

An Explainable AI Approach to Large Language Model Assisted Causal Model Auditing and Development

Yanming Zhang*
Stony Brook University

Brette Fitzgibbon†
College of Marin

Dino Garofolo‡
Stony Brook University

Akshith Kota§
Stony Brook University

Eric Papenhausen¶
Akai Kaeru LLC

Klaus Mueller||
Stony Brook University
Akai Kaeru LLC

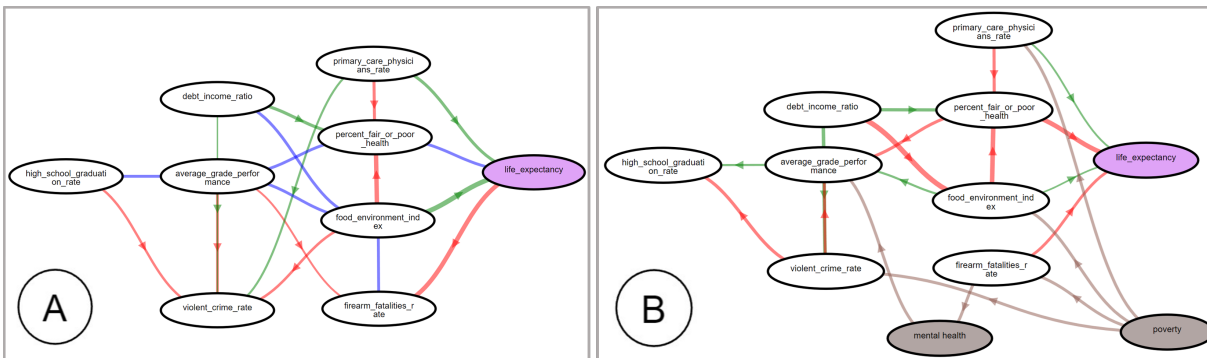


Figure 1: Our ChatGPT-powered Causal Auditor. (A) **Initial Causal Model** generated by the PC algorithm, some edges are undirected, others might be incorrect (B) **ChatGPT Refined Model** with all audited edges and new latent variables.

ABSTRACT

Causal networks are widely used in many fields, including epidemiology, social science, medicine, and engineering, to model the complex relationships between variables. While it can be convenient to algorithmically infer these models directly from observational data, the resulting networks are often plagued with erroneous edges. Auditing and correcting these networks may require domain expertise frequently unavailable to the analyst. We propose the use of large language models such as ChatGPT as an auditor for causal networks. Our method presents ChatGPT with a causal network, one edge at a time, to produce insights about edge directionality, possible confounders, and mediating variables. We ask ChatGPT to reflect on various aspects of each causal link and we then produce visualizations that summarize these viewpoints for the human analyst to direct the edge, gather more data, or test further hypotheses. We envision a system where large language models, automated causal inference, and the human analyst and domain expert work hand in hand as a team to derive holistic and comprehensive causal models for any given case scenario. This paper presents first results obtained with an emerging prototype.

1 INTRODUCTION

Causal networks are widely used in various fields, such as epidemiology [29], healthcare [9], biology [4], social sciences [7], etc. to model and understand complex systems. These networks typically

consist of a set of nodes and edges, where the nodes represent the variables and the edges represent the causal relationships between them. As the accuracy and reliability of causal networks heavily depend on the assumptions and data used in the modeling process, it is essential to verify that the causal relationships represented in the network are consistent with the actual real-world relationships.

One approach to assure the accuracy and reliability of causal networks is through auditing. Auditing involves evaluating the causal model to identify potential errors that may have been introduced during its development. Typically these activities require a qualified human domain expert who, however, might not be readily available to a model developer. As such, model auditing represents a significant bottleneck in causal model development.

We propose a potential solution to mediate the scarcity of expert knowledge in model auditing activities, namely, automated chatbots, such as ChatGPT. ChatGPT is a large language model (LLM) based on transformer neural networks trained on a massive corpus of text data, including web pages, books, and other sources. It has already demonstrated impressive capabilities in generating natural language responses to various prompts, such as completing sentences, answering questions, and even generating coherent paragraphs. We propose that these capabilities will also make it an ideal candidate to help with auditing and even developing causal networks.

But as has been widely reported ChatGPT can fall victim to hallucinations, and asking it a simple question like “Does A cause B”, even when posed within a sophisticated prompt [16], can result in erroneous statements. To lower the probability of false conclusions we have built a prompt template that asks the causal question in all possible polarities. We hypothesize that asking the same fundamental question from multiple angles can elevate confidence in the chosen relation and also expose conflicts associated with it.

Asking causal questions in multiple ways can also illuminate the broader context associated with this relation. In its justifications ChatGPT typically mentions confounder and mediator variables that surround the queried relation, but we can also specifically ask for it. Any new domain knowledge gained from this information can

*e-mail: yanming.zhang@stonybrook.edu

†e-mail: bconnolly6720@mycom.marin.edu

‡e-mail: dino.garofolo@stonybrook.edu

§e-mail: akkota@cs.stonybrook.edu

¶e-mail: epapenha@akaikaeru.com

||e-mail: mueller@cs.stonybrook.edu

then prompt targeted data collection and causal model refinement, or even splitting the model into several contextual models.

A downside of this general approach is that it generates a substantial amount of data, and to serve the mission of explainable AI the data are meant for human consumption to foster insight, confidence, and trust into the causal model. To aid this mission, we have devised a set of dedicated visualizations to make it easier for a human analyst to grasp the essence of this LLM-generated data. This paper reports on first studies we have conducted with this approach.

In the following, Sections 2 and 3 discuss background and related work, Sections 4 and 5 describe our methodology and present further examples and results. Finally, Section 6 offers a conclusion.

2 BACKGROUND

There are several strategies by which a causal directed acyclic graph (DAG) can be developed. The most rigorous one is to construct it from first principles via randomized controlled trials. This, however, is often not feasible due to cost, ethical, or practical constraints. It is also difficult to scale and restrictive on the number of researchers who can conduct such studies. A more scalable way is to *discover* a causal DAG from observational data. There are essentially two popular strategies. One is to enforce the constraint that two statistically independent variables are not causally linked and then run a series of *conditional independence tests* to construct a compliant DAG. Popular algorithms here are the PC [28] algorithm and the Fast Causal Inference (FCI) algorithm [27]. Another way is to greedily explore the space of possible DAGs [2]; one may begin with an empty graph and iteratively add directed edges to maximize a model fitness measure such as the Bayesian Information Criterion (BIC) [1, 24].

2.1 Shortfalls of Automated Causal Discovery

Causal discovery entails four common assumptions: (1) the causal structure can be represented by an acyclic DAG, (2) all nodes are Markov-independent from non-descendants when conditioned on their parents, (3) the DAG is faithful to the underlying conditional independences, and (4) the DAG is sufficient in that there is no pair of nodes with a common external cause. Unfortunately these conditions are rarely entirely fulfilled, mostly due to selection/sampling bias in the observational data. While the probability of getting a (partially) incorrect DAG can be made arbitrarily small by using enough data, it remains undefined how much data is actually needed [22, 25].

The manifestation of the problem is the discovery algorithm’s sensitivity to errors – specifically, errors in the orientation of edges. If one edge is oriented incorrectly (e.g. the DAG states that *A causes B*, when in reality *B causes A*) then it may lead to a cycle. Since causal DAGs must be acyclic, the orientation of subsequent, correctly oriented edges, may be reversed. In this way, errors in the orientation of edges tend to propagate throughout the DAG. While there has been recent work [23] that seeks to deal with this problem, progress so far has been limited. We therefore believe that the time is right to look for other solutions – the vast causal knowledge stored in text documents. While extracting causal relations from text documents is not new [36], thanks to ChatGPT they can now be conveniently unearthed with a simple prompt and for many application domains.

3 RELATED WORK

Early work on the visualization of causal relations used animated views [5, 14], while more recent research by Wang and Mueller [30] focused on visualizing causal relationships in form of Pearl’s DAG representation, with added support for analysts to interactively audit the DAG and gauge model quality with the BIC score. Their method used the PC algorithm to create the initial DAG. Conversely, Xie et al. [33] used a greedy score-based method, F-GES [21] to create the DAG and added mechanisms to enable overdraw-free graph layouts. The explanatory nature of causal DAGs readily affords What-if

analyses. Here, Xie et al. took a cohort-based approach, while Hoque and Mueller [12] allowed individuals to gain insight on how the various factors causally affected the outcome of an automated decision system and at what sensitivity. Guo et al. [11] used the Rubin model to allow analysts to assess the potential effects of treatments on population subgroups, while Wang and Mueller [31] created DAGs for the various subgroups to explain their outcomes. Along similar lines, Yan et al. [35] and Ghai and Mueller [8] used DAGs to expose fairness issues in data used for model training data whereby the latter system also allowed users to mitigate these issues. Other researchers focused more on temporal aspects of causality [13, 32] also in text narratives [3], while yet others [15, 34, 37] studied the perception of causality from a visual analyst’s point of view. Toolkits for causal inference have also become available [10, 26].

The literature on LLM-assisted causal network learning is rapidly emerging. The earliest documented attempt of using LLMs (here GPT-3) for causal analytics was published by Long et al. [19]. However, this was a rather preliminary study which focused on the structure of the prompt to reveal insight on the absence or presence of a directed edge. More recently, Kıcıman et al. [17] studied the subject far more extensively. They developed a set of elaborate prompts for GTP-3.5 and GPT-4 that would force yes/no answers to standard causal queries. They showed excellent success rates for benchmark datasets where the causal truth was known but they did not attempt to use the GPT-output to acquire additional causal and contextual knowledge. They, and other papers that followed since then (e.g. [6, 18, 20]), also did not attempt to visualize the acquired information in the realm of explainable and trustworthy AI.

4 METHODOLOGY

Fig. 2 illustrates the workflow of our method. Step 1 computes an initial causal model from the user-provided data. This model is displayed as a DAG in a suitable visual causal analytics interface, such as [31, 33]. This model can then be refined based on the outcome of a ChatGPT prompt with regards to a certain relation indicated by the user. Apart from listing the outcomes of the prompts in form of text, we also provide visualizations that seek to summarize the gist of these responses for the user.

We investigate causality at two levels. The first level seeks to identify whether a direct causal relation actually exists, while the second level dives deeper and checks for mediators and confounders. We could have used the slightly weaker term ‘covariate’ over ‘confounder’ but we found using ‘confounder’ produced richer results.

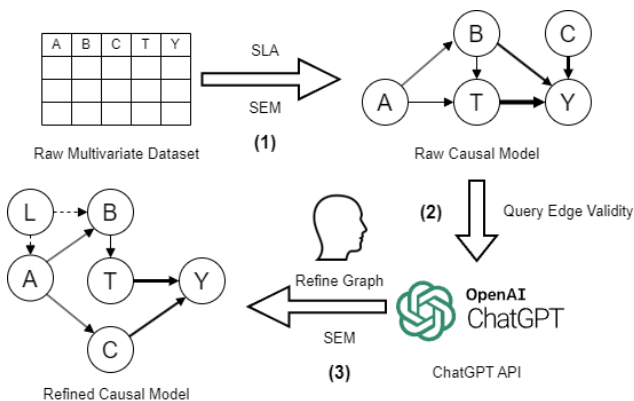


Figure 2: Workflow of our ChatGPT-powered Causal Auditor. (1) Algorithmic discovery of the initial (raw) causal model. (2) Query-driven ChatGPT-based edge commentary. (3) Analyst-initiated model refinement informed by the outcomes of steps 1 and 2.

4.1 Prompting for Direct Causal Relationships

The prompts we use are according to the following template:

Prompt: On a scale from 1 to 4, where 4 represents strong or most likely, rate the cause-and-effect relationship: higher/lower A/B causes higher/lower B/A.

This generates 10 distinct prompts, 5 each for A and B taking opposite roles, and within each of these two sets there are 4 combinations of A and B being (higher, lower) plus one relation that just asks this for a general case. For example, we might have:

Prompt: On a scale from 1 to 4, where 4 represents strong or most likely, rate the cause-and-effect relationship: *higher percent fair or poor health causes higher life expectancy.*

The below is an example for the general prompt:

Prompt: On a scale from 1 to 4, where 4 represents strong or most likely, rate the cause-and-effect relationship: *changing percent fair or poor health causes a change in life expectancy.*

This prompt resulted in the following response using GPT-4 (some part are redacted for brevity):

Response: I would rate the cause-and-effect relationship between the percentage of the population in fair or poor health and changes in life expectancy as a 4. This rating is based on the understanding that life expectancy is heavily influenced by the overall health status of a population. If a large percentage of the population is in fair or poor health, it's likely to negatively affect the average life expectancy. Poor health status could be linked to a range of factors including chronic diseases, infectious diseases, poor mental health, and lifestyle-related health conditions. All these factors can shorten individual lifespans, leading to a lower average life expectancy for the population as a whole. However, it's important to note that this is a general trend and doesn't necessarily hold in every specific case. Other factors can also impact life expectancy, such as healthcare access and quality, socioeconomic status, environmental conditions, and more...

4.2 Prompting for Confounders and Mediators

The prompt template we developed here also distinguishes among the 4 combinations that explore the effects of higher and lower levels plus one relation that just asks this for the general case. In the following we explain this template using the variables *food environment index (FEI)* and *violent crime rate* as an example.

Prompt: On a scale from 1 to 4, 4 represents strong or most likely, rate the cause-and-effect relationship 'For a county, lower *food environment index* causes higher *violent crime rate*'. Make a concise list of mediators in that relation and assign strengths to them (weak, medium, strong). Also make a concise list of confounders in that relation and assign strengths to them (weak, medium, strong).

We found that by specifying 'county' we obtained a much more concrete list that made better references to the data at hand.

This prompt resulted in the following response using GPT-4. After noting that the causal relation can be rated a 2, indicating a "some-what plausible link" it wrote (some part are redacted for brevity):

Response: GPT-4 suggested the following three mediators:

- **Poverty (Strong):** Poorer food environments (PFE) often correlate with higher poverty levels, which in turn can lead to increased crime due to socioeconomic frustrations, lack of opportunities, or the need to commit crime to survive.
- **Educational Attainment (Medium):** PFE may align with areas of lower educational attainment. Lower levels of education can contribute to higher crime rates due to fewer opportunities for gainful employment and lack of knowledge about alternatives.

- **Health Outcomes (Weak):** PFE can lead to negative health outcomes, which could potentially increase crime rates indirectly. Poor health may lead to higher levels of stress in the community, which can exacerbate social issues that contribute to crime.

GPT-4 also suggested the following three confounders.

- **Socioeconomic Status (Strong):** Lower socioeconomic status often coincides with both PFE and higher crime rates. This may make it seem like PFE cause higher crime rates when the real culprit is socioeconomic status.
- **Urban vs Rural Setting (Medium):** Urban areas tend to have both lower FEI and higher crime rates due to higher population density, greater economic disparities, and other factors. This can confound the relation between FEI and crime rates.
- **Public Policy (Weak):** Public policies can influence both the food environment (through regulations, subsidies, etc.) and crime rates (through law enforcement funding, criminal justice policies, etc.), acting as a confounding factor.

Last, GPT-4 reminded us that "correlation does not imply causation, and while these relationships can exist, they do not guarantee that a lower FEI will always lead to a higher crime rate. Other factors, such as strong community leadership, effective policing, and active non-profit organizations, can also have significant impacts on crime rates". This message hints on an important truth which applies to almost every causal relationship, even those with higher ratings. There are almost always mediators and covariates that play a role. They can contextualize the causal relation to certain sub-populations and they can also offer effective means for intervention. This has been a strong motivator for the visual interface we have designed.

4.3 Visualizing the GPT-4 Generated Text Responses

Once that one acknowledges that a binary score alone is not all that counts to judge a causal relation, the problem now becomes that there is large amount of textual information that GPT-4 generates. Even when instructed to just summarize the findings, this wealth of information can be overwhelming to a general user. In the following we describe a few dedicated visualizations we have designed to make browsing this information easier.

The first step in deriving a visualization is to identify and retrieve the key information from the text responses. For this we have made use of standard NLP techniques available in libraries like Spacy¹ and Allen NLP². Specifically, we have used these libraries to identify the various variables, mediators, and covariates, their importance ratings and their interrelations in the GPT-4 generated text responses.

In the following we describe the various charts we have created.

4.4 The Causal Debate Chart

Fig. 3 shows our summary visualization that contrasts the numerical outcomes of these 10 prompts. We call it *Causal Debate Chart* since it visually argues the strength of one variable being the cause of the other. The chart is a bidirectional bar chart where each side is headed by one of the two relation variables. In this case the left side is *Percent Fair or Poor Health* and the right side is *Life Expectancy*. The x-axis is the score assigned by ChatGPT and the length of each bar is mapped to that score. The grey bars are for the general prompt while the other bars are colored in red if the cause was a higher or increasing level of the variable or in blue if the cause was a lower or decreasing level (see color legend on the top right).

Let's evaluate the chart. We observe that for the first (grey) set of bars *Percent Fair or Poor Health* has a substantially longer bar (level 4) than *Life Expectancy* which has level 2 (level 2 is a rather weak cause in GPT-4 semantics). It means that the former wins

¹<https://spacy.io>

²<https://allenai.org/allennlp>



Figure 3: Causal Debate Chart for the relation *Percent Fair or Poor Health - Life Expectancy*, presenting an overwhelming belief that the former is the cause of the latter.




Figure 4: Causal Relation Environment Chart for the relation *Percent Fair or Poor Health - Life Expectancy*. The intensity of red and green encodes the strength of the mediators and covariates (weak, medium, strong), and the color of the cause and effect variables are alike those in Fig. 3; in this specific case they are grey.

the causal debate – it has causal dominance. *Percent Fair or Poor Health* seems to be a general cause of *Life Expectancy*

Now let us check out the other bars which play out specific level studies. There we see if ChatGPT is consistent in its logic (as opposed to hallucinating). We confirm that also here we see that high or increasing *Percent Fair or Poor Health* leads to low or decreasing *Life Expectancy*. The same is true for the opposite. Finally, the other two sets have low bars on both sides, which is to be expected if the relation as indicated by the other bars is taken as true. In some sense the Causal Debate Chart of Fig. 3 is a prime example for what we would expect from a steadfast causal relation.

4.5 The Causal Relation Environment Chart

The Causal Relation Environment Chart is a diagram that shows the complete causal relation, consisting of the two main relational variables along with a list of mediators and confounders/covariates. An example for the general *Percent Fair or Poor Health (PFPH) - Life Expectancy (LE)* relation is shown in Fig. 4.

It turns out that the more focused (low, high) relations, even those with low causal score, have quite similar mediators and covariates. But they vary in the sign. For instance to go from *low PFPH* to *high LE* positive levels of the mediators are cited, such as good access to healthcare and good health habits, while to go from *high PF-PH* to *low LE* the cited mediators are usually the opposite, like limited access to health care and poor health habits. These two cases are shown in Fig. 5 where the up and down arrows indicate the positive and negative levels, respectively.

In this context, it is interesting to examine the relation *high PFPH* to *high LE* which is deemed improbable (see Fig. 3). For this case, GPT-4 correctly responds that this relation is "counterintuitive" but then goes along with it and "considers it as a hypothetical scenario". Following this prelude, it suggests that the mediators need to be "improved" or "strong" to achieve the desired *high LE*. In essence,


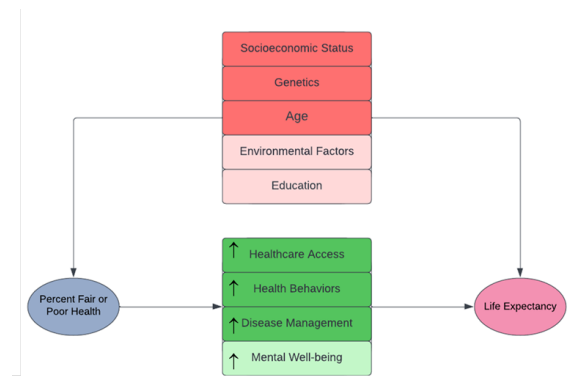


Figure 5: Causal Relation Environment Chart for two level combinations of the relation *Percent Fair or Poor Health - Life Expectancy*. The up and down arrows show the appropriate signs of the mediators.




Figure 6: Causal Relation Environment Chart for an improbable level combination of the relation *Percent Fair or Poor Health - Life Expectancy*, namely one where both variables have positive level. The up arrows in the mediators show how this improbable combination might be achieved, in form of interventions on the mediators in the direction of the arrows.

these are all potential interventions a policymaker could undertake to fix an undesirable *high PFPH* to *low LE* and make the improbable *high PFPH* to *high LE* a reality. Fig. 6 visualizes the Causal Relation Environment Chart for this case. However, given its semantics we might rename it to *Causal Intervention Chart*.

We have observed that for confounders the situation is not as clear-cut and we intend to study this further in future work.

4.6 The Confounder/Mediator Chart

The Causal Debate Chart can quickly communicate the numerical outcomes of the 10 different causal level combinations and the Causal Relation Environment Chart can focus on the confounders and mediators of one such causal relation at a time. To give a full overview we have designed a summary chart we call the *Confounder/Mediator Chart*. Fig. 7 shows an example for the causal relation *Food Environment Index (FEI)* and *Violent Crime Rate (VCR)*,

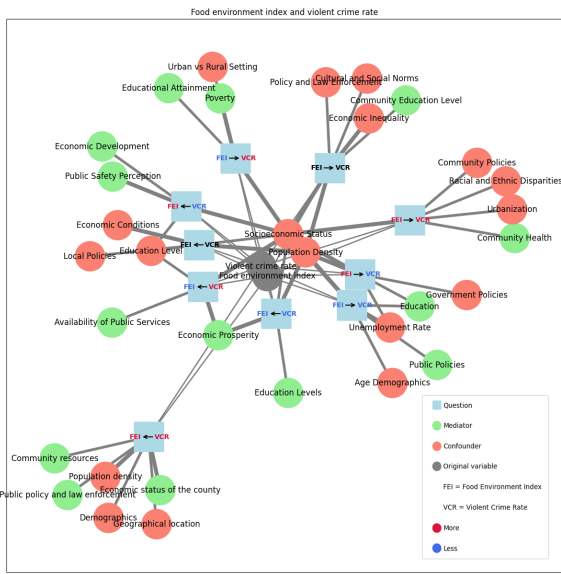


Figure 7: Confounder/Mediator Chart for the relation *Food Environment Index (FEI)* and *Violent Crime Rate (VCR)*. The legend on the bottom right explains the various symbols and colors used.

This chart represents each causal relation and level combination as a light-blue box (called a question) where the color of the inscribed text denotes the level of each of the two causal variables. The identified mediators (green disks) and confounders (red disks) emerge from their respective questions where the width of the connecting edge denotes the influential strength as identified by GPT-4.

A first observations we make in this chart is that the most common confounders are Socioeconomic Status and Population Density which are plotted at the center of the graph. These are the common themes under which this causal relation "lives".

By looking at and around the individual relations one can make out more specific scenarios only valid for one or a subset of relations. One of these is the relation on the right, high FEI and high VCR. Why could there be quality food but still high crime rate? The chart shows that racial and ethnic disparities (which usually occur in urban environments) and community policies can set the contextual stage for this (as confounders).

Interesting is also to investigate how high VCR can lead to low and high FEI, respectively. For the former, we see a mediator *Availability of Public Services*, or better, the lack thereof. Conversely, for the latter, we see mediators like *Public Policy and Law Enforcement* that ensure that FEI remains high despite high VCR.

5 FURTHER EXAMPLES AND RESULTS

We have prepared three further examples to demonstrate our method. For brevity we only show the respective Causal Debate Charts.

Food Environment Index and Violent Crime Rate. Fig. 8 shows the Causal Debate Chart for the relation between *Food Environment Index (FEI)* and *Violent Crime Rate (VCR)*. We observe that there only weak (level 1-2) support for a direct causal relation. But even though the direct causal relation is not strong, the Causal Debate Chart does not indicate that it is an absolute zero. And indeed, the Confounder/Mediator Chart in Fig. 7 brings up some interesting mediators as well as local contextual confounders which apply only for a subset of causal conditions. We have discussed some of these in the previous section.

Average Grade Performance and Food Environment Index Fig. 9 shows the Causal Debate Chart for the relation between *Average Grade Performance (AGP)* and *Food Environment Index*



Figure 8: Causal Debate Chart for the relation *Food Environment Index - Violent Crime Rate*. There does not seem a direct causal relation between these two variables for any polarity.




Figure 9: Causal Debate Chart for the relation *Average Grade Performance - Food Environment Index*. food environment index positively changes average grade performance, ChatGPT provides a medium score since a variety of factors makes the assessment complex; there is only a mild causation between these two attributes.

(FEI). There is a mild causal relation between the two, going from FEI to AGP. The chart provides coherent evidence for this. For FEI the grey bar is 2.5 times longer than the grey bar for AGP. And likewise, for both mixed (red-blue, high-low, opposite) relations the FEI bar is 2.5 or 1.5 times longer which is causally consistent, albeit not overly strong. Further analysis reveals that the relations is confounded by many other factors such as socioeconomic status, education level, cultural preferences, etc.

Average Grade Performance and High School Graduation Rate Fig. 10 shows the Causal Debate Chart for the relation between *Average Grade Performance (AGP)* and *High School Graduation Rate (HGR)*. There is a strong cause and effect relation among the two. This time there is positive causal relationship between the two variables, namely (high/low) AGP leading to (high/low) HGR.




Figure 10: Causal Debate Chart for the relation *Average Grade Performance - High School Graduation Rate*. We observe that Average Grade Performance positively changes High School Graduation rate.

5.1 Evolved Causal Model

We used the new insight gained from the GPT-4 based analysis to evolve the initial causal model we had computed via the PC-algorithm. The initial model is shown in Fig. 1A, and the evolved model is shown in Fig. 1B. Since the model evolution is not the main focus of this paper – the visualization of the GPT-4 text is – we will not delve much into it and only present some results.

In our initial experiment we used Poverty as a proxy for the various nuanced confounder variables we identified. As shown in Fig. 12 and Fig. 11, the audition with GPT-4 increased the overall graph fit by 70.69% (BIC score from -3673.4 to -6270.3), and for most attributes the BIC score improved.




Figure 11: Attribute BIC scores before and after various stages of the audit, presenting the fit score for each attribute (higher is better). The blue bars are the scores for the initial model; the orange bars are the scores for the audited model but without adding the data for the newly identified mediators and confounders; the green bars are the bars after recomputing the model with the data of the newly introduced variables. For some variables this last update caused a small score reduction but the reduction was only minor.




Figure 12: Overall graph BIC scores before and after various stages of the audit, presenting the fit score for each attribute (higher is better),

5.2 A Note on GPT-4 Accuracy

We found that GPT-4 was significantly more accurate than GPT-3.5, which has been reported widely. Concretely, in our experiments GPT-4 produced the correct direction 103/110 times (94%) and produced a number 109/110 times (99%). For the “greater/lower A/B causes greater/lower B/A” queries, the 42 queries that proposed the inverse of the correct relationship (e.g., “greater average grade performance causes greater violent crime rate”) had a range of 1 to 3 and a median of 1. The 68 queries that proposed the correct relationship, including “a change in A/B causes a change in B/A”, had a range of 1 to 4 and a median of 2. Based on these observations we feel that GPT-4 is a valuable causal assistant but still requires human supervision.

6 CONCLUSION

We feel that the integration of ChatGPT as a Causal Auditor represents a significant stride forward in the causal inference field. By facilitating interactive model editing through easily accessible textual causal knowledge, this approach holds great promise in advancing our capacity to build better causal models and expand the frontiers of our understanding across various disciplines.

To best utilize the vast knowledge GPT-4 can bring across we need techniques that allow humans to quickly get the gist of it and then delve in further. This is a typical “Overview first, zoom and filter, then details-on-demand” paradigm – the visual information seeking mantra. We have presented a first step toward this premise. In the future we plan to add more information richness to our Causal Relation Environment Chart such that it can better capture the intricate information conveyed in the text. Also, the Confounder/Mediator Chart is still too cluttered and better layout algorithms are needed. Finally, we also need to study the interface with actual human users and observe how they use the system in their model auditing and development activities.

REFERENCES

- [1] K. P. Burnham and D. R. Anderson. Multimodel inference: understanding aic and bic in model selection. *Sociological Methods & Research*, 33(2):261–304, 2004.
- [2] D. M. Chickering. Optimal structure identification with greedy search. *Journal of Machine Learning Research*, 3(Nov):507–554, 2002.
- [3] A. Choudhry, M. Sharma, P. Chundury, T. Kapler, D. W. Gray, N. Ramakrishnan, and N. Elmqvist. Once upon a time in visualization: Understanding the use of textual narratives for causality. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1332–1342, 2020.
- [4] T. N. Dang, P. Murray, J. Aurisano, and A. G. Forbes. Reactionflow: An interactive visualization tool for causality analysis in biological pathways. In *BMC Proceedings*, vol. 9, pp. 1–18, 2015.
- [5] N. Elmqvist and P. Tsigas. Causality visualization using animated growing polygons. In *IEEE Symposium on Information Visualization*, pp. 189–196, 2003.
- [6] J. Gao, X. Ding, B. Qin, and T. Liu. Is chatgpt a good causal reasoner? a comprehensive evaluation. *arXiv preprint arXiv:2305.07375*, 2023.
- [7] J. Gerring. Causation: A unified framework for the social sciences. *Journal of Theoretical Politics*, 17(2):163–198, 2005.
- [8] 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, 2022.
- [9] T. A. Glass, S. N. Goodman, M. A. Hernán, and J. M. Samet. Causal inference in public health. *Annual Review of Public Health*, 34:61–75, 2013.
- [10] G. Guo, M. Glenski, Z. Shaw, E. Saldanha, A. Endert, S. Volkova, and D. Arendt. Vaine: Visualization and ai for natural experiments. In *IEEE VIS*, pp. 21–25, 2021.
- [11] G. Guo, E. Karavani, A. Endert, and B. C. Kwon. Causalvis: Visualizations for causal inference. In *ACM CHI*, pp. 1–20, 2023.
- [12] 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, 2021.
- [13] 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.
- [14] N. R. Kadaba, P. P. Irani, and J. Leboe. Visualizing causal semantics using animations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1254–1261, 2007.
- [15] A. Kale, Y. Wu, and J. Hullman. Causal support: modeling causal inferences with visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1150–1160, 2021.
- [16] E. Kıcıman, R. Ness, A. Sharma, and C. Tan. Causal reasoning and large language models: Opening a new frontier for causality. *arXiv preprint arXiv:2305.00050*, 2023.

- [17] E. Kıcıman, R. Ness, A. Sharma, and C. Tan. Causal reasoning and large language models: Opening a new frontier for causality, 2023.
- [18] S. Long, A. Piché, V. Zantedeschi, T. Schuster, and A. Drouin. Causal discovery with language models as imperfect experts. *arXiv preprint arXiv:2307.02390*, 2023.
- [19] S. Long, T. Schuster, A. Piché, S. Research, et al. Can large language models build causal graphs? *arXiv preprint arXiv:2303.05279*, 2023.
- [20] A. Nam, C. Hughes, T. Icard, and T. Gerstenberg. Show and tell: Learning causal structures from observations and explanations. 2023.
- [21] J. Ramsey, M. Glymour, R. Sanchez-Romero, and C. Glymour. A million variables and more: the fast greedy equivalence search algorithm for learning high-dimensional graphical causal models, with an application to functional magnetic resonance images. *International Journal of Data Science and Analytics*, 3:121–129, 2017.
- [22] J. M. Robins, R. Scheines, P. Spirtes, and L. Wasserman. Uniform consistency in causal inference. *Biometrika*, 90(3):491–515, 2003.
- [23] R. Y. Rohekar, S. Nisimov, Y. Gurwicz, and G. Novik. Iterative causal discovery in the possible presence of latent confounders and selection bias. *NIPS*, 34:2454–2465, 2021.
- [24] G. Schwarz. Estimating the dimension of a model. *The Annals of statistics*, pp. 461–464, 1978.
- [25] C. Shalizi. *Advanced Data Analysis from an Elementary Point of View, chapter 22.6 Limitations on Consistency of Causal Discovery*. 2021.
- [26] Y. Shimoni, E. Karavani, S. Ravid, P. Bak, T. H. Ng, S. H. Alford, D. Meade, and Y. Goldschmidt. An evaluation toolkit to guide model selection and cohort definition in causal inference. *arXiv preprint arXiv:1906.00442*, 2019.
- [27] P. Spirtes. An anytime algorithm for causal inference. In *International Workshop on Artificial Intelligence and Statistics*, pp. 278–285, 2001.
- [28] P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. *Causation, Prediction, and Search, Second Edition*. MIT press, 2000.
- [29] J. P. Vandenbroucke, A. Broadbent, and N. Pearce. Causality and causal inference in epidemiology: the need for a pluralistic approach. *International Journal of Epidemiology*, 45(6):1776–1786, 2016.
- [30] 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.
- [31] J. Wang and K. Mueller. Visual causality analysis made practical. In *IEEE VAST*, pp. 151–161, 2017.
- [32] J. Wang and K. Mueller. Domino: Visual causal reasoning with time-dependent phenomena. *IEEE Transactions on Visualization and Computer Graphics*, 2022.
- [33] 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, 2020.
- [34] C. Xiong, J. Shapiro, J. Hullman, and S. Franconeri. Illusion of causality in visualized data. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):853–862, 2019.
- [35] J. N. Yan, Z. Gu, H. Lin, and J. M. Rzeszutarski. Silva: Interactively assessing machine learning fairness using causality. In *ACM CHI*, pp. 1–13, 2020.
- [36] J. Yang, S. C. Han, and J. Poon. A survey on extraction of causal relations from natural language text. *Knowledge and Information Systems*, 64(5):1161–1186, 2022.
- [37] C.-H. E. Yen, A. Parameswaran, and W.-T. Fu. An exploratory user study of visual causality analysis. In *Computer Graphics Forum*, vol. 38, pp. 173–184, 2019.