

Visual Analysis of Multi-outcome Causal Graphs

Mengjie Fan, Jinlu Yu, Daniel Weiskopf, Nan Cao, Huai-Yu Wang, and Liang Zhou*

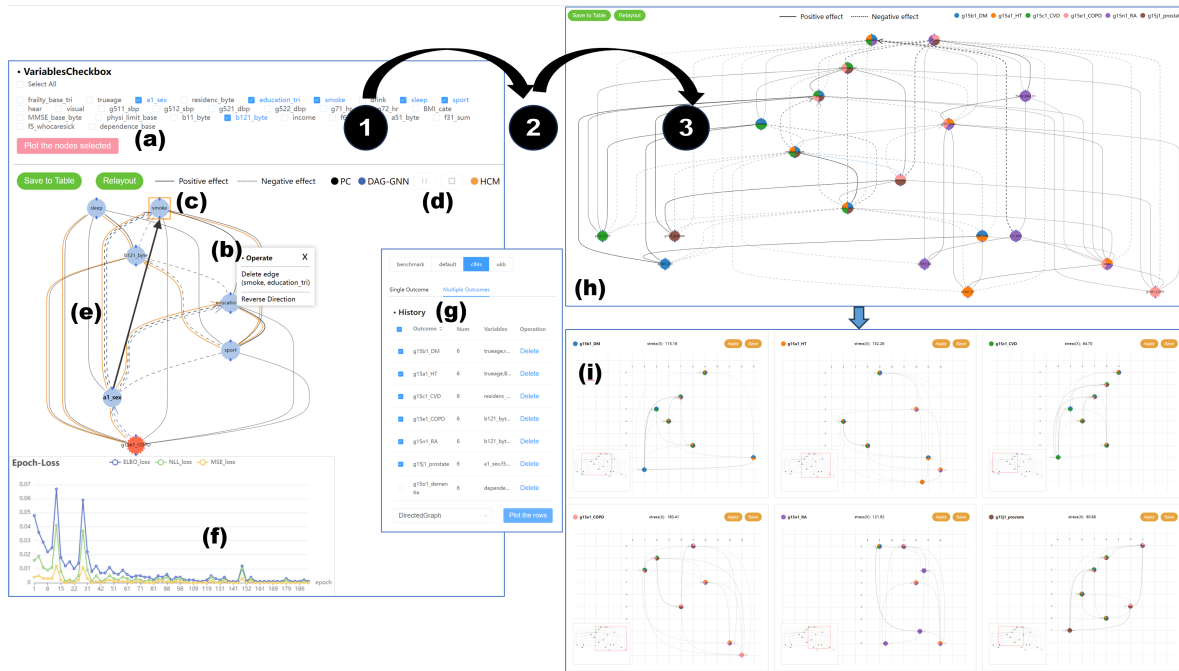


Fig. 1: The overall process of the visual analysis of multi-outcome causal graphs of a health research dataset [15]. Analysts use (1) the Single Graph View to analyze and edit the causal graph (a–c) of a single outcome aided by a set of new visualization techniques (d–f) to leverage the synergy of various causal discovery methods. The fine-tuned causal graphs of single outcomes are stored in (2) the History View (g). Next, causal graphs of interest are selected for visual comparison in (3) the Multi-outcome Graphs Comparison View (h, i), supported by our new graph layout and visual mappings (i).

Abstract—We introduce a visual analysis method for multiple causal graphs with different outcome variables, namely, multi-outcome causal graphs. Multi-outcome causal graphs are important in healthcare for understanding multimorbidity and comorbidity. To support the visual analysis, we collaborated with medical experts to devise two comparative visualization techniques at different stages of the analysis process. First, a progressive visualization method is proposed for comparing multiple state-of-the-art causal discovery algorithms. The method can handle mixed-type datasets comprising both continuous and categorical variables and assist in the creation of a fine-tuned causal graph of a single outcome. Second, a comparative graph layout technique and specialized visual encodings are devised for the quick comparison of multiple causal graphs. In our visual analysis approach, analysts start by building individual causal graphs for each outcome variable, and then, multi-outcome causal graphs are generated and visualized with our comparative technique for analyzing differences and commonalities of these causal graphs. Evaluation includes quantitative measurements on benchmark datasets, a case study with a medical expert, and expert user studies with real-world health research data.

Index Terms—Causal graph visualization and visual analysis, causal discovery, comparative visualization, visual analysis in medicine.

1 INTRODUCTION

- M. Fan is with the Institute of Medical Technology, Peking University Health Science Center, and the National Institute of Health Data Science (NIHDS), Peking University. E-mail: mengjiefan@bjmu.edu.cn.
- J. Yu is with Chalmers University of Technology and NIHDS, Peking University. E-mail: yjl.inlp@gmail.com
- L. Zhou is with NIHDS, Peking University. L. Zhou is the corresponding author. E-mail: zhouling@pku.edu.cn.
- D. Weiskopf is with the Visualization Research Center (VISUS), University of Stuttgart. E-mail: weiskopf@visus.uni-stuttgart.de
- N. Cao is with the iDV^x Lab, Tongji University. E-mail: nan.cao@gmail.com
- H. Wang is with the National Institute of Traditional Chinese Medicine Constitution and Preventive Treatment of Diseases, Beijing University of Chinese Medicine. E-mail: wanghuaiyuelva@126.com

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on

Understanding causality is fundamental in human intelligence [17] and is crucial in many applications, for example, guiding practitioners designing and implementing public interventions in policy-making [13], assisting legal case matching in legal processes [56], predicting how the molecular system responds to different interventions in biology [59], and assisting disease diagnosis in medicine [48]. Causal discovery refers to the process of learning graphical structures with a causal interpretation [68] and has become one of the core research topics of statistics, machine learning, and artificial intelligence [30, 50, 68]. These graphical structures are directed acyclic graphs (DAGs), namely causal graphs, where vertices are data variables and directed edges denote causal relationships between variables.

In medicine, experts also seek to discover causalities between po-

obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxx

tential risk factors and a specific disease, aiming to explore factors influencing the development of a particular disease and eventually intervene in the prognosis. Except for the typical case of studying the causality of a single disease (a single outcome variable), understanding the similarities and differences between different yet related causal relationships among different diseases, i.e., multiple outcome variables, is also important. For example, a study into the co-occurrence of more than one chronic condition in one person (multimorbidity [28]) and the combined effect of multiple conditions in one person (comorbidity [60]) is beneficial for proper interventions and treatments. Therefore, it is necessary to perform analysis and comparisons for causality of multiple outcomes. Unfortunately, most previous visual analysis works intend to address problems with only one particular outcome.

In this paper, we study the causalities of multiple outcome variables that have not been considered in visual analytics. Based on a requirement analysis in close collaboration with medical experts, we propose a visual analysis framework as shown in Fig. 1. The framework consists of two stages: the discovery and interactive fine-tuning of a causal graph of a single outcome (Fig. 1(1)), and the comparative analysis of causal graphs of multiple outcomes (Fig. 1(3)).

Causal graphs in the first stage are calculated using three causal discovery algorithms: a classic causal discovery method that delivers a basic graph, and two state-of-the-art methods that provide analysts with more suggestions. The resulting graphs are visualized progressively as node-link diagrams with a superposed comparative visualization for easy detection of differences. The underlying machine learning methods can be interactively controlled, e.g., by pausing and resuming, to reduce computation time. Users then interactively fine-tune the single outcome causal graph with reference to multiple views of multivariate data visualizations and with their domain knowledge.

Next, the user selects causal graphs of outcomes of interest for the second stage (Fig. 1(g)). Individual graphs of single outcomes are then visualized as line-ups with a graph layout method and visual designs that facilitate comparisons (Fig. 1(i)). Our graph layout method preserves the relative position of shared nodes of supergraphs (Fig. 1(h)) calculated by combining multiple single graphs while compressing nodes that may be far apart. We use visual designs of glyphs, the shape of nodes, and grids to assist the comparison task.

We evaluate our method with numerical measurements, a case study, and an expert user study with medical experts. The effectiveness of the analysis and visualization of single outcome graphs is demonstrated using benchmark datasets; the multi-outcome comparison design is evaluated with stress measurements. The overall visual analysis method is evaluated through a case study of the analysis of diseases that are usually seen in multimorbidity/comorbidity using two widely accepted health research datasets. Feedback from expert users indicates that our method is useful for supporting causal reasoning of single and multiple outcomes and provides insights with the comparison of multi-outcome causal graphs that are impossible with existing methods.

Our major contributions are as follows.

- Two-stage visual analysis of multi-outcome causal graphs.
- Progressive visualization and interactive controls of multiple causal discovery techniques including modern continuous causal discovery methods.
- Graph layout and visual designs that support effective comparison of multiple DAGs.
- A case study using real-world health research datasets with a medical expert.

Please note that our method aims to address a domain problem in healthcare but is also general enough to aid causal analysis of single and multiple outcomes. The source code of our visual analysis system is publicly available at: https://github.com/mengjiefan/multi_outcome/tree/vis_rev_sub.

2 RELATED WORK

Our work is related to three areas of research: causal discovery, visualization and visual analysis of causality, and comparative visualization.

2.1 Causal Discovery

Causal reasoning is a challenging and active topic in machine learning, and details of this topic are covered in textbooks [45, 47]. Causal discovery is the process of finding causal graphs of variables in a given dataset [68]. The state-of-the-art of causal discovery is summarized in surveys [62, 68]. The methods can be categorized into constraint-based, score-based, structural asymmetries-based, and those exploiting various forms of intervention. In terms of how the problems are solved, two classes of techniques are available: traditional combinatoric approaches and modern continuous methods.

Combinatoric methods search for DAGs satisfying criteria within the aforementioned four categories. Popular combinatoric methods include the Peter-Clark (PC) algorithm [55], Greedy Equivalence Search (GES) [16], satisfiability (SAT) solver [32], and LinGAM [53].

In contrast, continuous methods rely on continuous optimizations that are typically used in neural networks. NO TEARS (Non-combinatoric Optimization via Trace Exponential Augmented Lagrangian Structure learning) is considered the first to formulate the combinatorial graph search method as a continuous optimization problem [70]. A further improvement of the method that is faster and more accurate than NO TEARS is available [9]. DAG-GNN [67] is a graph neural network approach to solve the continuous optimization problem. Methods focused on categorical data are also available with continuous optimizations using interventional or reinforcement learning [12, 37].

One issue with causal discovery methods is that they typically assume that all variables have the same data type—either categorical or continuous. In real-world datasets, however, data types are mixed, for example, typical survey data in medicine includes continuous biological measures like body height, weight, and blood pressure, and binary/categorical attributes like gender and stages of diseases. Few, dedicated methods are available for mixed-type datasets [3, 4, 14, 40], while some others, for example, DAG-GNN [67], can handle such data inherently. A latent-PC algorithm is available for multi-dimensional mixed-type data that is typical in medical research [14].

2.2 Visualization of Causal Graphs and Visual Analysis of Causal Reasoning

The visualization of causal graphs is studied in several works [6, 19, 61, 66]. Different visual mappings of the graph, e.g., adjacent matrix and node-link diagram are compared in a user study [61]. For node-link diagrams, different graph layouts are also compared with user studies [6, 61]. Layered graph layouts are used for showing large causal graphs [66] or comparing different causal discovery results [19].

Visual analysis methods for causal reasoning are available to aid the understanding of causalities and support interactive decision-making. An interactive causal reasoning interface allows visually editing and verifying causal relationships with multiple linked views of the causal graph, and statistical parameters [63]. A more comprehensive extension of the method is available with the support of multidimensional data visualizations, a visual design of the causal graph with layered graph layout, and the capability of handling mixed-type data [64]. An exploratory causal analysis approach introduces the uncertainty-aware visualization of causal graphs and provides interactive means for exploring, validating, and applying causal relations for decision-making [66].

Some works focus on the understanding of causal inferences and counterfactuals [38]. A multi-view visualization tool is available to support iterative causal inference [26]. Counterfactuals are used to improve the visualization of relationships between variables [36] and causal inferences [10]. Visual analytics of time-dependent causality can be applied to event sequences with a user-modifiable Granger causality model [34]. Causal hypotheses of time-dependent events can be formulated and tested with time delay in a visual analysis method [65].

Visual causal analysis is also used to address domain-specific problems. The interpretation of algorithmic decision-making models is supported by a visual analysis method that exploits the explanatory ability of causal models [31]. Algorithm bias is tackled with a causality-based visual analysis method that allows users to audit the data and generate debiased data [23]. Questionnaire responses are analyzed using causal reasoning for question combinations with association mining

and causal sub-graphs of each question combination are visualized and explored in a visual analysis system [39]. Spatial-temporal urban data is detected and analyzed using the Granger causality test in multiple linked views [20].

However, none of the above works utilizes state-of-the-art continuous optimization methods nor studies the comparison of causal graphs of multiple outcome variables in a dataset. Our method differs from these works in two important aspects: we use state-of-the-art causal discovery methods and address the problem of causal reasoning and visual comparison of multiple outcome variables that are important for public health.

2.3 Comparative Visualization

It is a common yet challenging task to design methods supporting comparison in visualization [24]. Four considerations are presented for comparative visualization designing and evaluation [24]. A general taxonomy of visual comparisons classifies visual designs into three basic categories—juxtaposition, superposition, and explicit encodings [25]. These categories should be combined for effective comparative visualization tools and tasks. A qualitative user study shows that interactive techniques are important to visual comparison tasks [58]. A combined encoding of colors and positions enables analysts to easily view differences and changes in the comparison of task-driven topic models [1].

Research efforts have been made for the visual comparison of graphs. Visual comparison of two graphs is achieved by visualizing their similarities and differences in a merged graph [5]. Interactive graph matching [29] addresses the visual comparison of graphs and clustered graphs. Weighted graph comparison is explored in brain connectivity analysis [2], where the effectiveness of node-link diagrams and adjacency matrices is compared. A map-based abstraction of node-aligned graphs on a triangle mesh facilitates a clutter-reduced comparison of undirected graphs [33]. Layout optimization, graph embedding, and clustering and grouping algorithms are used for the comparison of two temporal graph datasets [69].

Similar to these works, one of our goals is to visually compare graphs. However, existing works mainly address the comparison of two undirected graphs, whereas our method is designed to compare more than two directed graphs.

3 REQUIREMENT ANALYSIS AND METHOD OVERVIEW

Our work is a long-term collaboration with medical experts, which started in 2022. The work was initially motivated by the needs of one of our collaborators who studies the coexistence of multiple clinical conditions or diseases, such as multimorbidity or comorbidity, which are important topics in clinical medicine and public health [27, 57]. Mechanisms underlying multimorbidity or comorbidity are complex, they may share some direct causation, and associated risk factors like aging and socioeconomic deprivation [44, 54, 60], but the same factors may play different roles in different diseases. The expert hopes to analyze the influencing factors of different diseases and compare those diverse yet correlated diseases or clinical outcomes aiming to gain insights into the development of diseases and improve prognosis.

Traditionally, medical experts use statistical methods such as multiple logistic regression to explore factors associated with a disease. However, it is well-known that “correlation does not necessarily imply causation” [41], and, therefore, we considered using causal analysis to address this domain problem of analyzing multiple outcomes.

Following the design study methodology [52] and the abstraction according to a nested model for visualization design [43], we hosted one in-person meeting in the beginning (July 2022) and several follow-up online discussions (roughly every 2–3 months depending on the progress) with three clinical medicine and public health experts from the university that the first/corresponding author affiliates to (one of them is a co-author who is now at a different university). First, we abstracted the domain problem into two data analysis stages as shown in Fig. 2 that explains our workflow.

S1 Single causal graph exploration and analysis. Experts hope to analyze the causality between a specific clinical outcome and its potential influencing factors efficiently from a large dataset.

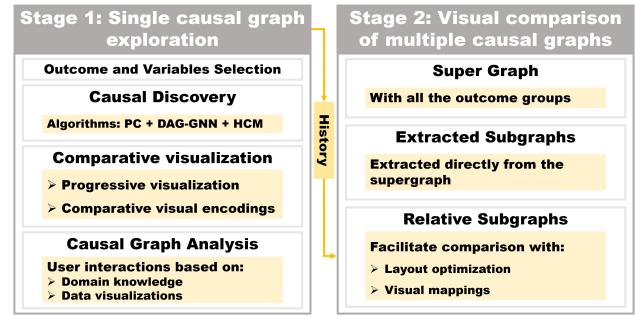


Fig. 2: The workflow of our visual analysis method.

The causal graph should be modified interactively by the user and the visualization should provide good interpretability for the underlying causal learning techniques.

S2 Multiple causal graph comparison and analysis. In this stage, experts hope to select interested clinical outcomes and compare causal graphs of the corresponding outcomes to guide interventions and treatments for different diseases.

We initially started with a design study on this specific healthcare problem and realized that aspects of this problem are general in causality visualization and visual analysis. For example, our problem involves the layout for causal graphs [6, 61, 66], interactive editing of causal graphs [20, 63, 64], comparison of causal learning methods [19], the interpretation of causal graphs [20, 31, 63, 64], and causal subgraphs [39]. Therefore, combining feedback from domain experts with design principles from previous causal-analysis-related visualization works, we further specify the following design requirements.

- R1 Support the interpretation of causal graphs.** The causal graphs should be interpreted by domain experts with ease. Clearly visualize causal graphs with a reasonable graph layout and provide effective visual encoding of the information of nodes and edges.
- R2 Visual comparison of causal graphs computed by various causal discovery methods.** None of the existing causal discovery methods always yields good results in all cases. Visualize results of multiple causal discovery methods with comparative visualization as a reference to analysts for constructing a reasonable causal graph.
- R3 Improve the efficiency and interpretability of causal discovery.** Long run time and poor interpretability are typical issues of causal discovery methods, e.g., modern deep-learning methods are too slow to converge for an interactive visual analysis system. Visualize intermediate results and diagnostic measures of the model quality to guide a user-controllable causal discovery process balancing the speed and quality.
- R4 Assist the selection of variables for analysis.** Datasets with multiple clinical outcomes typically contain a large number of variables. Including all variables would result in an overwhelmingly complex causal graph that prohibits understanding. Provide statistical and visualization methods combined with interactive approaches to help analysts decide possible related variables.
- R5 Incorporate human knowledge into causal analysis and causal graph editing.** To obtain real causal relationships with causal discovery algorithms is unlikely. Adding expert domain knowledge to the exploration and modification of causal graphs through interactions enhances the reliability of the results. Basic interactions should include adding and deleting nodes and edges, and inverting edges.
- R6 Allow intuitive comparison of causal relationships for different outcomes.** Different diseases may share the same influencing factors, and distinguishing the shared and different influencing factors can help medical experts rationally judge and intervene in specific influencing factors, and improve prognosis.

4 VISUALIZATION OF SINGLE CAUSAL GRAPH WITH VARIOUS CAUSAL DISCOVERY METHODS

In this section, we introduce the selected causal discovery methods, elaborate on the visualization of single causal graphs with the rationale of design choices, and explain the causal effect computation.

4.1 Causal Discovery

There are various causal discovery methods [62], but none is currently available as an all-around winner. Therefore, we aim to use recent causal discovery methods that inherently handle mixed-type data with categorical and continuous variables and could provide complementing results so that the user can combine them to get a good causal graph. Three causal discovery methods are adapted into our system: the classic Peter-Clark (PC) [55] algorithm, which is known to provide a good causality skeleton, the DAG-GNN method [67] as a representative modern continuous optimization method that can handle mixed-type data, and the hybrid causal learning method (HCM) [40], specifically designed for mixed-type data. More details of the methods are covered in the supplemental material.

(1) PC

The PC [55] algorithm is a well-known and frequently used causal discovery method that serves as the backbone approach. The method starts with a fully connected undirected graph and identifies a skeleton by evaluating the conditional independence of each variable pair. Next, the edges are oriented by identifying colliders/v-structures. The algorithm then orients additional edges in partially directed paths that would otherwise form more v-structures.

(2) DAG-GNN

The original combinatorial optimization problem of causal discovery can be transformed into a continuous optimization problem [70]. DAG-GNN [67] uses a generative graph neural network based on variational autoencoders with neural network functions and evidence lower bound to solve the problem. It extends the NO TEARS [70] method for nonlinear structural equation models and supports both discrete and continuous variables.

(3) HCM

HCM is a recent advancement for mixed-type data. It addresses the mixed-type data problem by formulating a mixed structure equation model [40] that treats continuous and categorical variables separately. The model is learned with skeleton learning using the PC-stable algorithm [18] with a new mixed-type conditional independence test; then the method finds causal graph directions via greedy search with a mixed information criterion; and the graph is further pruned using the conditional independence test.

4.2 Evaluation with Benchmark Datasets

The effectiveness of using multiple causal discovery methods is demonstrated using two mixed-typed benchmark datasets: healthcare¹ and prepd-bathia². Results of the three algorithms are visualized in Fig. 3(right), with our visualization method (Sec. 4.3).

The combined results of all three algorithms yield an (undirected) accuracy of 100% and 92% compared to the true graph (Fig. 3 (left)), respectively, at the cost of high false positive rates. More numerical measurements are summarized in the supplemental material. With our visualization, false positives are not a severe issue compared to a lower accuracy because removing a wrong edge is easier than discovering a new true edge—the presence of a wrong edge is easily observable given the domain knowledge of the analyst.

Moreover, with our visualization, patterns can be found that are otherwise difficult with numerical measures only. For example, the coexistence of edges discovered by different methods suggests that the possibility of a true edge is high. In both benchmark datasets, edges

¹<https://www.bnlearn.com/bnrepository/clgaussian-small.html#healthcare>

²<https://www.bnlearn.com/research/scirep17/>




Fig. 3: Visualization of the true causal graphs (left column) and results of the three causal discovery algorithms (right column).

shared by all three methods are all true edges. Therefore, the combination of these algorithms provides users with important additional information for creating a better causal graph than using any of them alone.

4.3 Comparative Progressive Visualization of Causal Discovery Algorithms

Our visualization method for single-outcome causal graph addresses the following three issues: (1) it visually encodes information of a causal graph (R1), (2) supports the visual comparison of results generated by the algorithms in Sec. 4.1 (R2), and (3) visualizes the progress of the algorithms with user-controllable pausing and resuming (R3).

4.3.1 Visual Encodings of Causal Graphs

An efficient layered method³ that takes advantage of several graph drawing works [8, 11, 21, 35] is used for the layout of the causal graph. In this layout, nodes are arranged top-down with decreasing out-degrees, and the cause-outcome relationship is inherently encoded by the vertical locations of two variables, i.e., the direction of causality naturally is from the upper node to the lower node of the two nodes connected by an edge. However, results of causal discovery methods sometimes include reversed causalities with the layered layout, i.e., cause nodes are below effect nodes. Such potentially self-contradictory cases require user attention. Arrows are commonly used to encode directional properties, so we choose to highlight the reversed causality using an edge with an arrow heading from the cause to the effect and keep other edges with top-down causal directions undirected to reduce clutter.

Color and shape—known as effective visual channels for categorical variables [42]—are used to distinguish outcome and other nodes. Edge thickness, a frequently used channel for weight encoding in a weighted node-link diagram [2], shows the magnitude of the causal effect (Sec. 4.4). Note that the causal effect is signed, and we use the edge style to encode the sign: solid lines for positive and dashed lines for negative effect values, respectively. Logarithmic transformation is applied to adapt the large range of the effect value.

Example visualizations are shown in Fig. 1(1), Fig. 4(g), and Fig. 7(4). The node encoding uses jagged circles in orange for the outcome variable, and regular circles in blue for other variables, and the name of each variable is drawn at the center of the circle. The two

³<https://github.com/dagrejs/dagre/wiki>

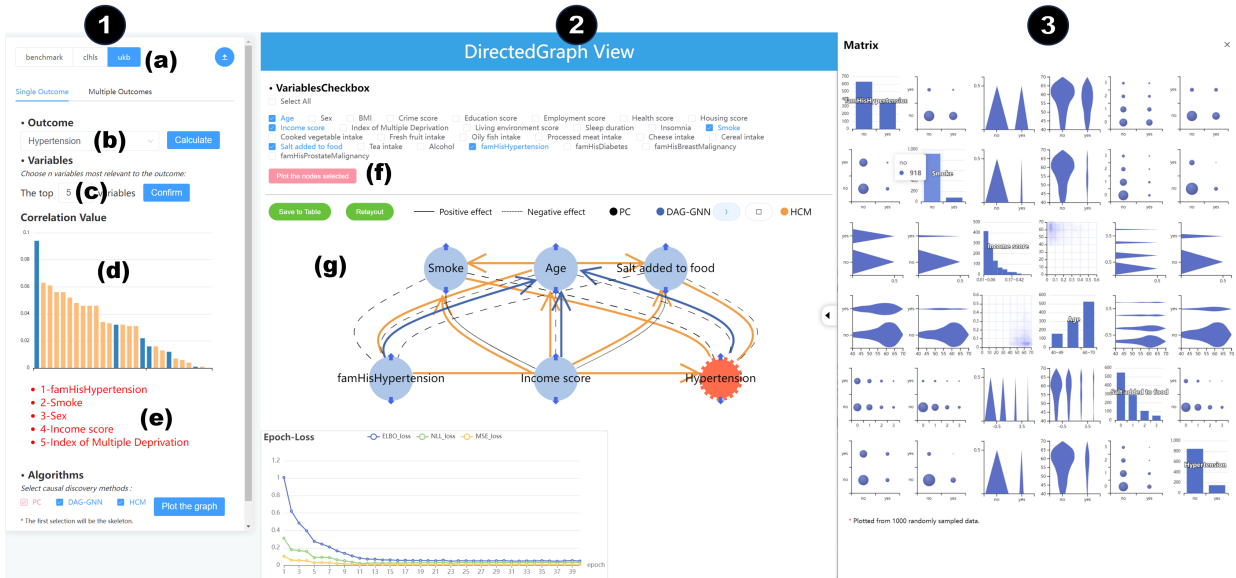


Fig. 4: The user interface of our visual analysis system for single outcome graphs consists of three views: (1) Dataset and Variables Selection View allowing for dataset (a) and variables selection (b, c) and correlation analysis (d, e), (2) Single Directed Graph View allowing for graph editing (f, g), and (3) Variables Matrix View providing insight for graph editing.

small widgets (in blue) attached to the top and bottom of a node are used for adding new edges.

4.3.2 Comparative Visualization of Causal Discovery Results

An appropriate visual comparison design for complementing causal discovery results could aid the editing of the final causal graph. We consider the comparison strategies [25] for our case that causal graphs have the same nodes but different edges. For this case, juxtaposition is slow and error-prone for comparison [2], and explicit encodings computing a number of pairwise relationships can be confusing and cluttered. In contrast, superposition that directly visualizes the presence/absence of an edge is an effective and space-efficient solution.

Our comparative visualization uses the superposition of the PC, DAG-GNN, and HCM results. Color coding is used for differentiation: black for PC, blue for DAG-GNN, and orange for HCM. Jittering is used to show the three results simultaneously without occlusion: the PC result is drawn at the center between nodes, the DAG-GNN result is offset to the right, and the HCM result is to the left. By default, all three algorithms are shown together. Each result can be highlighted by assigning it a high opacity while reducing the opacity of others with the legend shown on the top right of the view. Moreover, for showing a specific result or comparing any of the two, the rendering of each result can be conveniently switched on and off.

4.3.3 Progressive Visualization

The progressive visualization addresses the conflict between the long running time of causal discovery methods and the interactivity of visual analysis. For example, the DAG-GNN method by default runs for a large number of epochs, where each epoch takes a few seconds, making a naive use of it impractical. Instead, we progressively draw intermediate results at each epoch by updating the graph edges within the current layout, so that the running of the causal discovery methods does not prohibit the visual analysis.

Diagnostic measures of DAG-GNN are visualized as line charts in an epoch-loss view shown at the bottom of Fig. 1(f), Fig. 4(g), and Fig. 7(4). Specifically, the measurements—ELBO (Evidence Lower Bound) loss (blue), NLL (Negative Log-Likelihood) loss (green), and MSE (Mean Squared Error) loss (orange)—provide interpretability of the method and evidence for users to determine when to pause or stop the algorithm. Users can pause the process at any time to examine the current result and stop or restart the process depending on whether the result is satisfactory or not. Typically, the losses drop quickly after

a few dozen epochs and become stable, and similar is true for the associated causal graph result. Therefore, our visualization provides more insights into the DAG-GNN method than what was previously possible.

4.3.4 Multidimensional Visualization of Mixed-Typed Variables

To gain more insights into the variables of the causal graph, a multidimensional visualization is available in our method to show the distributions of variables of mixed types. As illustrated in Fig. 4(3), the visualization resembles the scatterplot matrix (SPLOM) showing distributions of every pair of variables. Unlike a SPLOM, we apply different visual encodings according to the type of variables according to the conventions of the domain experts. This design is suitable for our datasets (random sampling) as the visualizations are readable since the data are rather sparse and the encodings do not generate clutter.

For categorical vs. categorical, bubble plots that encode the number of samples with area, are used due to the familiarity of the domain experts. Violin plots—effective for comparing continuous distributions under different categories—are used for categorical vs. continuous. Scatterplots are used for continuous vs. continuous as they can clearly show distributions of two continuous variables. We use the whole matrix for the visualization as it is found more understandable and convenient for searching by domain experts than the triangular version.

With our comparative progressive visualization of the causal graph and the multidimensional visualization of variables involved, analysts can then interactively refine the causal graph of a single outcome with their domain knowledge.

4.4 Causal Effect Computation

The causal effect of a variable X on Y is defined as the expected change in Y due to a change in X . In our case, we apply the backdoor identification algorithm [46] and linear regression estimation to estimate the causal effect of each edge.

5 VISUALIZATION OF MULTI-OUTCOME CAUSAL GRAPHS

Once the individual causal graphs for individual outcomes have been analyzed and refined, our method supports the visual comparison of multi-outcome causal graphs composed by these individual graphs. In this section, we introduce the graph layout method for comparing multiple causal graphs and explain the associated visual designs facilitating comparisons.

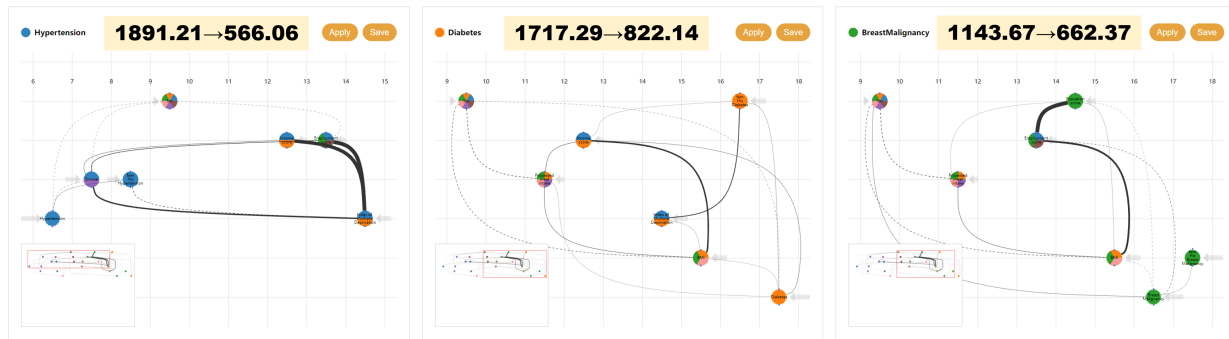


Fig. 5: Visualization of multi-outcome causal graphs with our new layout and visual mappings. Numbers on top of each subgraph are the horizontal stress of direct extraction (left of the arrow) and our method (right of the arrow), respectively.

5.1 Graph Layouts for Comparison

We first experimented with the direct comparison of subgraphs obtained using the aforementioned layout algorithm individually. We found that each subgraph reaches its optimal layout, but comparing these subgraphs is difficult. This is because the information of other subgraphs is missing during the layout computation and the subgraphs live in their own spaces that are not comparable. Therefore, it is necessary to include global information of all subgraphs for comparison. To this end, we propose two layouts for comparison, namely, the supergraph (for 2 to 3 subgraphs) and the comparable graph layout (for more subgraphs).

5.1.1 Supergraph of All Outcomes

We combine all single graphs of interest as one large causal graph, namely, the supergraph G_s . The supergraph is laid out using the aforementioned layered method and an example is shown in Fig. 1(h). According to the convention of the layout method, the vertical indices are called *ranks*, and the horizontal indices at each rank are *orders*.

The supergraph can show the similarities and differences between subgraphs because of the global image space and anchored common nodes. The supergraph properly supports comparing up to two or three outcomes but does not scale well for more outcomes perceptually. As shown in Fig. 1(h), comparing causal graphs of each outcome is difficult if not impossible with the supergraph of six outcomes. For more than three outcomes: in cases with few shared nodes, the supergraph becomes too wide; for many shared nodes, the association between nodes and subgraphs is difficult to perceive (shared nodes have several colors in our case). Instead, it is advantageous to compare individual graphs as line-ups (Fig. 1(i)) with comparable graph layouts.

A typical approach for subgraph line-ups would directly extract subgraphs of individual outcomes while preserving the locations of nodes and edges in the supergraph. This is the strategy used in the visual analysis of multiple biological interaction graphs [7]. However, subgraphs of individual outcomes extracted directly from the supergraph are rather sparse, as shown in Fig. 5.

5.1.2 Comparable Layout of Multiple Causal Graphs

Based on our experiments with various layout strategies, we decided to optimize the layout of subgraphs extracted from the supergraph to facilitate the comparison of more than two graphs. A key feature considered during the design of the layout algorithm is to preserve the relative node positions of the supergraph to make them comparable. Therefore, the objectives of the comparable layout are abstracted as: 1) preserving the relative node positions, 2) anchoring the common nodes, and 3) making the graphs compact.

The according constrained optimization problem reads for a given subgraph G_t of nodes V with input coordinates \mathbf{x} in the supergraph, output positions ξ , the set of indices Q of all common nodes V_Q shared

by any subgraphs, and the index set C_t of unique nodes of G_t :

$$\begin{aligned} \min_{V \in G_t} \sum_{i,j \in Q} w_i (||\mathbf{x}_i - \mathbf{x}_j|| - ||\xi_i - \xi_j||)^2 + \\ \sum_{k,l \in C_t} w_l (||\mathbf{x}_l - \mathbf{x}_k|| - ||\xi_l - \xi_k||)^2 \quad (1) \end{aligned}$$

subject to: V do not violate ranks and orders
 $V \in V_Q$ are anchored .

Since the supergraph already has the optimized layout [11], the problem of Eq. (1) can be approximated with a fast greedy algorithm. Note that the rank (the layer of nodes) indicates the out-degrees, and should be clearly separated and aligned between subgraphs for comparison. We decided to leave it intact after experimentation. Therefore, we just need to reduce the horizontal space by compressing the horizontal coordinates of nodes while maintaining the relative node positions.

Our horizontal compression algorithm first identifies common nodes shared by two or more subgraphs and compresses them horizontally to create anchored nodes V_Q . Then, for each subgraph G_t , unique nodes are relocated by compressing their horizontal distances to anchored nodes V_Q . Finally, the compressed subgraph is achieved by merging compressed unique nodes and the anchored nodes of G_t . The detailed algorithm can be found in the supplemental material.

Examples of subgraphs with the horizontal compression layout (namely, relative subgraphs) can be seen in Fig. 5. Compared to subgraphs extracted directly from the supergraph, our new layout provides better usage of screen space by removing much of the empty space while the relative positions are preserved. Moreover, we numerically evaluate the effectiveness of our method by computing the horizontal stress (i.e., $\text{stress}(X)$) of each subgraph [22]. A smaller stress number indicates a better approximation of the graph's theoretical distance. As shown on top of each subgraph, our layout method yields smaller $\text{stress}(X)$ numbers compared to directly extracted subgraphs. This suggests that our layout leads to edge distances closer to the graph's theoretical distances than directly extracted graphs.

5.2 Visual Mappings

The node encoding (jagged and regular disks) of the single graph is extended for the comparison of multiple causal graphs. Unlike the single graph, node color here encodes different outcome graphs. A node shared by multiple graphs has multiple colors filling the disk area evenly divided by the number of shared graphs. For example, some nodes shown in Fig. 6 are shared by various numbers of graphs. The visual encoding on edges (of effect values and causal directions) is the same as in Sec. 4.3.1. Rather than all three causal discovery results, only the final user-edited causal graph is shown.

With our layout method, non-anchored nodes at different original positions may land on similar or the same locations in subgraphs after the compression. Therefore, we augment the compressed graphs with reference grids and glyphs for accurate comparisons. Reference grids with coordinate labels can aid the comparison of the global positions of

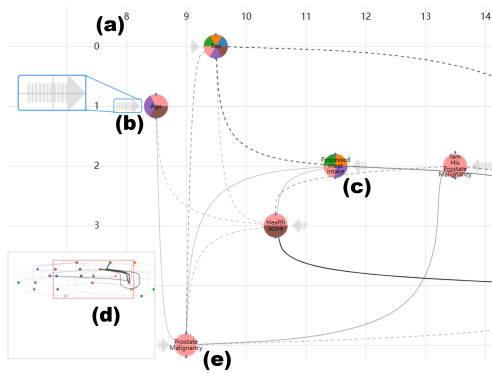


Fig. 6: Visual mappings for comparison of multi-outcome causal graphs.

nodes. The grids can be zoomed in and dragged for better comparison. Grids with coordinates provide quantitative information for comparison but require the shift of attention from nodes to the grid, which may potentially interfere with the mental map. Therefore, we design arrow glyphs (Fig. 6(b, c)) to directly encode the extent and directions of compression, i.e., left or right, on each node. The number of bars on the arrow glyph (the inset of Fig. 6) encodes the moved distances relative to the position in the supergraph. The arrow on the left (Fig. 6(b)) of a node indicates that the node has moved to the right relative to its original position after graph compression, and vice versa for arrows on the right (Fig. 6(c)).

To further indicate the global position of the subgraph, a thumbnail (Fig. 6(d)) is provided at the bottom left of each subgraph. The thumbnail view shows the bounds of the subgraph (as a red box) within the supergraph for a quick lookup.

6 VISUAL ANALYSIS OF MULTI-OUTCOME CAUSAL GRAPHS

The visual analysis of multi-outcome causal graphs is divided into two phases: single causal graph exploration and multi-outcome causal graph comparison and analysis (Fig. 2). This section introduces the components and details of the functionality of our visual analysis system, describes the typical analysis processes, and reports on the implementation. The user interface of our visual analysis system is shown in Fig. 1 and Fig. 4. Please refer to the supplemental video for the visual analysis process.

6.1 Single Graph Analysis

To realize our single causal graph exploration and analysis, we design the user interface of *single outcome graph analysis* (Fig. 4), which consists of three key views: (1) Dataset and Variables Selection View (Fig. 4(1)), (2) Single Directed Graph View (Fig. 4(2)), and (3) Variables Matrix View (Fig. 4(3)). A typical visual analysis of the causal graph of a single outcome is described as follows.

6.1.1 Variable Selection

The user should first choose the dataset (Fig. 4(a)) and then select the outcome of interest via a drop-down box (Fig. 4(b)). To start with, the user can choose the top n variables that have the highest correlation coefficients with the outcome via an input box (Fig. 4(c)), which helps reduce the complexity of the analysis to some extent (R4). The user can then refer to the correlation plot (Fig. 4(d)) (blue for positive correlations and orange for negative) sorted by correlation coefficients of the selected outcome with all other variables in descending absolute value order to decide whether to add or delete some nodes for causal analysis. Any variable of the dataset can be included or removed for the causal analysis by selections in the variable panel (Fig. 4(f)).

6.1.2 Causal Exploration and Construction of Causal Graphs

We consider the results calculated by the PC algorithm as the skeleton, and the results of DAG-GNN and HCM as references. Therefore, in our current design, all operations on nodes or edges are based on the results

of PC, and users can refer to the results of the other two algorithms to manipulate the current single causal graph.

Variables selected in Sec. 6.1.1 are input to the three causal discovery algorithms, and an initial causal graph (R1) is obtained using the layered graph layout (Sec. 4.3.1) based on the result computed by PC. In the meantime, causal graphs calculated by DAG-GNN and HCM are updated progressively in real-time as the two algorithms run in the backend and are shown with our comparative visualization (example in Fig. 1(1)) (R2, R3). Users can decide whether to pause, resume, or stop with the widgets (Fig. 1(d)) to control the progress of DAG-GNN.

The superposition of edges of the three algorithms (Fig. 1(e)) and the “epoch-loss” plot showing the model quality measures of DAG-GNN (Fig. 1(f)) can improve the interpretability of causal discovery (R2, R3) and support the user to determine a reasonable causal relation (R5). For example, if an edge exists in all three results, it is likely a true causality. The user can find the Variables Matrix View (Fig. 4(3)) that shows the multidimensional visualization of variables (including the outcome) obtained in the Dataset and Variables Selection View (Fig. 4(1)). By analyzing the distributions, the user can gain more insights into the relationships between any pair of selected variables (R4).

Combining domain knowledge with insights gained from the visualizations, the user can modify the causal graph interactively (R5). The variables checkbox (Fig. 1(a)) allows adding or removing nodes. By clicking on an edge, a menu pops up that allows deleting or reversing an edge (Fig. 1(b)). By connecting the blue widget of one node to the blue widget of another node, the user can add an edge (Fig. 1(c)). A re-layout of the causal graph can be initiated after edits for better visualization.

After iterative interactions with all these functionalities, a more reasonable causal graph for the outcome of interest can be obtained. The edited graph can be stored in the history for the visual comparison of multi-outcome causal graphs in the second stage.

6.2 Multi-outcome Graphs Analysis

The user repeats the process in Sec. 6.1 to obtain individual causal graphs for different outcomes and store them in the history panel (Fig. 1(g)). After that, the user can select graphs of interest for comparison from the history panel (R6). By manipulating the different tabs, the user can visualize the multi-outcome graphs in various styles including the supergraph (*Super Graph*), our new comparable layout for subgraphs (*Relative Subgraphs*), and subgraphs directly extracted from the supergraph (*Extracted Subgraphs*) serving as an uncompressed reference to our method. In all cases, each subgraph can be interactively manipulated with zooming and panning for in-depth comparisons. Aided by visualization, the user can gain insight into the similarities and differences between any two subgraphs (i.e., disease outcomes) regarding influencing factors and the extent of the effects of causality.

6.3 Implementation

Our visual analysis system is implemented using JavaScript and Python. The front-end functionality is realized with the Vue2 framework and data visualization libraries Echarts and D3.js. The graph layout computation is aided by the dagre library. The causal analysis and data processing in the back-end are implemented in Python. Multi-threading is used to support the progressive visualization of single causal graphs. The causal effect estimations are calculated using the causal inference library DoWhy.

7 CASE STUDY

We demonstrate the usability of our multi-outcome analysis system with a case study investigating the causality of various medical conditions. We applied our approach to two health research datasets (permissions obtained), CLHLS [15] (a longitudinal survey focusing on the health and longevity of elderly people in China, collecting a wide range of health, social and economic data) and UKB⁴ (a database of extensive biomedical data on approximately 500,000 UK residents, covering genetic, lifestyle, and health information).

⁴<https://www.ukbiobank.ac.uk>

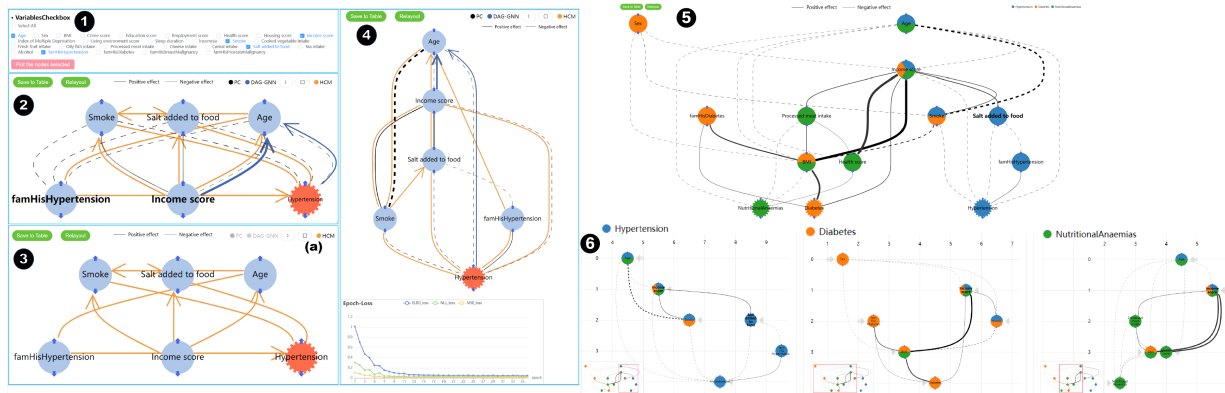


Fig. 7: The case study of UKB data with a medical expert. Single causal graph analysis (1–4) and multi-outcome causal graphs comparison (5, 6).

We invited one domain expert (**E1**, a co-author who was involved in the design study) to participate in our case study hosted in our laboratory. **E1** is an assistant professor at a medical university with more than 10 years of clinical medicine experience. She holds a clinical medicine license and performs health data research. The study began with a comprehensive 20-minute presentation on the background of causal discovery methods, and our two-stage system design and visualization. Subsequently, the expert dedicated an hour to freely explore our system and expressed her thoughts through a think-aloud method. The observers listened and talked to the expert to answer any questions. Finally, a 30-minute post-study interview was conducted to gather subjective comments, questions, and feedback from the expert on the method and the system. High-resolution images of the case study can be found in the supplemental material.

7.1 Causality in UKB

We used the preprocessed original UKB dataset to obtain 8 frequently observed diseases as outcome variables and 27 influencing variables related to physiological condition, education level, lifestyle, diet, family history, etc., which may affect these outcomes. The final dataset we used in our research contains records of 98,530 people with at least one out of the 8 outcomes. We report on two representative insights gained by the expert in this section.

Single causal graph of hypertension. The expert first chose hypertension (Fig. 4(b)) as the outcome and selected the top 5 (Fig. 4(c)) most relevant variables as potential influencing factors (Fig. 4(e)) to explore the causal relationship between these variables. Compared to “Sex” and “Index of Multiple Deprivation”, she thought that “Age” and “Salt added to food” were more likely to affect the occurrence of hypertension, so she changed the nodes through the variables checkbox (Fig. 7(1)) and then recalculated the DAG (Fig. 7(2)). Users can switch on or off the legend of each algorithm (Fig. 7(a)) to show a specific result (results with HCM alone are shown in Fig. 7(3)). Observing these three results, **E1** commented that “Each algorithm can calculate causal relationships that are partially consistent with domain knowledge, but none is perfectly correct,” which supports our choice to apply multiple algorithms to complement each other. Combining domain knowledge and insights gained from the three algorithms, the expert modified the original DAG (Fig. 7(2)), and the final causal graph approved by the domain expert is shown in Fig. 7(4). Some causal relationships, such as the causal relationship of the negative effect value of “Age \rightarrow Hypertension” and “Smoke \rightarrow Hypertension,” seem inconsistent with the first impression, however, **E1** considered it may be due to the reality that “fewer heavy smokers can live old enough.”

Causal comparison of multiple outcomes. After analyzing the causal graph of hypertension, the expert further explored the causal graph of diabetes and nutritional anemia for further comparison (Fig. 7(5), (6)), as she anticipated that these different diseases might coexist in the same patient and tend to share some influencing factors, but the same factors have different mechanisms on different diseases. Hypertension and diabetes share “Income score” and “Smoke,” hy-

pertension and nutritional anemia share “Age” and “Income score,” diabetes and nutritional anemia share “Income score” and “BMI.” **E1** found that a supergraph is easy for comparing two different outcomes. However, when three or more outcomes are compared, it becomes difficult with a supergraph (Fig. 7(5)), in which case our method can help with better comparisons (Fig. 7(6)). Note the anchored shared nodes “Income score,” “Smoke,” and “Age” denoted with grids in Fig. 7(6).

For the shared node “Income score,” positive and negative causal effect values of “Income score \rightarrow Hypertension” (negative) and “Income score \rightarrow Diabetes” (positive) are found. **E1** thought it is reasonable that people with a higher “Income score” have stronger health awareness, and therefore are less likely to develop hypertension. However, the other links that the wealthy have higher BMI and are more likely to develop diabetes are unusual; **E1** commented that “It may be associated with the sampled people in this data” and further investigation is needed.

The effect of “Income score” on anemia is not direct. Smoking has a direct causal effect on the occurrence of hypertension, but the causal effect on the occurrence of diabetes is not quite clear, suggesting that smoking cessation behavior may be more beneficial to the control of hypertension, which is in line with the expectations of the expert. Similarly, the causal relationships in the graph suggest that salt restriction may be better for hypertension control and weight loss may be better for diabetes control. Obese people (high BMI) are more likely to develop diabetes and less likely to develop nutritional anemia. Insights gained from the comparison of different outcomes are in line with the expectations of **E1** and are beneficial to guide the intervention in the prognosis of various diseases.

7.2 Causality in CLHLS

The CLHLS data was preprocessed to have 13,860 records with 10 outcome variables and 28 influencing variables related to physiological condition, education level, lifestyle, social and economic status, etc. We asked **E1** to explore this dataset for cross-validation using the same outcome variables, hypertension, and diabetes, as in the exploration of the UKB dataset described above.

E1 got similar insights from this dataset. For example, “*f31_sum*” is an indicator of economic sources, and a higher value indicates a better income situation, and thus a lower possibility of having hypertension and diabetes; obese people (high BMI) are more likely to develop hypertension and diabetes, etc.

Interesting differences to the UKB data were also found: for example, people in the CLHLS are older than those in the UKB, **E1** commented that “it is likely that the older, the lower rate of hypertension and diabetes, since those with diseases will not live long enough.”

8 EXPERT USER EVALUATION

We invited three domain experts (**E1**, **E2**, **E3**) all with more than 8 years of clinical or public health experience to evaluate our system (**E2** and **E3** are not involved in the design study). Since the case study took

a long time (more than 2 hours), **E2** and **E3** were only involved in the evaluation phase due to their tight schedule. **E2** is an associate professor of public health, and **E3** is a doctoral student of health data science, both at the university that the first/corresponding author is affiliated with. **E2** studies cancer epidemiology, data science of epidemiology, and potential risk factors of common diseases, while **E3** has extensive research experience in chronic diseases. The evaluation took place in our lab using a laptop computer extended with a 27-inch display. Two researchers (authors) first introduced the method and the tool and demonstrated the use of the tool with the CLHLS. The experts were then asked to explore the UKB data by themselves. The think-aloud protocol was used during the evaluation: the experts made comments and asked questions while the researchers answered their questions.

All experts rated the two-stage analysis framework positively, as it allows them to first obtain the causality of one specific outcome and then compare those different yet correlated diseases, giving them insights to better understand conditions like multimorbidity or comorbidity and eventually help improve the prognosis. They agreed that the comparative visualization of different algorithms is beneficial for guiding the editing of causal graphs (**R2**, **R5**). **E1** commented: “If all three algorithms calculate an edge, I would be inclined to assume that there is a high probability that the edge exists, which helps to edit the causal relationship reasonably when there is doubt about some causal relationship.” In addition, the progressive visualization of DAG-GNN improves the interpretability of the results (**R3**), according to **E2**: “I will choose to pause the algorithm when the line chart of epoch-loss tends to be stable, which indicates that the algorithm has obtained relatively satisfactory results and can be referred to.” Experts commented that the visualization of multivariate variables (Fig. 4(3)) and correlation coefficients (Fig. 4(d, e)) are helpful in the selection of variables (**R4**).

They also commented on the visualization and encoding design of our system. The layered layout from top to bottom makes it easy to intuitively determine the direction of causality (**R1**). All experts agreed that the visual encoding of nodes and edges is intuitive, and the operation process is relatively simple. **E3** commented that “it is easy for me to edit the causal graph combining with domain knowledge” (**R5**). Regarding multi-outcome causal graph comparison (**R6**), all experts agreed that it is convenient with the *history view* (Fig. 1(g)). They appreciated both forms of visual comparison: the supergraph (Fig. 1(h)) and relative subgraphs (Fig. 1(i)) for different situations. They commented that the supergraph is preferred for up to three outcomes, while the relative subgraphs should be used for more outcomes.

9 DISCUSSION

In our current analysis system, users are asked to select a limited number of potential variables to reduce the algorithm running time, partly supported by calculating the correlation coefficients. Note that the datasets we used contain both categorical and continuous variables, suggesting we should apply different methods to calculate the correlation coefficients [51]. However, we argue that correlation coefficients calculated by different methods are incomparable and, therefore, confusing. Moreover, we only use the correlation coefficient to assist in the selection of variables, and the final selection of variables is more based on expert domain knowledge, therefore, we choose to use the Pearson correlation coefficient uniformly. Domain experts commented that providing more advanced representations of influencing factors would be useful, for example, by calculating multivariate logistic regression, or through data-driven predictive models instead of the correlation coefficients. We will explore this possibility in future work.

Both CLHLS and UKB data used in our case are cross-sectional, which means that the direction of causality between variables cannot be fully determined. For example, people who smoke or are heavy on salt are more likely to develop hypertension, but in cross-sectional data, it is also possible that they are told to quit smoking and limit salt intake because of high blood pressure. If cohort data are available, the sequence of events in time can be used to further determine the causal direction. In addition, as mentioned in Sec. 7, our current analysis is based on partial variables extracted from two datasets. More interesting discoveries may be made if we can access more comprehensive data.

The computational scalability [49] of our method is mainly limited by the causal discovery methods. We conducted a scalability study with varying number of variables for the CLHLS and UKB datasets. The study suggests that our method is feasible for the causal analysis of large datasets with less than ten variables per outcome with reasonable computational time. Although we did not perform a formal visual scalability study, our experience shows that the visualization becomes cluttered if more than ten variables are involved per outcome. Experts commented that focusing on the most important factors, typically less than ten, was sufficient to solve most of the specific domain problems. For seven variables, PC took 6 seconds (for the CLHLS dataset) and 52 seconds (UKB dataset) to finish, DAG-GNN took 2 minutes (CLHLS) and 3 minutes (UKB) to be stable, and HCM took 3 minutes (CLHLS) and 4 minutes (UKB), all on a machine with Intel i7 2.60 GHz CPU and 16 GB RAM. To address the issue of long computation times for DAG-GNN and HCM, we designed our system to perform their computations progressively to avoid interference with user interactions. Details of the study are documented in the supplemental material.

Even though the method aims to resolve specific medical problems, several new techniques in our method address general issues in causal analysis and visual comparison. For example, causal discovery is enhanced by the synergy of various learning algorithms, including state-of-the-art techniques; the deep-learning-based causal discovery method is integrated into interactive visual analysis with progressive visualization that also improves interpretability; comparative visualizations enable effective comparisons of multiple causal graphs and multiple causal learning techniques. Our method is thus possible to be applied to other datasets that contain more than one outcome for multi-outcome comparison. For example, in business, purchasing decisions, brand loyalty, shopping experience, etc. can be regarded as different outcomes, and there may be the same or different influencing factors among these outcomes. Analyzing the causal differences between these outcomes can help enterprises optimize marketing strategies and improve product design according to various goals.

We attempted to compare our method with existing visual analysis methods of causal graphs [31, 63, 64, 66]. However, none of them provides readily usable tools or source codes. Therefore, we could only refer to the papers and associated video demos. After examining these sources, we believe that these methods can partly address the single outcome causal graph analysis problem but do not support the creation and the subsequent comparative analysis of multiple causal graphs with different outcome variables. Moreover, each of these works uses only one classical causal discovery method, e.g., GES, F-GES, PC, to generate causal graphs.

10 CONCLUSION AND FUTURE WORK

In conclusion, we have introduced the visual analysis of multi-outcome causal graphs that were previously not studied in visual analytics. Inspired by a domain problem in healthcare, the method is designed in collaboration with medical experts. Our method consists of two stages: starting from a single outcome, analysts explore, analyze, and edit the associated causal graph discovered by a combination of causal discovery algorithms aided by our progressive visualization technique; causal graphs of single outcomes are then visually compared using a new graph layout and visual mappings as line-ups or as a supergraph. The usefulness of our method was evaluated through a case study and an expert user study with medical experts using widely accepted medical datasets.

In the future, we would like to improve the guidance of initial variable selection using more advanced data analysis techniques than correlation analysis. We would also like to support the analysis of causal graphs with external evidence in addition to the data, such as well-accepted knowledge graphs and conclusions learned from state-of-the-art publications on the domain problem. Another direction of improvement is to study the temporal aspect of cohort data as it provides vital information for understanding causality, especially in healthcare problems.

ACKNOWLEDGMENTS

This work was supported in part by the NSFC grant (No. 62372012), the National Key R&D Program of China (No. 2022YFC2705104), and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project-ID 449742818.

REFERENCES

- [1] E. Alexander and M. Gleicher. Task-driven comparison of topic models. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):320–329, 2016. doi: [10.1109/TVCG.2015.2467618](https://doi.org/10.1109/TVCG.2015.2467618) 3
- [2] B. Alper, B. Bach, N. Henry Riche, T. Isenberg, and J.-D. Fekete. Weighted graph comparison techniques for brain connectivity analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, p. 483–492. Association for Computing Machinery, New York, NY, USA, 2013. doi: [10.1145/2470654.2470724](https://doi.org/10.1145/2470654.2470724) 3, 4, 5
- [3] B. Andrews, J. Ramsey, and G. F. Cooper. Scoring Bayesian networks of mixed variables. *International Journal of Data Science and Analytics*, 6(1):3–18, 2018. doi: [10.1007/s41060-017-0085-7](https://doi.org/10.1007/s41060-017-0085-7) 2
- [4] B. Andrews, J. Ramsey, and G. F. Cooper. Learning high-dimensional directed acyclic graphs with mixed data-types. In *Proceedings of Machine Learning Research*, vol. 104 of *Proceedings of Machine Learning Research*, pp. 4–21. PMLR, 2019. 2
- [5] K. Andrews, M. Wohlfahrt, and G. Wurzing. Visual graph comparison. In *13th International Conference Information Visualisation*, pp. 62–67, 2009. doi: [10.1109/IV.2009.108](https://doi.org/10.1109/IV.2009.108) 3
- [6] J. Bae, T. Helldin, and M. Riveiro. Understanding indirect causal relationships in node-link graphs. *Computer Graphics Forum*, 36(3):411–421, 2017. doi: [10.1111/cgf.13198](https://doi.org/10.1111/cgf.13198) 2, 3
- [7] A. Barsky, T. Munzner, J. Gardy, and R. Kincaid. Cerebral: Visualizing multiple experimental conditions on a graph with biological context. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1253–1260, 2008. doi: [10.1109/TVCG.2008.117](https://doi.org/10.1109/TVCG.2008.117) 6
- [8] W. Barth, M. Jünger, and P. Mutzel. Simple and efficient bilayer cross counting. In M. T. Goodrich and S. G. Kobourov, eds., *Revised Papers from the 10th International Symposium on Graph Drawing*, pp. 130–141. Springer-Verlag, Berlin, Heidelberg, 2002. doi: [10.1007/3-540-36151-0_13](https://doi.org/10.1007/3-540-36151-0_13) 4
- [9] K. Bello, B. Aragam, and P. Ravikumar. DAGMA: learning DAGs via M-matrices and a log-determinant acyclicity characterization. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, number 598 in NIPS '22, pp. 8226–8239. Curran Associates Inc., Red Hook, NY, USA, 2024. doi: [10.5555/3600270.3600868](https://doi.org/10.5555/3600270.3600868) 2
- [10] D. Borland, A. Z. Wang, and D. Gotz. Using counterfactuals to improve causal inferences from visualizations. *IEEE Computer Graphics and Applications*, 44(1):95–104, 2024. doi: [10.1109/MCG.2023.3338788](https://doi.org/10.1109/MCG.2023.3338788) 2
- [11] U. Brandes and B. Köpf. Fast and simple horizontal coordinate assignment. In P. Mutzel, M. Jünger, and S. Leipert, eds., *Graph Drawing*, pp. 31–44. Springer, Berlin, Heidelberg, 2002. doi: [10.1007/3-540-45848-4_3](https://doi.org/10.1007/3-540-45848-4_3) 4, 6
- [12] P. Brouillard, S. Lachapelle, A. Lacoste, S. Lacoste-Julien, and A. Drouin. Differentiable causal discovery from interventional data. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, eds., *Proceedings of the 34th International Conference on Neural Information Processing Systems*, vol. 33, pp. 21865–21877. Curran Associates, Inc., Red Hook, NY, USA, 2020. doi: [doi/abs/10.5555/3495724.3497558](https://doi.org/10.5555/3495724.3497558) 2
- [13] S. Busetti. Causality is good for practice: policy design and reverse engineering. *Policy Sciences*, 56(2):419–438, 2023. doi: [10.1007/s11077-023-09493-7](https://doi.org/10.1007/s11077-023-09493-7) 1
- [14] Z. Cai, D. Xi, X. Zhu, and R. Li. Causal discoveries for high dimensional mixed data. *Statistics in Medicine*, 41(24):4924–4940, 2022. doi: [10.1002/sim.9544](https://doi.org/10.1002/sim.9544) 2
- [15] Center for Healthy Aging and Development Studies. The Chinese Longitudinal Healthy Longevity Survey (CLHLS)-Longitudinal Data (1998-2018), 2020. doi: [10.18170/DVN/WBO7LK](https://doi.org/10.18170/DVN/WBO7LK) 1, 7
- [16] D. M. Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3:507–554, 2003. 2
- [17] Y.-L. Chou, C. Moreira, P. Bruza, C. Ouyang, and J. Jorge. Counterfactuals and causability in explainable artificial intelligence: Theory, algorithms, and applications. *Information Fusion*, 81:59–83, 2022. doi: [10.1016/j.inffus.2021.11.003](https://doi.org/10.1016/j.inffus.2021.11.003) 1
- [18] D. Colombo and M. H. Maathuis. Order-independent constraint-based causal structure learning. *Journal of Machine Learning Research*, 15(1):3741–3782, 2014. 4
- [19] J. A. Cottam, M. Glenski, Z. H. Shaw, R. Rabello, A. Golding, S. Volkova, and D. L. Arendt. Graph comparison for causal discovery. In *Proceedings of the Sixth Symposium on Visualization in Data Science*, VDS 2021, 2021. 2, 3
- [20] Z. Deng, D. Weng, X. Xie, J. Bao, Y. Zheng, M. Xu, W. Chen, and Y. Wu. Compass: Towards better causal analysis of urban time series. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1051–1061, 2022. doi: [10.1109/TVCG.2021.3114875](https://doi.org/10.1109/TVCG.2021.3114875) 3
- [21] E. Gansner, E. Koutsofios, S. North, and K.-P. Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–230, 1993. doi: [10.1109/32.221135](https://doi.org/10.1109/32.221135) 4
- [22] E. R. Gansner, Y. Koren, and S. North. Graph drawing by stress majorization. In J. Pach, ed., *Proceedings of the 12th International Conference on Graph Drawing*, pp. 239–250. Springer-Verlag, Berlin, Heidelberg, 2004. doi: [10.1007/978-3-540-31843-9_25](https://doi.org/10.1007/978-3-540-31843-9_25) 6
- [23] B. Ghai and K. Mueller. D-BIAS: A causality-based human-in-the-loop system for tackling algorithmic bias. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):473–482, 2023. doi: [10.1109/TVCG.2022.3209484](https://doi.org/10.1109/TVCG.2022.3209484) 2
- [24] M. Gleicher. Considerations for visualizing comparison. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):413–423, 2018. doi: [10.1109/TVCG.2017.2744199](https://doi.org/10.1109/TVCG.2017.2744199) 3
- [25] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, 2011. doi: [10.1177/1473871611416549](https://doi.org/10.1177/1473871611416549) 3, 5
- [26] G. Guo, E. Karavani, A. Ender, and B. C. Kwon. Causalvis: Visualizations for causal inference. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, number 462 in CHI '23, pp. 1–20. Association for Computing Machinery, New York, NY, USA, 2023. doi: [10.1145/3544548.3581236](https://doi.org/10.1145/3544548.3581236) 2
- [27] P. Hanlon, B. I. Nicholl, B. D. Jani, D. Lee, R. McQueenie, and F. S. Mair. Frailty and pre-frailty in middle-aged and older adults and its association with multimorbidity and mortality: a prospective analysis of 493737 UK Biobank participants. *The Lancet Public Health*, 3(7):e323–e332, 2018. doi: [10.1016/S2468-2667\(18\)30091-4](https://doi.org/10.1016/S2468-2667(18)30091-4) 3
- [28] C. Harrison, M. Fortin, M. van den Akker, F. Mair, A. Calderon-Larranaga, F. Boland, E. Wallace, B. Jani, and S. Smith. Comorbidity versus multimorbidity: Why it matters. In *Journal of multimorbidity and comorbidity*, vol. 11. England, 2021. doi: [10.1177/2633556521993993](https://doi.org/10.1177/2633556521993993) 2
- [29] M. Hascoët and P. Dragicevic. Interactive graph matching and visual comparison of graphs and clustered graphs. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, AVI '12, p. 522–529. Association for Computing Machinery, New York, NY, USA, 2012. doi: [10.1145/2254556.2254654](https://doi.org/10.1145/2254556.2254654) 3
- [30] A. Holzinger, G. Langs, H. Denk, K. Zatloukal, and H. Müller. Causability and explainability of artificial intelligence in medicine. *WIREs Data Mining and Knowledge Discovery*, 9(4):e1312, 2019. doi: [10.1002/widm.1312](https://doi.org/10.1002/widm.1312) 1
- [31] M. N. Hoque and K. Mueller. Outcome-Explorer: A causality guided interactive visual interface for interpretable algorithmic decision making. *IEEE Transactions on Visualization and Computer Graphics*, 28(12):4728–4740, 2022. doi: [10.1109/TVCG.2021.3102051](https://doi.org/10.1109/TVCG.2021.3102051) 2, 3, 9
- [32] A. Hyttinen, P. O. Hoyer, F. Eberhardt, and M. Järvisalo. Discovering cyclic causal models with latent variables: a general SAT-based procedure. In *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence*, UAI '13, p. 301–310. AUAI Press, Arlington, Virginia, USA, 2013. doi: [doi/10.5555/3023638.3023669](https://doi.org/10.5555/3023638.3023669) 2
- [33] Z. Jin, N. Chen, Y. Shi, W. Qian, M. Xu, and N. Cao. TrammelGraph: visual graph abstraction for comparison. *Journal of Visualization*, 24:365–379, 2021. doi: [10.1007/s12650-020-00706-2](https://doi.org/10.1007/s12650-020-00706-2) 3
- [34] Z. Jin, S. Guo, N. Chen, D. Weiskopf, D. Gotz, and N. Cao. Visual causality analysis of event sequence data. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1343–1352, 2021. doi: [10.1109/TVCG.2020.3030465](https://doi.org/10.1109/TVCG.2020.3030465) 2
- [35] M. Jünger and P. Mutzel. *2-Layer Straightline Crossing Minimization: Performance of Exact and Heuristic Algorithms*, pp. 3–27. World Scientific, 2002. doi: [10.1142/9789812777638_0001](https://doi.org/10.1142/9789812777638_0001) 4
- [36] S. Kaul, D. Borland, N. Cao, and D. Gotz. Improving visualization interpretation using counterfactuals. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):998–1008, 2022. doi: [10.1109/TVCG.2021.3114779](https://doi.org/10.1109/TVCG.2021.3114779) 2
- [37] N. R. Ke, O. Bilaniuk, A. Goyal, S. Bauer, H. Larochelle, B. Schölkopf,

- M. C. Mozer, C. Pal, and Y. Bengio. Learning neural causal models from unknown interventions, 2020. doi: [10.48550/arXiv.1910.01075](https://doi.org/10.48550/arXiv.1910.01075) 2
- [38] D. Lewis. Causation. *The Journal of Philosophy*, 70(17):556–567, 1973. doi: [10.2307/2025310](https://doi.org/10.2307/2025310) 2
- [39] R. Li, W. Cui, T. Song, X. Xie, R. Ding, Y. Wang, H. Zhang, H. Zhou, and Y. Wu. Causality-based visual analysis of questionnaire responses. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):638–648, 2024. doi: [10.1109/TVCG.2023.3327376](https://doi.org/10.1109/TVCG.2023.3327376) 3
- [40] Y. Li, R. Xia, C. Liu, and L. Sun. A hybrid causal structure learning algorithm for mixed-type data. *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, 36(7):7435–7443, 2022. doi: [10.1609/aaai.v36i7.20707](https://doi.org/10.1609/aaai.v36i7.20707) 2, 4
- [41] X. S. Liang and X.-Q. Yang. A note on causation versus correlation in an extreme situation. *Entropy*, 23(3):316, 2021. doi: [10.3390/e23030316](https://doi.org/10.3390/e23030316) 3
- [42] J. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions Graphics*, 5(2):110–141, 1986. doi: [10.1145/22949.22950](https://doi.org/10.1145/22949.22950) 4
- [43] T. Munzner. A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):921–928, 2009. doi: [10.1109/TVCG.2009.111](https://doi.org/10.1109/TVCG.2009.111) 3
- [44] J. E. Olson, P. Y. Takahashi, and J. M. St Sauver. Understanding the patterns of multimorbidity. In *Mayo Clinic Proceedings*, vol. 93, pp. 824–825. England, 2018. doi: [10.1016/j.mayocp.2018.05.016](https://doi.org/10.1016/j.mayocp.2018.05.016) 3
- [45] J. Pearl. *Causality*. Cambridge University Press, 2nd ed., 2009. doi: [10.1017/CBO9780511803161](https://doi.org/10.1017/CBO9780511803161) 2
- [46] J. Pearl. An introduction to causal inference. *The International Journal of Biostatistics*, 6(2), 2010. doi: [10.2202/1557-4679.1203](https://doi.org/10.2202/1557-4679.1203) 5
- [47] J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. The MIT Press, 2017. 2
- [48] J. G. Richens, C. M. Lee, and S. Johri. Improving the accuracy of medical diagnosis with causal machine learning. *Nature Communications*, 11:3923, 2020. doi: [10.1038/s41467-020-17419-7](https://doi.org/10.1038/s41467-020-17419-7) 1
- [49] G. Richer, A. Pister, M. Abdelaal, J.-D. Fekete, M. Sedlmair, and D. Weiskopf. Scalability in visualization. *IEEE Transactions on Visualization and Computer Graphics*, 30(7):3314–3330, 2024. doi: [10.1109/TVCG.2022.3231230](https://doi.org/10.1109/TVCG.2022.3231230) 9
- [50] P. Sanchez, J. P. Voisey, T. Xia, H. I. Watson, A. Q. O’Neil, and S. A. Tsaftaris. Causal machine learning for healthcare and precision medicine. *Royal Society Open Science*, 9(8):220638, 2022. doi: [10.1098/rsos.220638](https://doi.org/10.1098/rsos.220638) 1
- [51] P. Schöber, C. Boer, and L. A. Schwarte. Correlation coefficients: Appropriate use and interpretation. *Anesthesia & Analgesia*, 126(5):1763–1768, 2018. doi: [10.1213/ANE.0000000000002864](https://doi.org/10.1213/ANE.0000000000002864) 9
- [52] M. Sedlmair, M. Meyer, and T. Munzner. Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2431–2440, 2012. doi: [10.1109/TVCG.2012.213](https://doi.org/10.1109/TVCG.2012.213) 3
- [53] S. Shimizu, P. O. Hoyer, A. Hyvärinen, and A. Kerminen. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research*, 7:2003–2030, 2006. doi: [doi/10.5555/1248547.1248619](https://doi.org/10.5555/1248547.1248619) 2
- [54] S. T. Skou, F. S. Mair, M. Fortin, B. Guthrie, B. P. Nunes, J. J. Miranda, C. M. Boyd, S. Pati, S. Mtenga, and S. M. Smith. Multimorbidity. *Nature Reviews Disease Primers*, 8:48, 2022. doi: [10.1038/s41572-022-00376-4](https://doi.org/10.1038/s41572-022-00376-4) 3
- [55] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. Springer-Verlag, New York, 1993. doi: [10.1007/978-1-4612-2748-9](https://doi.org/10.1007/978-1-4612-2748-9) 2, 4
- [56] Z. Sun, J. Xu, X. Zhang, Z. Dong, and J.-R. Wen. Law article-enhanced legal case matching: A causal learning approach. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR ’23, p. 1549–1558. Association for Computing Machinery, New York, NY, USA, 2023. doi: [10.1145/3539618.3591709](https://doi.org/10.1145/3539618.3591709) 1
- [57] C. Tazzeo, D. Rizzuto, A. Calderón-Larrañaga, A. Roso-Llorach, A. Marengoni, A.-K. Welmer, G. Onder, C. Trevisan, and D. L. Vetrano. Multimorbidity patterns and risk of frailty in older community-dwelling adults: a population-based cohort study. *Age and Ageing*, 50(6):2183–2191, 2021. doi: [10.1093/ageing/afab138](https://doi.org/10.1093/ageing/afab138) 3
- [58] C. Tominski, C. Forsell, and J. Johansson. Interaction support for visual comparison inspired by natural behavior. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2719–2728, 2012. doi: [10.1109/TVCG.2012.237](https://doi.org/10.1109/TVCG.2012.237) 3
- [59] S. Triantafillou, V. Lagani, C. Heinze-Deml, A. Schmidt, J. Tegner, and I. Tsamardinos. Predicting causal relationships from biological data: Applying automated causal discovery on mass cytometry data of human immune cells. *Scientific Reports*, 7:12724, 2017. doi: [10.1038/s41598-017-08582-x](https://doi.org/10.1038/s41598-017-08582-x) 1
- [60] J. M. Valderas, B. Starfield, B. Sibbald, C. Salisbury, and M. Roland. Defining comorbidity: implications for understanding health and health services. *The Annals of Family Medicine*, 7(4):357–363, 2009. doi: [10.1370/afm.983](https://doi.org/10.1370/afm.983) 2, 3
- [61] D.-B. Vo, K. Lazarova, H. C. Purchase, and M. McCann. Visual causality: Investigating graph layouts for understanding causal processes. In A.-V. Pietarinen, P. Chapman, L. Bosveld-de Smet, V. Giardino, J. Corter, and S. Linker, eds., *Diagrammatic Representation and Inference*, pp. 332–347. Springer International Publishing, Cham, 2020. doi: [10.1007/978-3-030-54249-8_26](https://doi.org/10.1007/978-3-030-54249-8_26) 2, 3
- [62] M. J. Vowels, N. C. Camgoz, and R. Bowden. D’ya Like DAGs? A survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4):1–36, article no. 82, 2022. doi: [10.1145/3527154](https://doi.org/10.1145/3527154) 2, 4
- [63] J. Wang and K. Mueller. The visual causality analyst: An interactive interface for causal reasoning. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):230–239, 2016. doi: [10.1109/TVCG.2015.2467931](https://doi.org/10.1109/TVCG.2015.2467931) 2, 3, 9
- [64] J. Wang and K. Mueller. Visual causality analysis made practical. In *2017 IEEE Conference on Visual Analytics Science and Technology*, VAST, pp. 151–161, 2017. doi: [10.1109/VAST.2017.8585647](https://doi.org/10.1109/VAST.2017.8585647) 2, 3, 9
- [65] J. Wang and K. Mueller. DOMINO: Visual causal reasoning with time-dependent phenomena. *IEEE Transactions on Visualization and Computer Graphics*, 29(12):5342–5356, 2023. doi: [10.1109/TVCG.2022.3207929](https://doi.org/10.1109/TVCG.2022.3207929) 2
- [66] X. Xie, F. Du, and Y. Wu. A visual analytics approach for exploratory causal analysis: Exploration, validation, and applications. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1448–1458, 2021. doi: [10.1109/TVCG.2020.3028957](https://doi.org/10.1109/TVCG.2020.3028957) 2, 3, 9
- [67] Y. Yu, J. Chen, T. Gao, and M. Yu. DAG-GNN: DAG structure learning with graph neural networks. In K. Chaudhuri and R. Salakhutdinov, eds., *Proceedings of the 36th International Conference on Machine Learning*, vol. 97 of *Proceedings of Machine Learning Research*, pp. 7154–7163. PMLR, 2019. doi: [10.48550/arXiv.1904.10098](https://doi.org/10.48550/arXiv.1904.10098) 2, 4
- [68] A. Zanga, E. Ozkirimli, and F. Stella. A survey on causal discovery: Theory and practice. *International Journal of Approximate Reasoning*, 151:101–129, 2022. doi: [10.1016/j.ijar.2022.09.004](https://doi.org/10.1016/j.ijar.2022.09.004) 1, 2
- [69] L. Zhang, X. Wang, Y. Liu, G. Zhang, M. Jing, and J. Yu. Visual representation and layout optimization for comparison of dynamic graph. In *2022 IEEE Smartworld, Ubiquitous Intelligence & Computing, Scalable Computing & Communications, Digital Twin, Privacy Computing, Metaverse, Autonomous & Trusted Vehicles (SmartWorld/UIC/ScalCom/DigitalTwin/PriComp/Meta)*, pp. 1091–1098, 2022. doi: [10.1109/SmartWorld-UIC-ATC-ScalCom-DigitalTwin-PriComp-Metaverse56740.2022.00169](https://doi.org/10.1109/SmartWorld-UIC-ATC-ScalCom-DigitalTwin-PriComp-Metaverse56740.2022.00169) 3
- [70] X. Zheng, B. Aragam, P. Ravikumar, and E. P. Xing. DAGs with NO TEARS: continuous optimization for structure learning. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, p. 9492–9503. Curran Associates Inc., Red Hook, NY, USA, 2018. doi: [doi/10.5555/3327546.3327618](https://doi.org/10.5555/3327546.3327618) 2, 4