

CAUSALENS: A System for Summarizing Causal DAGs

Noam Chen
noamchen@campus.technion.ac.il
Technion
Haifa, Israel

Anna Zeng
annazeng@mit.edu
CSAIL, MIT
Boston, USA

Michael Cafarella
michjc@csail.mit.edu
CSAIL, MIT
Boston, USA

Batya Kenig
batyak@technion.ac.il
Technion
Haifa, Israel

Markos Markakis
markakis@mit.edu
CSAIL, MIT
Boston, USA

Oren Mishali
omishali@cs.technion.ac.il
Technion
Haifa, Israel

Brit Youngmann
brity@technion.ac.il
Technion
Haifa, Israel

Babak Salimi
bsalimi@ucsd.edu
University of California, San Diego
San Diego, USA

Abstract

Causal inference aids researchers in discovering causal relationships, leading to scientific insights. Pearl’s causal model uses causal DAGs to estimate causal effects, so DAG correctness is essential for reliable causal conclusions. However, for high dimensional data, the causal DAGs are often complex beyond human verifiability. Graph summarization is a logical next step, but current methods for general-purpose graph summarization are inadequate for causal DAG summarization, as they are not designed to preserve causal information. In this demonstration, we present a system called CAUSALENS that summarizes a given causal DAG, and balances graph simplification for better understanding and retention of essential causal information for reliable inference directly on the summary DAG. We illustrate that causal inference on the summary DAG is more robust to misspecification in the initial causal DAG compared to performing inference directly on the initial causal DAG, thereby enhancing the robustness of causal inference. We will demonstrate the utility of CAUSALENS for generating useful summary causal DAGs by interacting with the SIGMOD’25 participants, who will act as data analysts aiming to perform causal analysis on high dimensional datasets.

CCS Concepts

• **Information systems** → *Data management systems*; Graph-based database models.

Keywords

Causal Inference, Graph Summarization

ACM Reference Format:

Noam Chen, Anna Zeng, Michael Cafarella, Batya Kenig, Markos Markakis, Oren Mishali, Brit Youngmann, and Babak Salimi. 2025. CAUSALENS: A

System for Summarizing Causal DAGs. In *Companion of the 2025 International Conference on Management of Data (SIGMOD-Companion '25)*, June 22–27, 2025, Berlin, Germany. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3722212.3725086>

1 Introduction

Causal inference is central to informed decision-making in economics, sociology, and medicine. It has become increasingly relevant in machine learning, supporting multiple tasks, including algorithmic fairness and explainable AI. Causal inference has also become a major theme in data management, enhancing tasks like query results explanation [16], and data discovery and cleaning [2, 10].

The core idea behind causal inference is that computing the effect of a variable T (the *treatment* variable) on another variable O (the *outcome* variable) should be aware of the causal structure of the system of which they are a part. This structure is often represented as a Causal Directed Acyclic Graph (DAG) [8], where each directed edge represents possible causal influence. Graphical criteria, such as the backdoor criterion [8], can then be applied to a causal DAG to find a sufficient set of *confounding variables* to adjust for. This ensures sound causal inference, even if the underlying data is purely *observational* (i.e., passively collected).

The reliability of causal inference thus depends on high-quality causal DAGs, which are often unavailable. Such DAGs are often constructed using domain knowledge, a process that is error-prone and time-consuming. Alternatively, multiple causal discovery methods have been proposed [3], which attempt to recover a causal DAG from available data alone. While useful, often they do not perform well on real-world data and require significant human effort for verification [13]. This challenge is exacerbated in high-dimensional data, underscoring the need for efficient methods to simplify and validate causal DAGs. We illustrate this need via an example:

Example 1.1. Consider a dataset about query performance on a cloud-based data warehouse service¹ that includes performance metrics and query features. This dataset can help answer key causal questions for performance optimization, such as understanding



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGMOD-Companion '25, Berlin, Germany*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1564-8/2025/06
<https://doi.org/10.1145/3722212.3725086>

¹Amazon Redshift Serverless: <https://aws.amazon.com/redshift/redshift-serverless/>




Figure 1: CAUSALENS UI showing example causal DAGs for analyzing the Redshift dataset on query performance. Left: Manual causal DAG with color-coded semantic relationships. Right: Generated summary causal DAG based on user-defined size and similarity constraints.

how caching impacts latency, or analyzing the effect of join complexity on query planner performance. Fig. 1 (left) shows a small possible causal DAG for this task, constructed manually. For example, the edge from NumColumns to ExecTime suggests that the number of columns referenced in a query may affect its execution time. In this example DAG, QueryTemplate is a confounder that must be adjusted for in order to answer the causal queries above in an unbiased way: it affects both performance metrics (e.g., ElapsedTime, PlanTime) and the mechanisms being analyzed (e.g., ResultCacheHit, NumJoins). If the DAG was mis-specified and some of the relevant edges from QueryTemplate were missing, we would reach inaccurate conclusions. This underscores the importance of domain expert verification for edges, a task that can be overwhelming even in this simple example.

Causal DAGs might be generated by an automated method or by a different team. In either case, the DAG is very complicated, so graph summarization is a logical next step, making it easier for users to verify. Graph summarization has been extensively studied [4, 5, 12, 15], with techniques such as node grouping based on similarity [14], and bit reduction [9]. While these general-purpose methods excel at handling large graphs, they are unsuitable for summarizing causal DAGs, as they focus on objectives like minimizing reconstruction errors or enhancing visualizations rather than preserving causal information crucial for reliable inference.

In this work, we demonstrate CAUSALENS, a system that simplifies an input causal DAG into a *summary causal DAG*, making the verification process easier. In [17], we show that the generated summary DAG can be used for sound causal inference and is more resilient to errors compared to the original DAG. CAUSALENS thereby enhances *interpretability*, *verifiability*, and *robustness* in causal inference. In CAUSALENS, the user provides an input causal DAG or generates one using a causal discovery algorithm via the

UI. They can then set a bound on the number of nodes in the summary DAG. To ensure semantic coherence, CAUSALENS estimates the semantic similarity between variables (visualized through color codes in Fig. 1), allowing the user to set a threshold for the maximum semantic distance between nodes within the same cluster in the summary DAG. In the demonstration, we will show how the summary DAG can be used for reliable causal inference and how it is more robust to errors in the input DAG.

2 Technical Background

2.1 Background on Causal Inference

We consider a single-relation database over a schema $\mathbb{A}$. We use upper case letters for a variable from $\mathbb{A}$ and bold symbols for variable sets. The broad goal of causal inference is to estimate the effect of a *treatment variable* $T \in \mathbb{A}$ (NumJoins) on an *outcome variable* $O \in \mathbb{A}$ (ElapsedTime). The gold standard approach is performing a *randomized controlled experiment*, where the population is randomly divided into a *treated* group (denoted $\text{do}(T = 1)$ for a binary treatment) and a *control* group ($\text{do}(T = 0)$). Randomization mitigates the effect of *confounding factors*, i.e., variables that would have affected both T and O , were the treatments not explicitly assigned.

One popular measure of a causal estimate is the *Average Treatment Effect* (ATE). In a randomized experiment, the ATE is the difference in the average outcomes of the treated and control groups [8]:

$$\text{ATE}(T, O) = \mathbb{E}[O \mid \text{do}(T = 1)] - \mathbb{E}[O \mid \text{do}(T = 0)] \quad (1)$$

However, experiments where treatments are assigned at random cannot be done in many practical scenarios due to, e.g., economic or ethical issues. In such cases, we can still account for the confounding attributes to get an unbiased causal estimate *from observational data*, by using Pearl’s model [8]. To achieve this, we must replace the *do*-operator in Equation 1 with appropriate conditional probabilities.




Table 1: The recursive bases of the DAGs in Figure 2

Graph	Recursive Basis
$\mathcal{G}$	$(C \perp\!\!\!\perp B A), (D \perp\!\!\!\perp A BC), (E \perp\!\!\!\perp ABC D)$
$\mathcal{H}_1$	$(E \perp\!\!\!\perp ABC D)$
$\mathcal{H}_2$	$(E \perp\!\!\!\perp AC BD)$

Finding the right conditional probabilities to substitute can be achieved using the rules of the *do*-calculus framework [8], which requires access to a causal DAG. A causal DAG is a specific type of a Bayesian network, where nodes represent random variables and edges signify *potential* causal influence. Starting from such a graph, Pearl developed the *backdoor criterion* [8], which allows us to graphically identify a sufficient set of confounding variables.

2.1.1 Causal Information. A causal DAGs encodes causal information as Conditional Independence (CI) relationships among the variables. These relationships can be read off the DAG using a graphical criterion called *d-separation* [8], which determines whether two sets of nodes X and Y are conditionally independent, given a third set Z . The *Recursive Basis* (RB) for a causal DAG consists of up to n CIs, where n is the number of nodes, indicating each node's independence from its non-descendant nodes, given its parents. This set of CIs holds significance, as it can be used to derive all other CIs encoded in the DAG using the semi-graphoid axioms. As an example, consider the causal DAG shown in Fig. 2(a). An example CI it encodes is $B \perp\!\!\!\perp C | A$. Its RB is given in Table 1.

2.1.2 CIs & Missing Edges. In causal DAGs, the information encoded by *missing* edges implies the set of CIs the DAG represents. Namely, removing edges can undermine the causal model as it implies additional CIs that do not necessarily hold. On the other hand, existing edges only indicate *potential* causal influence. This implies that adding edges to a causal DAG, provided acyclicity is maintained, does not necessarily compromise validity [8].

2.2 The Causal DAG Summarization Problem

CAUSALENS takes an input causal DAG and simplifies it for better interpretability while preserving its usefulness for causal inference. We begin by defining summary causal DAGs.

2.2.1 Summary Causal DAG. A *summary causal DAG* is created through *node contraction*, where only acyclicity-preserving contractions are allowed. Contracting nodes U, V produces a graph $\mathcal{H}$, replacing U and V with a single node $C=\{U, V\}$. C retains the neighbors of U and V , while any edge between U and V is removed.

A causal DAG $\mathcal{G}$ is said to be *compatible* with a summary DAG $\mathcal{H}$, if, there exists a partition of the nodes $V(\mathcal{G})$ among the nodes $V(\mathcal{H})$, such that: If $(U, V) \in E(\mathcal{G})$, then either U and V are in the

same node in $V(\mathcal{H})$ or there is an edge from the cluster node of U to the cluster node of V in $E(\mathcal{H})$. The RB of a summary DAG is defined in a similar manner. The only difference is that in a summary DAG, a node may represent a set of nodes of the original DAG. Table 1 presents the recursive bases of the example causal DAG and two possible summary DAGs for it, as illustrated in Fig. 2.

2.2.2 Problem Formulation. Our objectives are to maintain as much causal information as possible while simplifying the input DAG. Causal Information Preservation: Given two summary DAGs derived from the same DAG $\mathcal{G}$, we prefer the one that preserves a larger number of the CIs encoded in $\mathcal{G}$. We compare summary DAGs based on their RBs. Given two summary DAGs $\mathcal{H}_1$ and $\mathcal{H}_2$, we say that $\mathcal{H}_1$ is *superior* to $\mathcal{H}_2$ if the RB of $\mathcal{H}_2$ is implied by that of $\mathcal{H}_1$. Namely, all the CIs encoded by $\mathcal{H}_2$ can be deduced from $\mathcal{H}_1$.

Example 2.1. Consider the causal DAG $\mathcal{G}$ shown in Fig. 2, with two possible summary DAGs $\mathcal{H}_1$ and $\mathcal{H}_2$. Despite $\mathcal{H}_1$ having only three nodes, it surpasses $\mathcal{H}_2$. From $\mathcal{H}_1$, for example, we can infer that $(E \perp\!\!\!\perp_d A | D)$, while this CI cannot be deduced from $\mathcal{H}_2$.

Additionally, a summary DAG should not introduce any spurious CIs that the original causal DAG does not imply. We refer to this property as an I-Map. Our goal is to find a summary DAG that is an I-Map for the input causal DAG and its RB is maximal.

Size Constraint The summary DAG should be concise to reduce cognitive load for analysts. We therefore allow users to impose a size constraint on the number of nodes in the summary DAG.

Given a causal DAG $\mathcal{G}$ and a bound $k > 1$, CAUSALENS finds a summary causal DAG $\mathcal{H}$ s.t. (i) the number of nodes in $\mathcal{H}$ is $\leq k$; (ii) $\mathcal{G}$ is compatible with $\mathcal{H}$, is an I-Map for $\mathcal{G}$ and its RB is maximal.

In [17], we show that the rules of *do*-calculus are sound and complete in summary causal DAGs. This is vital to ensure that the summary causal DAGs are effective formats that support causal inference. In a nutshell, we show that contracting nodes is akin to adding edges to the input causal DAG. Based on this connection, we can think of node contraction as the operation of adding edges to the input causal DAG. As mentioned, adding edges to a causal DAG does not undermine the causal model, but it reduces the amount of CIs it encodes. Our goal is then to find the summary DAG that results in the minimal number of added edges.

2.3 Causal DAG Summarization Algorithm

We can show that the causal DAG summarization problem is NP-hard. We, therefore, employ a heuristic, scalable algorithm that meets the size constraints. CAUSALENS's underlying algorithm follows a previous line of work (e.g., [14]), where a bottom-up greedy approach is used to identify promising node pairs for contraction. Its main contribution lies in *how* it estimates merge costs, to preserve the causal interpretation of the graph: It counts the number of edges to be added to the original causal DAG for each node pair (a proxy for the effect of summarization on the RB, per Section 2.2.2). In each iteration, the algorithm contracts the node pair resulting in the minimal number of additional edges.

Semantic Constraint. The user can ensure that only semantically related variables are merged, thereby maintaining semantic coherence in the summary DAG. To achieve this, CAUSALENS evaluates

pairwise semantic similarity scores using GPT-4 [7]. The scores are displayed to the user by color-coding the nodes. For instance, in Fig. 1, the semantic scores of NumTable, NumJoins, and NumColumns are relatively high, indicated by their similar yellow shades. Users can specify a semantic threshold. Two nodes are eligible for contraction only if their similarity score exceeds the threshold. The color coding of the nodes in the summary DAG is consistent with the colors used in the input DAG. For example, if two yellow shade nodes are merged, the cluster node will also be colored yellow.

3 System & Demonstration

We implemented CAUSALENS in Python. To generate semantic similarity scores, we used GPT-4 [7]. We used the DoWhy [11] package to compute causal effects and Streamlit² for the UI.

System overview: The UI of CAUSALENS is shown in Fig. 1. The analyst provides a dataset with a corresponding causal DAG. If a causal DAG is not provided, CAUSALENS can leverage a causal discovery algorithm to infer one. Next, the analyst specifies a size constraint and optionally a semantic threshold. CAUSALENS then constructs a summary DAG that satisfies the size and semantic constraints. If no summary DAG can be found (e.g., when k is low, and the semantic threshold is high), CAUSALENS notifies the user.

In this demonstration, we will show that the obtained summary causal DAGs are valuable for reliable causal inference. We will also highlight how they effectively help users identify and address errors in their input causal DAGs, as error detection is significantly easier in summary DAGs compared to the original causal DAGs.

Demo scenarios: We demonstrate the operation of CAUSALENS over six datasets, including the Redshift dataset (12 nodes and 23 edges in the causal DAG), the German Credit dataset [1] (21 nodes, 43 edges), and a Car Accidents Data [6] (41 nodes, 368 edges). The participants will act as data analysts, aiming to explore and summarize their input causal DAG, before using it for causal analysis.

Introducing CAUSALENS: In this part of the demonstration, attendees will explore pre-loaded causal DAGs, such as the one in Figure 1 for REDSHIFT. We will invite the attendees to select a dataset and load its corresponding causal DAG. We will then explain the available configuration knobs in CAUSALENS, setting them to initial values and obtaining a summary causal DAG. Attendees can then examine how the configuration knobs affect the summary DAG, as well as estimate causal effects over the summary and original DAGs to demonstrate the summary DAG’s utility for causal inference. Specifically, we will compare ATEs estimated from the original DAG and the summary DAGs, each with a 95% confidence interval. While confounding variable sets may differ between the summary and original DAGs, resulting in some variation, we will show significant overlap in the intervals.

Demonstrating Robustness: The participants will then be invited to change the input causal DAGs using the UI. They can either modify the input DAG by adding or removing edges, or use a causal discovery algorithm to derive a new DAG. This scenario simulates a real-life task where an analyst tries to verify the correctness of a causal DAG. Participants will then inspect the generated summary DAG to see the robustness of the summary DAG against errors in

the input DAG. In particular, we will demonstrate that the ATEs calculated from summary causal DAGs are robust to errors in the input DAGs. This highlights their practical importance in real-world scenarios where a perfect input causal DAG may not be available.

Additionally, we will present some flawed input causal DAGs and ask participants to identify the errors (e.g., missing/redundant edges). Participants will be invited to inspect the input causal DAG directly, or examine the corresponding summary DAG. They will quickly realize that reviewing a summary DAG is a simpler task, as it requires inspecting fewer edges and nodes, thereby facilitating the verification process of the input causal DAG.

Looking under the hood: Participants will be invited to examine how various parameters affect quality and performance. For example, we will run an exhaustive approach that considers all summary causal DAGs, and show that while the summary DAG generated by CAUSALENS is similar to that of the exhaustive approach, CAUSALENS is at least one order of magnitude faster, and the difference in runtime increases as the size of the causal DAG increases.

Acknowledgments

We appreciate the support from DARPA ASKEM Award HR00112220 042, the ARPA-H Biomedical Data Fabric project, and Liberty Mutual. We also want to thank the NSF Award IIS-2340124, NIH Grant U54HG012510, and the ISF Grant No. 1442/24 for their support.

References

- [1] Arthur Asuncion and David Newman. 2007. UCI machine learning repository.
- [2] Sainyam Galhotra, Yue Gong, and Raul Castro Fernandez. 2023. Metam: Goal-oriented data discovery. In *ICDE*. IEEE, 2780–2793.
- [3] Clark Glymour, Kun Zhang, and Peter Spirtes. 2019. Review of causal discovery methods based on graphical models. *Frontiers in genetics* 10 (2019), 524.
- [4] Kyuhan Lee, Hyeonsoo Jo, Jihoon Ko, Sungsu Lim, and Kijung Shin. 2020. Ssum: Sparse summarization of massive graphs. In *SIGKDD*. 144–154.
- [5] Antonio Maccioni and Daniel J Abadi. 2016. Scalable pattern matching over compressed graphs via dedensification. In *SIGKDD*. 1755–1764.
- [6] Sobhan Moosavi, Mohammad Hossein Samavatian, Srinivasan Parthasarathy, Radu Teodorescu, and Rajiv Ramnath. 2019. Accident risk prediction based on heterogeneous sparse data: New dataset and insights. In *SIGSPATIAL*. 33–42.
- [7] OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- [8] Judea Pearl. 2000. *Causality : models, reasoning, and inference*. Cambridge University Press.
- [9] Ryan A Rossi and Rong Zhou. 2018. Graphzip: a clique-based sparse graph compression method. *Journal of Big Data* 5, 1 (2018), 1–14.
- [10] Babak Salimi, Luke Rodriguez, Bill Howe, and Dan Suciu. 2019. Interventional fairness: Causal database repair for algorithmic fairness. In *SIGMOD*. 793–810.
- [11] Amit Sharma and Emre Kiciman. 2020. DoWhy: An End-to-End Library for Causal Inference. *arXiv preprint arXiv:2011.04216* (2020).
- [12] Kijung Shin, Amol Ghoting, Myunghwan Kim, and Hema Raghavan. 2019. Sweg: Lossless and lossy summarization of web-scale graphs. In *WWW*. 1679–1690.
- [13] Karamjit Singh, Garima Gupta, Vartika Tewari, and Gautam Shroff. 2018. Comparative benchmarking of causal discovery algorithms. In *CODS-COMAD. Association for Computing Machinery*, 46–56.
- [14] Yuanyuan Tian, Richard A Hankins, and Jignesh M Patel. 2008. Efficient aggregation for graph summarization. In *SIGMOD*. 567–580.
- [15] Quinton Yong, Mahdi Hajiabadi, Venkatesh Srinivasan, and Alex Thomo. 2021. Efficient graph summarization using weighted lsh at billion-scale. In *SIGMOD*. 2357–2365.
- [16] Brit Youngmann, Michael Cafarella, Amir Gilad, and Sudeepa Roy. 2024. Summarized Causal Explanations For Aggregate Views. *PACMMOD* 2, 1 (2024), 1–27.
- [17] Anna Zeng, Michael Cafarella, Batya Kenig, and Babak Youngmann, Brit Salimi. 2025. Causal DAG Summarization. *Proc. VLDB Endow.* (2025).

²<https://streamlit.io/>