Characterizing the Performance of Counterfactual and Correlation Guidance via Dataset Perturbations

A. Z. Wang[1], D. Borland[2], and D. Gotz[1]

[1]The University of North Carolina at Chapel Hill, USA [2]RENCI, The University of North Carolina at Chapel Hill, USA

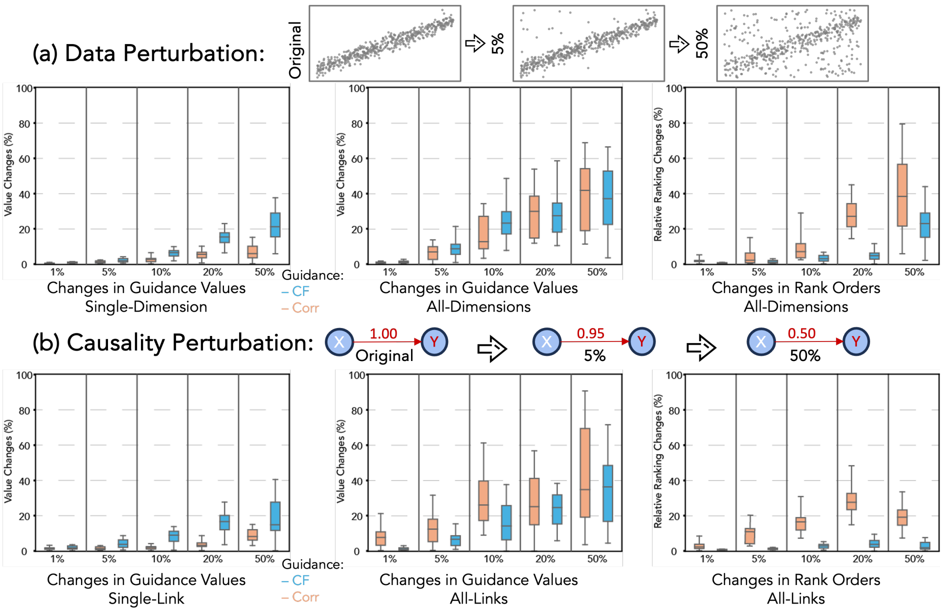

Figure 1: *Results of (a) data perturbation, and (b) causality perturbation, on counterfactual (CF), and correlation (Corr), guidance measures, characterized by percent change in the guidance values, and by changes in suggested variable rank order.*

Abstract

Guidance methods are often employed in visual analytics systems to help users navigate complex datasets and discover meaningful insights. Guidance based on correlation is a common method that can steer users towards closely related variables. However, recent work has shown that guidance based on counterfactual subsets can more effectively capture and surface causal relationships. In this work we further explore these guidance methods by characterizing their performance by systematically introducing perturbations in both the data points generated from a ground truth causal graph, and the causal relationships in the graph itself. Our results indicate that while both guidance types exhibit similar sensitivity to global data point perturbations, counterfactual guidance can better capture perturbations affecting only a single dimension, and more effectively reflect changes in causal link strengths, indicating an improved ability to capture narrow data changes and causal relationships.

1. Introduction

Guidance techniques are often used in visual analytics to suggest interesting data subsets, relationships, or views [CGM*16], with correlation-based guidance the most widely used method, e.g., [BHT20]. Recently, guidance based on counterfactual subsets, rooted in causal inference, has emerged as an alternative [WBG25]. Counterfactual guidance enables the examination of "what-if" scenarios, assessing how an outcome might change if certain factors were different, relying on creating counterfactual data subsets based on the similarity of user-defined data subsets to certain excluded data. This approach aims to effectively capture causal relationships rather than relying on statistical associations.

Although this initial previous work demonstrated the conceptual differences between the two approaches, along with some of the advantages of counterfactual guidance, more work is necessary to fully characterize them. Real-world data is often noisy, incomplete, or subject to shifts over time. Furthermore, the underlying causal structures of the data may not be perfectly known or stable. Therefore, understanding how these different guidance types react to variations in the data and the underlying causal model is crucial for VA system designers and analysts relying on such guidance.

This poster addresses this gap by empirically evaluating and comparing the performance characteristics of counterfactual and correlation guidance under controlled perturbations. We utilize synthetic datasets generated from known causal graphs, enabling us to precisely manipulate both the resulting data points and the underlying causal relationship strengths. Our findings highlight the distinct behaviors of correlation and counterfactual guidance, demonstrating the ability of counterfactual guidance to capture narrow data changes as well as larger structural changes, shedding light on when and how different guidance types might be more appropriate.

2. Guidance Computation

We briefly introduce how correlation and counterfactual guidance are typically conceptualized and computed. For the purpose of this study, we assume the user filtered the dataset based on some initial variables of interest. In practice, a guidance value is computed for each non-filter variable in a dataset to determine which variables may be most useful to explore in the context of the filter variables. The data subsets were computed via point- and set-based similarity. For more details, please see the employed *Co-op* library [WBG24].

Correlation Guidance aims to identify variables that exhibit strong linear or monotonic associations within the dataset. It is computed by measuring the correlation coefficients between the user-filtered data subset IN and the excluded subset EX.

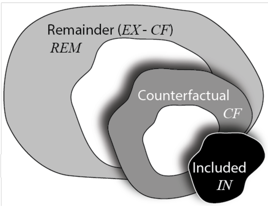

Figure 2: *Relations of the employed data subsets.*

Counterfactual Guidance leverages principles from causal inference to provide insights based on "what-if" questions, using the user-filtered subset IN as reference, and calculating a counterfactual subset CF from EX whose data points are similar to IN based on the other variables in the data. The remainder of the data points, $EX - CF$, is termed REM. The guidance value is calculated by measuring the similarities between IN, CF, and REM.

3. Method

Following prior work [WBG25], we randomly generated 20 causal graphs, each consisting of 20 data variables, where 19 variables were assigned a random causal strength in [0, 1] to the last outcome variable. These graphs provide ground truth for the causal relationships and enable controlled manipulations. We synthesized 20 datasets using continuous Gaussian 2D distributions with these causal graphs via CausalSynth [WWBG24]. To simulate real-world user behavior, we generated random filters based on randomly selecting one to five of the dataset variables to create an IN subset for each dataset, and computed the **original guidance values** for both correlation and counterfactual approaches.

Data perturbation was performed by randomly selecting a percentage of data points for the discrete steps [1%, 5%, 10%, 20%, 50%] methods, and randomizing the values of those points for specified variables. Two scopes were used for variable specification. We calculated the proportional data using simple random sampling from all data points. **One-dimensional** perturbations were applied only for a single randomly selected variable. **All-dimensional** perturbations were applied independently to all variables in the dataset.

Causality perturbation reduced the causal link strengths of the causal graphs, simulating changes in the underlying causal structure. We varied the same percentage changes applied to the causal link strengths from 1% to 50%. We also employed two perturbation scopes; **Single-Link** perturbation was applied only to the strength of a single randomly selected causal link in the causal graph, **All-Links** perturbations were applied simultaneously to the strengths of all causal links in the graph.

Each of the 20 perturbations was performed 50 times. For each perturbation we compared the guidance value computed for each variable to the guidance computed from the corresponding original dataset. We measured **changes in guidance value** by the absolute difference of the computed values, and **changes in rank order** by the relative differences of guidance ranking offset distance [WBG25, RY98]. We discuss changes in guidance values for one-dimensional and single-link perturbations, and changes in both guidance values and rank orders for all-dimensional and all-link perturbations.

4. Results

One-Dimensional Data Perturbation (Figure 1 (a)): The counterfactual (CF) guidance values changed more obviously than the correlation (Corr) guidance values even with low percentages of perturbed points. This highlights the ability of CF guidance to capture data changes for specific dimensions.

All-Dimensional Data Perturbation (Figure 1 (a)): For both methods, the magnitude of changes in guidance values generally increased with the percentage of perturbed points. This suggests that both methods are similarly capable of detecting global noise or shifts distributed across the entire dataset. But for the rank orders, we observed that CF had fewer changes, especially for smaller perturbation percentages. This may imply that CF guidance can robustly infer causal relationships in the presence of noise.

Single-Link Causality Perturbation (Figure 1 (b)): We found similar observations to those for the one-dimensional data perturbation. The CF guidance values changed more obviously, indicating potential ability to identify the impact of specific changes in causal graphs, which Corr guidance may possible miss.

All-Links Causality Perturbation (Figure 1 (b)): CF guidance largely preserved the relative rank order of the causal links based on their influence, however the guidance values changed significantly to reflect the altered strengths. Corr guidance, however, exhibited greater changes with a larger variance in guidance values, even for low percentage perturbations, and was more prone to altering the rank order of variable pairs. These findings imply that CF guidance better adapts to the new underlying causal strengths while maintaining the guidance ranking.

5. Conclusion

Our results indicate that CF guidance can capture more nuanced data characteristics than Corr guidance when faced with both narrow data perturbations and changes in the groundtruth causal graphs. This suggests that in scenarios where understanding causal strengths or detecting changes in specific variables is important, CF guidance offers advantages over traditional Corr-based approaches. However, the analysis is not comprehensive; future work investigate more generalized insights.

Acknowledgments

This research is made possible in part by NSF Award #2211845.

References

[BHT20] BATT S., HARMON O. R., TOMOLONIS P.: Learning tableau: A data visualization tool. *The Journal of Economic Education 51*, 3-4 (2020), 317–328. doi:10.2139/ssrn.3438993. 1

[CGM*16] CENEDA D., GSCHWANDTNER T., MAY T., MIKSCH S., SCHULZ H.-J., STREIT M., TOMINSKI C.: Characterizing guidance in visual analytics. *IEEE Transactions on Visualization and Computer Graphics 23*, 1 (2016), 111–120. 1

[RY98] RISTAD E. S., YIANILOS P. N.: Learning string-edit distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence 20*, 5 (1998), 522–532. 2

[WBG24] WANG A. Z., BORLAND D., GOTZ D.: A framework to improve causal inferences from visualizations using counterfactual operators. *Information Visualization* (2024). doi:10.1177/14738716241265120. 2

[WBG25] WANG A. Z., BORLAND D., GOTZ D.: Beyond correlation: Incorporating counterfactual guidance to better support exploratory visual analysis. *IEEE Transactions on Visualization and Computer Graphics* (2025). 1, 2

[WWBG24] WANG Z., WANG A. Z., BORLAND D., GOTZ D.: Causalsynth: An interactive web application for synthetic dataset generation and visualization with user-defined causal relationships. In *IEEE VIS Posters*. 2024. 2