

Visual Analytics for Causal Reasoning from Real-World Health Data

Arran Zeyu Wang*
UNC Chapel Hill

David Borland†
RENCI, UNC Chapel Hill

David Gotz‡
UNC Chapel Hill

ABSTRACT

The increasing capture and analysis of large-scale longitudinal health data offer opportunities to improve healthcare and advance medical understanding. However, a critical gap exists between (a)—the observation of patterns and correlations, versus (b)—the understanding of true causal mechanisms that drive outcomes. An accurate understanding of the underlying mechanisms that cause various changes in medical status is crucial for decision-makers across various healthcare domains and roles, yet inferring causality from real-world observational data is difficult for both methodological and practical challenges. This Grand Challenge advocates increased Visual Analytics (VA) research on this topic to empower people with the tool for sound causal reasoning from health data. We note this is complicated by the complex nature of medical data—the volume, variety, sparsity, and temporality of health data streams make the use of causal inference algorithms difficult. Combined with challenges imposed by the realities of health-focused settings, including time constraints and traditional medical work practices, existing causal reasoning approaches are valuable but insufficient. We argue that advances in research can lead to new VA tools that augment human expertise with intuitive and robust causal inference capabilities, which can help realize a new paradigm of data-driven, causality-aware healthcare practices that improve human health outcomes.

Index Terms: Causality, Visual Analytics, Healthcare, Outcome, Treatment, Temporal Data

1 INTRODUCTION

The increasingly embraced medical data revolution provides many promising opportunities to improve health outcomes through data-driven practices [5], but also introduces significant new challenges. Visual analytics has been seen as a powerful tool to address many of these challenges [8], and much progress has been made. However, one crucial aspect remains underexplored (for both health and visual analysis discipline in general [2, 11]): *visual analytic tools for causal reasoning*.




Figure 1: A traditional bar chart visualization of treatment A (T_A) and treatment B (T_B) for kidney stone treatment data [4].

Causal understanding and inference are central to many medical decisions. Doctors choose interventions with the aim of causing

improvements to a patient’s condition. Patients choose treatments with the goal of causing benefits to their health status. Public health professionals choose policies with the aim of causing improvements to the health outcomes for a given population. Yet many visualizations of health data rely on representations of patterns and correlations that leave any causal inferences to the user without proper visual support.

For example, consider a bar chart showing data for two kidney stone treatments, T_A and T_B , where T_B exhibits a higher success rate (Figure 1). A decision maker looking at this chart might understandably come to the conclusion that the data support the use of T_B . However, omitted from the design of this chart is a critical, high-stakes question: Why? Was it the treatment itself that caused an improvement for T_B patients, or perhaps T_B patients were different from those that received T_A in a meaningful way (younger, healthier, milder disease state)? In this case, we see that T_B is correlated with better outcomes, but is it the cause? **Incorrectly mistaking correlation for causation** can lead to misguided decisions that worsen rather than improve health outcomes.

In the previous example, a closer examination of the kidney stone treatment data (from a real-world case [13]) shows that the populations receiving the two treatments T_A and T_B were in fact quite different in how their condition presented (small vs. large stone size). In a classic case of Simpson’s Paradox, the data show that T_B had a higher efficacy rate overall while T_A performed better for both small and large stone populations. This simple but compelling example provides a warning regarding the potential of confounding variables to mislead us as we seek data-based evidence for our decisions.

Table 1: Success rates of kidney stone treatments (T_A and T_B) for small and large stones [4]. Values in **bold** imply that T_B is more effective overall while T_A is more effective for both small and large stones.

Stone Size	T_A	T_B
Small	93% (81/87)	87% (234/270)
Large	73% (192/263)	69% (55/80)
All	78% (273/350)	83% (289/350)

Meanwhile, medical data can be enormously complex, with potentially hundreds of thousands of non-independent variables to consider (the ICD-10 system alone has approximately 70,000 distinct codes), Disentangling the causal relationship between any single event (e.g., a particular treatment decision) and a desired outcome becomes enormously challenging. This becomes even more difficult with the introduction of advanced statistical and AI models which can make the connection between specific factors and outcomes even more opaque.

Visual Analytics (VA) offers a powerful way forward, not just for data exploration, but as a crucial bridge for **insight-driven causal inference and validation**. VA is more than just data visualization; if designed to support causal reasoning, VA could serve as a bridge between users and the underlying data and algorithms. It could enable a collaborative process where users (e.g., clinicians, patients, policymakers) could partner with intelligent systems via a “causal conversation” with the data. This position paper calls upon the visual analytics community to tackle one of the most profound challenges in data-driven healthcare: **uncovering causal relationships from real-world data to improve health decision-making**.

*e-mail: zeyuwang@cs.unc.edu

†e-mail: borland@renci.org

‡e-mail: gotz@unc.edu

2 THE FRAILITY OF STATISTICAL CAUSAL MODELS

Existing visualization methods [20, 21, 9, 6] heavily rely on the use of expert-defined directed acyclic causal graphs [15]. However, this approach could often face challenges in real-world healthcare:

Computation: Automatically discovering a causal graph from data is an *NP-hard* problem [7]. For a typical clinical dataset, this is not just a theoretical concern. Consider a moderately complex dataset from an EHR system with thousands of dimensions and sometimes yields 1 million [8] per patient—a mix of demographics, diagnoses, lab results, medications, and vital signs taken over time. The number of potential relationships between these variables explodes into the trillions, making the exhaustive computation of a reliable causal graph practically impossible.

Comprehension: Even if we could compute such a causal graph, what would a patient, a clinician, or a health analyst do with it? A visualization of a causal graph with a number of nodes and a number of directional causal links could possibly be an incomprehensible “hairball” general users, who are usually not experts in statistics. It therefore may fail to provide the actionable insight that is the ultimate goal of the analysis [2]. The challenge is to *invent novel visual analytics paradigms that open up this black box*, allowing users to scrutinize, understand, and even steer the process of creating comparable cohorts, thus understanding causal graphs.

Future Direction

- **Transparent Visual Causal Modeling.** A key opportunity for VAHC is to create visual analytics tools that make the causal modeling process transparent and interpretable—allow users to interactively build causal models, visualize the impact of certain variables, and understand the trade-offs made by statistical methods.

3 LIMITATIONS BEYOND CAUSAL GRAPHS

Statistical techniques like propensity scoring [10] have proven useful for causal inference by creating comparable groups from observational data. With this idea, counterfactual visualizations [17, 18, 19] were proposed to help users reason causality by constructing “what if” data subsets. They lower the barrier to entry by not requiring causal models and statistical expertise. However, the processes and assumptions often remain opaque and challenging to clinical users.

Temporality: Causality in healthcare is fundamentally tied to sequences and time—all use cases in health data can be relevant to time [1]. While seemingly straightforward, effectively visualizing temporal data to support causal reasoning is exceptionally difficult. In such conditions, a meaningful counterfactual question is rarely “What if the patient had been given Treatment B?” but rather, “What if Treatment B had been administered after the fever spiked on day 3, instead of Treatment A?” Initial effort [12] has been made to deal with event sequence data; however, it heavily relied on even more complex causal models, and may implicitly rely on *temporal precedence as a proxy for causation*. For example, such a visualization might clearly show that patients receiving a new therapy were less likely to be readmitted to the hospital. However, it may fail to visually account for a critical hidden confounder: perhaps the new therapy was only given to healthier patients in the first place. The visualization, by merely ordering events, can unintentionally reinforce a fallacious causal conclusion.

Dynamic real-world scales: While promising, many proposed techniques have only been demonstrated on simplified, low-dimensional datasets. With a 1000+ dimension patient record, a simple user interface may become unusable, e.g., which of the thousand features should the user be prompted to change? How can the system guide the user to plausible or meaningful alternative scenarios without overwhelming them with infinite choices? Existing

systems lack the guidance and navigational functionalities needed to *make causal inference processes feasible and meaningful at this scale*.

Future Direction

- **Visualizing Causal Effects in Temporal Treatment Pathways.** An urgent challenge for VAHC is to possibly connect visual causal inference with event sequence visualization, moving from a logic of “*this followed that*” to “*this happened instead of that, considering these other factors*”.

4 CHALLENGES FROM THE EMERGING ERA

Beyond the foundational issues of scalability and temporal representation, the rapid evolution of emerging technologies presents a new set of challenges.

Healthcare AI models The deployment of complex, often opaque AI models in clinical settings is accelerating [16]. These models may recommend treatments or predict risks, but they rarely explain their reasoning in a causally meaningful way for general users [3]. One of the key challenges for VA is to serve as a *causal auditor and explainer for AI systems’ healthcare scenarios*.

Multi-modal data The complexity of healthcare data is no longer just about the number of variables; it’s about the variety of data types [8]. A patient’s story is told through structured EHR events, continuous time-series from ICU monitors, high-resolution pathology images, static but massive genomic profiles, and narrative clinical notes. A pressing new challenge is to develop VA frameworks that *support causal inference across these fundamentally different modalities*.

Collaborative inference Historically, VA has focused on helping users analyze data under the assumption that either no reliable causal model exists [14], or the causal model is entirely pre-designed by experts [9]. A forward-looking challenge is to flip this paradigm: can we use VA to connect general users and experts to collaboratively build and refine causal models? In such a system, the large models could analyze the data to suggest potential causal links based on statistical evidence. The human users would then use their own knowledge or simply priors to visually confirm, reject, prune, or modify these links, adding directionality and context that the data alone cannot provide—resulting in a *direct manipulation of causal models*.

Future Direction

- **Causal Auditing and Explainability for Clinical AI.** A critical role for VA is to serve as a causal auditor for clinical AI. Instead of just explaining what features an AI used, VAHC research can create interfaces for “causal sanity checks” and support the development of causally-informed machine learning models that are more robust and reliable in real-world clinical practice.

5 THE FUTURE OF CAUSAL INFERENCE-ENABLED VAHC

We have outlined the profound need to use visualization methods to move beyond simple correlations and toward a causality-aware understanding of health data. We argue that the path forward is not to build automated “causality-finding machines” that replace human experts, but rather that we need to achieve a more ambitious goal: augmenting human intelligence to support causal reasoning through a new generation of visual analytics tools.

ACKNOWLEDGMENTS

We thank the reviewers for their insightful comments. This work was supported by the National Science Foundation under Grant No. 2211845.

REFERENCES

- [1] D. Borland, I. Brain, K. Fecho, E. Pfaff, H. Xu, J. Champion, C. Bizzon, and D. Gotz. Enabling longitudinal exploratory analysis of clinical covid data. In *IEEE Workshop on Visual Analytics in Healthcare (VAHC)*, pp. 19–24, 2021. [2](#)
- [2] 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 [1](#), [2](#)
- [3] M. Cascella, J. Montomoli, V. Bellini, and E. Bignami. Evaluating the feasibility of chatgpt in healthcare: an analysis of multiple clinical and research scenarios. *Journal of medical systems*, 47(1):33, 2023. [2](#)
- [4] C. R. Charig, D. R. Webb, S. R. Payne, and J. E. Wickham. Comparison of treatment of renal calculi by open surgery, percutaneous nephrolithotomy, and extracorporeal shockwave lithotripsy. *Brit. Med. J.*, 1986. [1](#)
- [5] S. Dash, S. K. Shakyawar, M. Sharma, and S. Kaushik. Big data in healthcare: management, analysis and future prospects. *Journal of big data*, 6(1):1–25, 2019. [1](#)
- [6] M. Fan, J. Yu, D. Weiskopf, N. Cao, H.-Y. Wang, and L. Zhou. Visual analysis of multi-outcome causal graphs. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):656 – 666, 2025. [2](#)
- [7] O. Giménez and A. Jonsson. The complexity of planning problems with simple causal graphs. *Journal of Artificial Intelligence Research*, 31:319–351, 2008. [2](#)
- [8] D. Gotz and D. Borland. Data-driven healthcare: challenges and opportunities for interactive visualization. *IEEE Computer Graphics and Applications*, 36(3):90–96, 2016. [1](#), [2](#)
- [9] G. Guo, E. Karavani, A. Endert, and B. C. Kwon. Causalvis: Visualizations for causal inference. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–20. ACM, New York, 2023. doi: 10.1145/3544548.3581236 [2](#)
- [10] J. S. Haukoos and R. J. Lewis. The propensity score. *Jama*, 314(15):1637–1638, 2015. [2](#)
- [11] J. Hullman and A. Gelman. Designing for interactive exploratory data analysis requires theories of graphical inference. *Harvard Data Science Review*, 3(3):10–1162, 2021. doi: 10.1162/99608F92.3AB8A587 [1](#)
- [12] 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, 2020. doi: 10.1109/TVCG.2020.3030465 [2](#)
- [13] S. A. Julious and M. A. Mullee. Confounding and Simpson’s paradox. *Bmj*, 309(6967):1480–1481, 1994. [1](#)
- [14] 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, 2021. doi: 10.1109/TVCG.2021.3114779 [2](#)
- [15] J. Pearl. *Causality*. Cambridge University Press, Cambridge, 2009. doi: 10.1017/CBO9780511803161 [2](#)
- [16] J. Qiu, K. Lam, G. Li, A. Acharya, T. Y. Wong, A. Darzi, W. Yuan, and E. J. Topol. Llm-based agentic systems in medicine and healthcare. *Nature Machine Intelligence*, 6(12):1418–1420, 2024. [2](#)
- [17] A. Z. Wang, D. Borland, and D. Gotz. An empirical study of counterfactual visualization to support visual causal inference. *Information Visualization*, 23(2):197–214, 2024. doi: 10.1177/14738716241229437 [2](#)
- [18] A. Z. Wang, D. Borland, and D. Gotz. A framework to improve causal inferences from visualizations using counterfactual operators. *Information Visualization*, 2024. doi: 10.1177/14738716241265120 [2](#)
- [19] A. Z. Wang, D. Borland, and D. Gotz. Beyond correlation: Incorporating counterfactual guidance to better support exploratory visual analysis. *IEEE Transactions on Visualization and Computer Graphics*, 31(1):776 – 786, 2025. [2](#)
- [20] 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, 2015. doi: 10.1109/TVCG.2015.2467931 [2](#)
- [21] J. Wang and K. Mueller. Visual causality analysis made practical. In *2017 IEEE Conference on visual analytics science and technology (VAST)*, pp. 151–161. IEEE, 2017. [2](#)