An Intelligent Assistant for Mediation Analysis in Visual Analytics

Chi-Hsien Yen
Yu-Chun Yen
Wai-Tat Fu
Computer Science Department
University of Illinois at Urbana-Champaign, Urbana, USA
{cyen4,yyen4,wfu}@illinois.edu

ABSTRACT

Mediation analysis is commonly performed using regressions or Bayesian network analysis in statistics, psychology, and health science; however, it is not effectively supported in existing visualization tools. The lack of assistance poses great risks when people use visualizations to explore causal relationships and make data-driven decisions, as spurious correlations or seemingly conflicting visual patterns might occur. In this paper, we focused on the causal reasoning task over three variables and investigated how an interface could help users reason more efficiently. We developed an interface that facilitates two processes involved in causal reasoning: 1) detecting inconsistent trends, which guides users' attention to important visual evidence, and 2) interpreting visualizations, by providing assisting visual cues and allowing users to compare key visualizations side by side. Our preliminary study showed that the features are potentially beneficial. We discuss design implications and how the features could be generalized for more complex causal analysis.

CCS CONCEPTS

• **Human-centered computing → Visualization systems and tools**.

KEYWORDS

Intelligent Visualization Tool, Mediation Analysis, Causal Reasoning

ACM Reference Format:
Chi-Hsien Yen, Yu-Chun Yen, and Wai-Tat Fu. 2019. An Intelligent Assistant for Mediation Analysis in Visual Analytics. In *24th International Conference on Intelligent User Interfaces (IUI '19), March 17–20, 2019, Marina del Rey, CA, USA*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3301275.3302325

IUI '19, March 17–20, 2019, Marina del Rey, CA, USA

ACM ISBN 978-1-4503-6272-6/19/03…$15.00
https://doi.org/10.1145/3301275.3302325

1 INTRODUCTION

Interactive visualizations are effective in finding extreme values, examining data distributions, or discovering trends. However, when using visualizations to reason about causal relationships, many pitfalls might exist and thus mislead users. Imagine a journalist, Alice, who is interested in investigating whether gender bias exists in the admission process of a school, uses a common visualization tool to analyze relevant data. She first uses *Gender* and *Admission* as variables for visualizations and finds that the female applicants have a lower admission rate than the male applicants do. Without further analysis, the visual pattern might lead Alice to believe that gender bias exists. But, the relationship between *Gender* and *Admission* may be *mediated* by a third variable, the *Department* an applicant applies to. For example, female students might apply to more competitive departments, which have lower admission rates regardless of the gender of an applicant.

Now, assuming that Alice is aware of the possible effects of *Department*, she visualizes the trends between *Gender* and *Department* and between *Department* and *Admission* and finds that both of the trends support the aforementioned mediation hypothesis. However, even when these pairwise correlations exist, it is still possible that the gender bias exists, as the correlation between *Department* and *Admission* might be caused by the confounding effect of *Gender*. Therefore, more thorough visual analysis is needed to draw causal inferences.

To perform such mediation analysis, data experts commonly run statistical models, such as regressions or Bayesian analysis, by program scripts or specialized tools. However, existing visualization tools for the general public mostly focus on making the process of generating a chart simple and fast, while lacking effective assisting in causal analysis. As open data and visualization tools are increasingly available, the lack of support poses great risks as more and more people are making decisions based on visualizations, and similar reasoning situations as described in our examples frequently occur in the real world. Our first example resembles the well-known Simpson's Paradox phenomenon in Berkeley Graduate Admission study [3, 11]; other real-world examples include racial bias in California lawsuits [9], effectiveness of education expenditures [5], or baseball performance [10]. As such analysis is widely needed, we aim to design an intelligent visualization tool that supports mediation analysis.

As shown in the example, a single visualization is not sufficient to infer causal relationships, but multiple visualizations have to be interpreted and compared to gain a full understanding. Therefore, we propose a visualization interface design that not only facilitates the visualizing process but also seeks relevant visualizations and helps users compare the trends across them, which is less investigated in prior research. Specifically, our design facilitates two crucial processes in causal reasoning: first, it automatically detects whether the trend shown in a visualization is consistent or not in other visualizations; second, the system provides visual cues to help users interpret and compare visualizations. A proof-of-concept study is conducted to show preliminary user feedback. We then discuss design implications for future visualization interface that assists causal analysis.

2 RELATED WORK

Causal analysis has been studied intensively in statistics, psychology, and health science. In the past decades, several mathematical frameworks for causal inference have been developed, such as regression-based approaches [2, 7, 18], or Bayesian network analysis pioneered by Judea Pearl [8, 12]. Program libraries or tools based on these frameworks were also developed [15, 16]. However, using these frameworks and tools require expert knowledge in statistics; therefore, people without statistics background mainly rely on visualizations to gain insights from data, and this motivates us to investigate how visualization interfaces could be designed to support causal reasoning.

To assist exploratory data analysis, various ideas of interface design have been proposed, such as recommending important visualizations or detecting cognitive biases [4, 14]. For example, Vartak et al. [13] presented a system that recommends 'interesting' visualizations, which is defined by how much deviation is shown in a visualization. However, in the context of causal reasoning, the algorithm may not be suitable because the lack of a deviation could actually be important evidence to refute an inference.

To address the issues of Simpson's Paradox (SP), which frequently occurs in causal reasoning, Guo et al. [6] developed algorithms to detect SP within large-scale datasets. Armstrong and Wattenberg [1] designed a comet chart to visualize SP. Our work shares the same goal of helping users detect and avoid reasoning pitfalls, while we consider not only SP, but also other possible causal models in mediation analysis. In addition, we propose interface features that facilitate searching relevant evidence and interpretation across multiple visualizations, which could be extended to more complex causal analysis not limited to SP.

3 MEDIATION ANALYSIS USING VISUALIZATION

In this section, we examine how visualizations could be used in a mediation analysis and derive design guidelines for our interface. A mediation model hypothesizes that an independent variable (X) influences a mediating variable (M), which in turn influences a dependent variable (Y). To verify the hypothesis, three direct causal relationships are required to be inspected: 1) whether X influences M, 2) whether M influences Y, and 3) whether X directly influences Y.

Considering all combinations of whether each of the relationships exists, 8 ($= 2 \times 2 \times 2$) different causal models can be generated. We exclude two causal models where there is no direct or indirect relationship between X and Y, but keep the one where all relationships do not exist for baseline model comparison. When being visualized, datasets with different causal models would show various patterns across different visualization settings. Figure 1 summarizes whether each correlation or conditional correlation exists for each causal model. A green check mark means the correlation exists and thus a trend would be seen on the corresponding visualization; on the other hand, a red cross sign means the correlation would not be seen. The names of each causal model are commonly used in the statistics field. Examples of visualizations for each causal model are provided. Note that we illustrate the correlation by showing a decreasing bar chart for the sake of simplicity. Depending on the actual dataset, it may be an increasing bar chart, an increasing trend line on a scatter plot, or other visual patterns that show correlation.

The first row shows the existence of the correlation between X and M (denoted $X - M$). Because Y does not influence these two variables, whether one would see a correlation between X and M matches whether the causal relationship from X to M exists.

Two key patterns in the rows of $(M - Y)$ and $(X - Y)$ are worth noting. First, for causal model 3, although M does not directly influence Y, one would see a correlation between them in the visualization when X is not controlled. Such spurious correlation occurs when two variables (M and Y in this case) are influenced by a common cause (X, also called confounding factor). Therefore, to reason whether M actually influences Y, i.e., to differentiate model 3 and model 6, one needs to visualize the trend of M and Y when X is controlled (denoted $(M - Y|X)$). This can be achieved by partitioning the data using X first and visualizing the trend between Y and M within each subgroup. By comparing whether a correlation is shown in each subgroup, model 3 and 6 can be distinguished.

Second, the $(X - Y)$ visualization of causal model 5 would show a correlation despite that X does not directly influence Y, represents the classic mediation model. To reason about whether X directly influences Y, i.e., to differentiate model 5 and model 6, one needs to examine the trend of X and Y when M is controlled ($(X - Y|M)$) using the same visualization for $(M - Y|X)$ as described above. Here, one should compare the values *across* groups. If the values do not differ across groups, it means Y does not change with respect to X when M is fixed, which supports model 5; otherwise, model 6 is supported.

The pattern comparisons described here implies an important design guideline for assisting mediation analysis in visualization systems: the visualization with all three variables is a necessity, because the consistency status of correlations $(M - Y)$ and $(X - Y)$ are the key to distinguish various possible causal models. If a correlation is inconsistent, it implies the underlying data may support causal model 3 or 5 (where Simpson's Paradox occurs); otherwise, a consistent trend implies other models are supported. The key reasoning point motivates us to design an interface that 1) automatically detects the (in)consistency of the correlations and 2) facilitates the interpretation process by allowing users to compare important visualizations side by side with assisting visual cues.

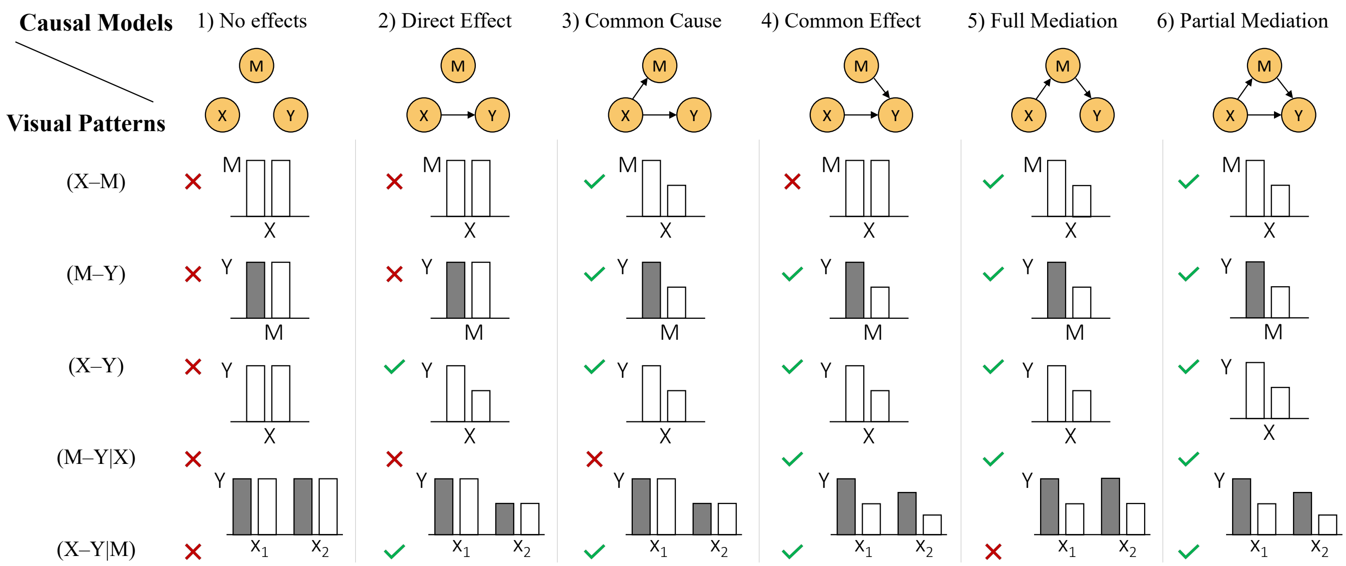

Figure 1: The table summarizes whether each correlation and conditional correlation exists for each causal models. Visualization examples are provided to illustrate how they can be used to distinguish causal models.

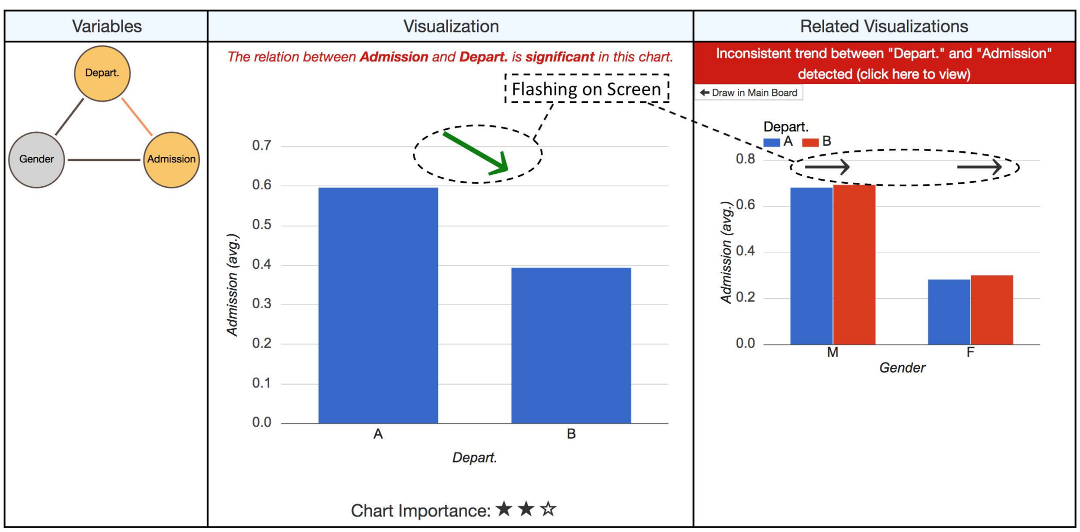

Figure 2: The interface design of our visualization system. The related visualization panel shows an inconsistent trend compared to the main visualization, guiding user's attention to important visual patterns.

4 SYSTEM DESIGN

In this section, we describe our interface design. The features could guide users to attend to the right visualizations, recognize critical patterns, integrate across visualizations to make important inferences. While we focus on mediation analysis here, the features could be generalized to more complex causal reasoning tasks, which will be discussed later.

4.1 Interface Design

Figure 2 shows a screenshot of our developed visualization system. The interface consists of three main panels: 1) Variable Panel, which shows the three variables in the dataset. The positions of the variables are organized in the same way as in Figure 1. Users can directly click on the variables to select or unselect the variables to be visualized, or click on an edge to quickly visualize the associated pair of variables.

2) Visualization Panel: when a user selects a set of variables, the system automatically generates a visualization using the selected variables. We put the variable that might be influenced by others on the Y-axis, and the other variable(s) on the X-axis, as it is a common practice to put the explanatory variable on X-axis. Figure 2 shows a user selects *Department* and *Admission* variables (highlighted in yellow). The bar chart in the visualization is then plotted to show the difference in admission rates between departments. When only one variable is selected, a histogram or an aggregated value is plotted. Additional visual cues such as the arrow, chart importance score, and helping text shown on the figure will be explained shortly.

3) Related Visualizations Panel: when a visualization is plotted in the Visualization Panel, additional visualizations are recommended on the rightmost panel. In Figure 2, the visualization using all of the three variables is recommended, with a red warning header stating that an inconsistent trend is detected. When a consistent trend is found, it would also be shown in the panel with green headers, which provides supporting evidence that might increase users' confidence. Users can click on the header to hide or show the recommended visualization, allowing them to focus on the middle main visualization or compare the two visualizations side by side. How the system recommends a visualization is explained below.

4.2 Detecting (In)Consistent Trends

When more than one variables are selected (and automatically plotted), regressions are run in the background to detect consistency or inconsistency of the visualized trend in other visualizations. Specifically, when two variables are selected (such as X and Y, or M and Y), we first run a regression using the two variables to see whether a correlation exists. Then, we run a regression that includes the third variable (regress Y on both X and M) to see whether the significance status of the first regression result remains the same when the third variable is controlled. The visualizations corresponding to the additional regressions are provided in the Related Visualizations Panel. In Figure 2, the system found that the trend between *Department* and *Admission* disappears when *Gender* is controlled. Based on these trends, one could infer that only the common causal model (model 3) is supported (refer to the table in Figure 1). Note that when X and M are selected, we do not regress M on X with Y controlled, because it may actually produce misleading results (e.g., an explaining-away phenomenon would occur in common effect model [17]) and does not help in differentiating causal models.

On the other hand, when all of the three variables are selected and plotted, in addition to running a regression on the three variables (regress Y on X and M), we also run separate regressions (regress Y on M and regress Y on X) to see whether any of the trends is inconsistent. This helps differentiating, for example, causal model 2 and 3.

4.3 Generating Visual Cues to Assist Interpretation

Besides recommending important related visualizations, the interface also provides visual cues that help interpretation. First, arrows are drawn to show the trends based on regression results. For example, the green downwards arrow in Figure 2 shows that the decreasing of the bars is statistically significant. However, the black horizontal arrows on the recommended visualization show that the admission rates are not significantly different within each subgroup. When a recommended visualization is shown, the arrows that are being compared will be flashing, which is useful to draw attention because other non-related arrows might also be drawn in some cases. A helping text is also generated based on the results of the regressions. Natural language sentences are generated by templates "The relation between [variable] and [variable] is (in)significant in the chart". The text is shown on the top of the Visualization Panel, which is spatially closer to the headers in the Related Visualization Panel; therefore, users can compare the text easier.

Furthermore, at the bottom of Visualization Panel, a rating of chart importance is provided by the system. The goal of the rating is to help users prioritize the visualizations. We consider two factors in scoring: the number of variables being visualized and the number of inconsistent trends detected. A higher score is given when the number of variables being visualized is higher, as the trend shown in the visualization is more likely "correct" when more variables are controlled. Also, if an inconsistent trend is found in another visualization with more variables, as the case in Figure 2, the importance score is lower.

5 EVALUATION AND DISCUSSION

A preliminary user study is conducted to get user feedback on the usability and effectiveness of the developed visualization interface. Five graduate students in a research university were recruited through social networks and participated in-lab interviews individually.

We adopted the within-participant experiment design to understand how participants behave with and without the assisting features. In the control condition, a baseline interface was used, which disabled and hid all of the assisting features, including the Related Visualizations Panel and visual cues (arrows, helping text, and chart importance score). The Variables Panel remained the same and visualizations were also automatically plotted in Visualization Panel. On the other hand, all of the described features were enabled in the treatment condition. The order of the conditions was counterbalanced.

The participants were asked to infer causal relationships for two datasets, one in each condition. We chose model 3 and 5 as the underlying causal model in the two datasets as they are commonly found in real-world situations. After the participants completed the reasoning tasks, they gave answers on whether X influences M, M influences Y, and X directly influences Y, along with their confidence level for each answer (5-point Likert scale).

The preliminary showed that, in the control condition, participants made 3 wrong answers in total, while 2 are made in the treatment condition. The average confidence level of their answers was 3.8 in the control condition, and 4.1 in the treatment condition. The small number of participants does not allow us to statistically compare the performance; however, our interface shows potential to decrease inference errors and increase the confidence level. Note that the goal of this short paper is not to perform statistical tests to evaluate its effectiveness at this point; rather the goal is to demonstrate the proof-of-concept system and how useful features can be developed for users in such scenarios, which is missing from existing systems. A large-scale rigorous user study is required to understand how performance is influenced by each feature and other factors, such as confirmation biases.

At the end of the interview, we asked the participants which interface they preferred and what were the most/least useful features. As expected, interface with assisting features was preferred by all participants as more information was provided. In addition, most of them stated that the ability to compare conflicting visualizations side by side is particularly useful. Two of them also mentioned that the chart importance score helped them prioritize reasoning efforts.

Future work includes extending the proposed interface features for more variables and more complex causal reasoning tasks. For example, the Variable Panel is currently a predefined graph with only three variables. The panel can be extended to an editable directed graph, where users can specify how each variable influences others based on their hypotheses. As directed graphs are widely used to represent causal models, it is intuitive to draw even when many variables are analyzed. Second, based on the specified causal graph, path analysis or structural equation modeling could be used to extend our regression-based approach. The key idea remains to be providing visualizations that contain inconsistent trends to the visualized trends. The system could also suggest alternative models and provide supporting visual patterns across visualizations.

6 CONCLUSION

We develop an intelligent visualization interface for mediation analysis on three variables. Based on the insights from comparisons among causal models, we propose interface features that facilitate two critical processes in causal reasoning: 1) detecting inconsistent or consistent trends across multiple visualizations, and 2) assisting the interpretation process of visualizations by providing visual cues and allowing users to compare conflicting patterns side by side. Our preliminary user study showed the potential of the interface. The proposed visualization interface features could be extended to assist more complex causal analysis as discussed.

REFERENCES

[1] Zan Armstrong and Martin Wattenberg. 2014. Visualizing statistical mix effects and simpson's paradox. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 2132–2141.

[2] Reuben M Baron and David A Kenny. 1986. The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of personality and social psychology* 51, 6 (1986), 1173.

[3] Peter J Bickel, Eugene A Hammel, and J William O'Connell. 1975. Sex bias in graduate admissions: Data from Berkeley. *Science* 187, 4175 (1975), 398–404.

[4] David Gotz, Shun Sun, and Nan Cao. 2016. Adaptive contextualization: Combating bias during high-dimensional visualization and data selection. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM, 85–95.

[5] D Guber. 1999. Getting what you pay for: The debate over equity in public school expenditures. *Journal of Statistics Education* 7, 2 (1999).

[6] Yue Guo, Carsten Binnig, and Tim Kraska. 2017. What you see is not what you get!: Detecting simpson's paradoxes during data exploration. In *Proceedings of the 2nd Workshop on Human-In-the-Loop Data Analytics*. ACM, 2.

[7] Andrew F Hayes. 2009. Beyond Baron and Kenny: Statistical mediation analysis in the new millennium. *Communication monographs* 76, 4 (2009), 408–420.

[8] Judea Pearl. 2009. *Causality*. Cambridge university press.

[9] Michael L Radelet and Glenn L Pierce. 1991. Choosing those who will die: Race and the death penalty in Florida. *Fla. L. Rev.* 43 (1991), 1.

[10] Ken Ross. 2007. *A mathematician at the ballpark: Odds and probabilities for baseball fans*. Penguin.

[11] Edward H Simpson. 1951. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society. Series B (Methodological)* (1951), 238–241.

[12] Peter Spirtes, Clark N Glymour, and Richard Scheines. 2000. *Causation, prediction, and search*. MIT press.

[13] Manasi Vartak, Sajjadur Rahman, Samuel Madden, Aditya Parameswaran, and Neoklis Polyzotis. 2015. S ee DB: efficient data-driven visualization recommendations to support visual analytics. *Proceedings of the VLDB Endowment* 8, 13 (2015), 2182–2193.

[14] Emily Wall, Leslie M Blaha, Lyndsey Franklin, and Alex Endert. 2017. Warning, bias may occur: A proposed approach to detecting cognitive bias in interactive visual analytics. In *IEEE Conference on Visual Analytics Science and Technology (VAST)*.

[15] Jun Wang and Klaus Mueller. 2016. The visual causality analyst: An interactive interface for causal reasoning. *IEEE transactions on visualization and computer graphics* 22, 1 (2016), 230–239.

[16] Jun Wang and Klaus Mueller. 2017. Visual causality analysis made practical. In *Proc. IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE.

[17] Timothy D Wilson and Daniel T Gilbert. 2008. Explaining away: A model of affective adaptation. *Perspectives on Psychological Science* 3, 5 (2008), 370–386.

[18] Xinshu Zhao, John G Lynch Jr, and Qimei Chen. 2010. Reconsidering Baron and Kenny: Myths and truths about mediation analysis. *Journal of consumer research* 37, 2 (2010), 197–206.