parish S, Sleight P, Peto R. MRC/BHF Heart Protection Study of cholesterol lowering with simvastatin in 20 536 high-risk individuals: A randomised placebo-controlled trial. The Lancet 2002;**360**:7-22 doi: 10.1016/S0140-6736(02)09327-3.

50. Scott R, Brien R, Fulcher G, et al. Effects of Fenofibrate Treatment on Cardiovascular Disease Risk in 9,795 Individuals With Type 2 Diabetes and Various Components of the Metabolic Syndrome. *Diabetes Care* 2009;**32**(3):493 doi: 10.2337/dc08-1543.
51. Intensive Blood Glucose Control and Vascular Outcomes in Patients with Type 2 Diabetes. *New England Journal of Medicine* 2008;**358**(24):2560-72 doi: 10.1056/NEJMoa0802987.
52. The Diabetes Control and Complications Trial Research G. The Effect of Intensive Treatment of Diabetes on the Development and Progression of Long-term Complications in Insulin-dependent Diabetes Mellitus. *RETINA* 1994;**14**(3)
53. Hsu C-y. Benazepril was effective and safe for advanced chronic kidney disease without diabetes. *ACP Journal Club* 2006;**145**(1):19 doi: 10.7326/ACPJC-2006-145-1-019.
54. Brenner BM, Cooper ME, de Zeeuw D, et al. Effects of Losartan on Renal and Cardiovascular Outcomes in Patients with Type 2 Diabetes and Nephropathy. *New England Journal of Medicine* 2001;**345**(12):861-69 doi: 10.1056/NEJMoa011161.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [Supplements.docx](#)