

VAINE: Visualization and AI for Natural Experiments

Grace Guo*

Georgia Institute of Technology

Emily Saldanha§

Pacific Northwest National Laboratories

Maria Glenski†

Pacific Northwest National Laboratories

Alex Endert¶

Georgia Institute of Technology

Dustin Arendt**

Pacific Northwest National Laboratories

ZhuanYi Shaw‡

Pacific Northwest National Laboratories

Svitlana Volkova¶

Pacific Northwest National Laboratories

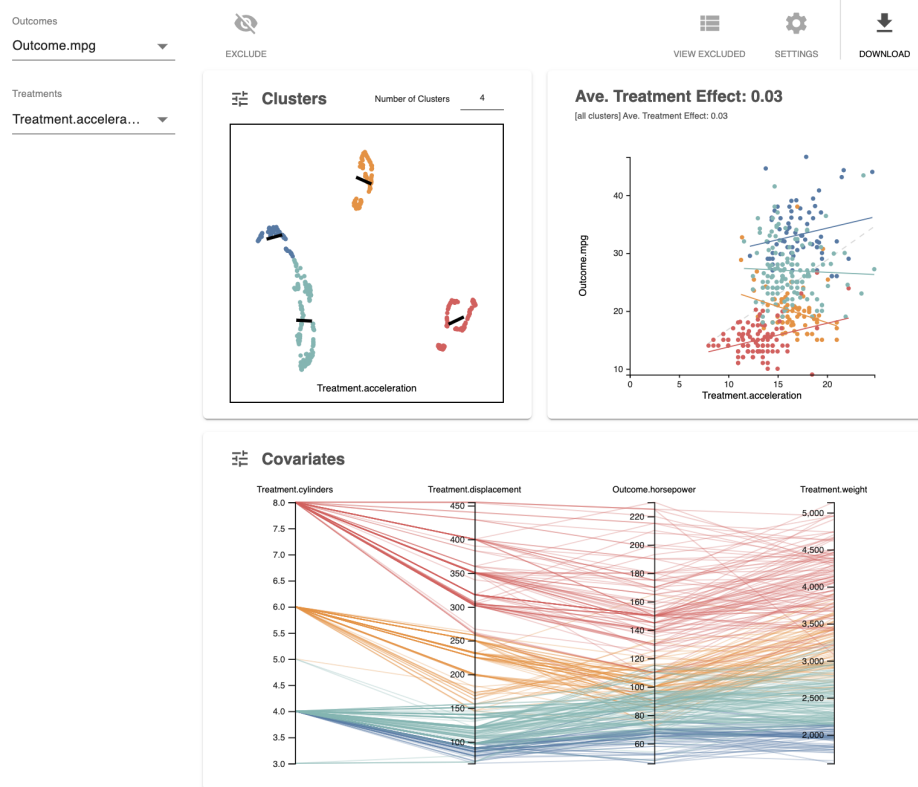


Figure 1: The VAINE (Visualization and AI for Natural Experiments) system interface. The three main views — the *Clusters* view, the *Average Treatment Effect* view, and the *Covariates* view — enable analysts to estimate the effect size of the treatment on the outcome variable of interest. Interactive menu options can be used to adjust clusters, exclude outliers, and inspect covariates.

ABSTRACT

Natural experiments are observational studies where the assignment of treatment conditions to different populations occurs by chance “in the wild”. Researchers from fields such as economics, health-care, and the social sciences leverage natural experiments to conduct hypothesis testing and causal effect estimation for treatment and outcome variables that would otherwise be costly, infeasible, or

unethical. In this paper, we introduce VAINE (Visualization and AI for Natural Experiments), a visual analytics tool for identifying and understanding natural experiments from observational data. We then demonstrate how VAINE can be used to validate causal relationships, estimate average treatment effects, and identify statistical phenomena such as Simpson’s paradox through two usage scenarios.

Index Terms: Human-centered computing—Visualization—Visualization systems and tools ; Human-centered computing—Visualization—Visualization application domains—Visual analytics

1 INTRODUCTION

Identifying and quantifying causal relationships are essential tasks in many domains to develop theories, build models, and guide policy. Accurate causal inference depends on running controlled experiments to investigate the variables in question. This requires the random allocation of treatment conditions via randomized controlled trials (RCTs). RCTs allow scientists to account for confounding factors, and are considered the “gold standard” for supporting causal

*e-mail: gguo31@gatech.edu

†e-mail: maria.glenski@pnnl.gov

‡e-mail: yi.shaw@pnnl.gov

§e-mail: emily.saldanha@pnnl.gov

¶e-mail: endert@gatech.edu

¶e-mail: svitlana.volkova@pnnl.gov

**e-mail: dustin.arendt@pnnl.gov

claims. However, RCTs may not be possible, especially in situations where manipulating treatment variables is impractical or unethical.

For example, during the COVID-19 pandemic, understanding causal relationships between non-pharmaceutical interventions and infection transmission would have helped policy makers decide which interventions to implement and when. However, it is both unethical and impractical to conduct RCTs by dictating different health guidelines to arbitrary subsets of populations. Researchers must thus rely on observational data available post-hoc to look for indirect clues about causality. Beyond global pandemics, similar approaches are needed to understand the causal relationships that explain human behavior, information operations (what *causes* successful disinformation campaigns to go viral online?), and model evaluations (what *causes* a model to fail?).

Causal claims can be tested using observational data if a *natural experiment* has occurred. This is an uncontrolled process where the assignment of treatment conditions to different groups has occurred “as if at random”. Such natural experiments have long been used by researchers in fields such as economics, healthcare, and the social sciences. By analyzing observational data, researchers are able to conduct hypothesis testing and causal effect estimation that would otherwise be costly, infeasible, or unethical [7].

There are many challenges and caveats to working with observational data. In particular, the well known adage that *correlation doesn't imply causation* duly applies. Correlations between treatment and outcome variables may be driven by underlying confounding variables that aren't observed. Making a claim of causation thus requires showing that correlation still exists when controlling for all other variables, i.e., *covariates*. As such, while algorithmically generated claims of causation can be used to explore potential causalities in observational data, they shouldn't be taken at face value and need to be interpreted and contextualized by a domain expert. When working with natural experiment data, experts would want to know whether all possible covariates have been accounted for during data collection, how much outliers influence the result, and what characterizes subgroups experiencing stronger effects. The need to contextualize, explore, and interpret the results of natural experiments thus motivates the development of a visual analytics tool to better support human-in-the-loop causal analysis.

In this paper, we introduce VAINE (Visualization and AI for Natural Experiments) (Fig. 1), a visual analytics system that enables users to identify natural experiments in observational data and estimate the impact of treatment variables on outcomes. We also demonstrate how VAINE can be used to validate causal relationships, estimate average treatment effects, and identify statistical phenomena such as Simpson's paradox through two usage scenarios.

2 RELATED WORK

Causal inference from observational data has long been a challenge in statistics and data analysis. Much work has been done in the field of mathematical causal discovery [14, 17] and causal structure analysis [16]. More recently, a number of visualization techniques have been developed to support causality representation [5, 9] and inference [19, 20]. However, causality inference from observational data remains a challenging task, particularly because *correlation does not imply causation*. For instance, a study by Shalizi et al. on social network data has found that the use of regression coefficients, even taking into consideration temporal data, cannot distinguish causal factors from confounding variables [18].

Despite the difficulties of causal inference, once a causal relationship has been identified, it becomes possible to conduct hypothesis testing by identifying natural experiments when they occur in observational data. Natural experiments have long been used by researchers in fields such as economics [2, 13], healthcare [10] and the social sciences [3]. Observational data can thus be successfully used to identify natural experiments when they occur, allowing an-

alysts to estimate the effect size of the treatment variable on the outcome variable of interest. However, natural experiments remain an underexplored area for visual analytics systems. In this paper, we aim to address this gap through the VAINE system.

3 VAINE

During the design and implementation of VAINE, the first author was embedded within a research team of domain experts applying natural experiments to observational data. In addition to causal effect estimation, the results had to be interpreted and contextualized. Based on these requirements and prior related work, the following design goals (DG) were derived collaboratively:

- DG1: Control for confounding variables.** VAINE should control for confounding variables by automatically grouping data items such that items in the same group have similar values for all covariates other than the treatment and outcome variables.
- DG2: Estimate effect sizes.** The primary outcome of VAINE should be an estimate of the effect of the treatment on the outcome variables. Users should be able to obtain this estimate, and understand how it is derived from the data clusters.
- DG3: Incorporate user feedback.** Users should be able to incorporate domain expertise or contextual information to guide the algorithm. For example, by adjusting clustering parameters, removing outliers, and inspecting cluster details.
- DG4: Support responsive UI.** Interactions should automatically trigger rapid recomputation of analytics and update visualizations.

VAINE¹ is a domain-agnostic tool implemented as a python package and widget for the popular Jupyter notebook² computational environment. Our target users are research analysts who are familiar with using such computational environments for causal analysis. The front-end is built in JavaScript, using the React³ framework.

3.1 Interface

The VAINE system displays the current treatment and outcome variables in the far left panel. The interface has three main views: *Clusters*, *Average Treatment Effect*, and *Covariates* views (Fig. 1). The same cluster color coding is used to coordinate across views. Clusters with an insignificant ($p > 0.05$) regression coefficient are deselected by default (grayed out) across all views.

The *Clusters* view displays all data instances clustered based on their covariates, with a default of 10 clusters. Data instances in the same cluster have similar covariate values, and only vary by treatment and outcome values. This controls for confounding variables, and isolates the effect of the treatment on the outcome variable (**DG1**). Each cluster also has an overlaid line indicating the linear regression of the outcome variable on the treatment variable for all data instances in that cluster (**DG2**). Note that this regression line is based on the data values of the treatment and outcome variables, not the dimensionality reduced coordinates of data points.

The *Average Treatment Effect* view displays data points in a linear model plot based on their treatment and outcome values. Each cluster is overlaid with a line of the linear regression of the outcome variable on the treatment variable. This view includes the weighted average treatment effect for selected clusters, an estimate of the impact of the treatment on the outcome variable (**DG2**). The overall regression line for the entire dataset (without clustering) is depicted with a faint dashed line, providing context to interpret subgroup regressions relative to that of the entire dataset.

¹<https://github.com/pnnl/vaine-widget>

²<https://jupyter.org>

³<https://reactjs.org>

The *Covariates* view summarizes covariate values of all data instances in a parallel coordinates plot, with only the first five covariates shown by default. Visual clustering is used to enhance patterns in the data and reduce visual clutter [21]. In this view, analysts can explore how the clustering controls for covariate values (**DG1**).

3.2 Interactions

Treatment and outcome variables of interest can be selected using the drop-down menus on the left. VAINE then automatically treats all other variables in the dataset as covariates, and re-clusters data instances based on the new covariates (**DG4**).

In the *Clusters* view, analysts can manually adjust selected clusters using a drop-down menu, or by clicking on their regression line overlay (**DG3**). In order to coordinate across views, cluster selections propagate to all three views. Additionally, the *Clusters* view also allows analysts to adjust the number of clusters created (**DG3**). When this value is changed, VAINE automatically re-clusters all data instances and updates all three views. Clicking on a cluster name in the drop-down menu opens a dialogue box where users can change cluster name and color, and inspect the distribution of its covariates (**DG1**). In the *Covariates* view, analysts can manually select which covariates to inspect using the drop-down menu (**DG1**). This allows users to prioritize covariates based on their relative importance.

Brushing and linking is used to coordinate analysis across all views. Users can brush over any of the views to select a subset of data points. Selected points, such as outliers, can be excluded from analysis using the *Exclude* button.

3.3 Algorithm

During preprocessing, the system computes a 2-D embedding using UMAP [15] for each treatment with the covariates and remaining treatments as input features (**DG1**). The tool then performs hierarchical clustering of each 2-D embedding and outputs a dendrogram that is passed to the system front-end using the approach described by Arendt et al. [4]. On load, the number of clusters defaults to 10. When the number of clusters is adjusted, the dendrogram from the preprocessing phase is used to re-cluster the n data points in $O(n)$ time, allowing VAINE to quickly respond to interactions (**DG4**).

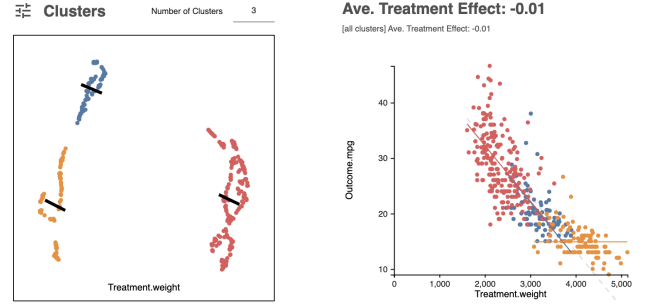
The average treatment effect (*ATE*) value estimates the size of the impact of the treatment on the outcome variable, and allows analysts to perform causal effect estimation. The *ATE* is calculated using the weighted average of the regression coefficients of all selected clusters (**DG2**). In order to do so, linear regression of the outcome variable on the treatment variable is first performed for all clusters. Linear regression was chosen as the default algorithm because it is simple and can be quickly calculated on the front end, allowing for rapid user adjustments (**DG4**). For a dataset with M selected clusters and a sum of N data instances across these clusters, each cluster has a size of n_i (such that $n_1 + n_2 + \dots + n_M = N$) and a regression coefficient of b_i . The calculated *ATE* of this dataset is:

$$ATE = \frac{1}{N} \sum_{i=1}^M n_i \cdot b_i \quad (1)$$

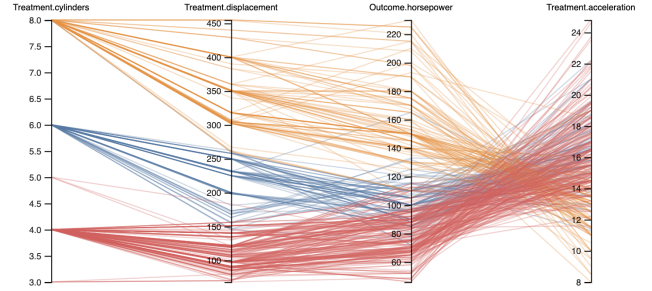
When analysts exclude data points from the dataset (described in 3.2), the clusters will not change. However, regression equations will be recalculated for affected clusters, and the *ATE* will be updated using the new regression coefficients and cluster sizes (**DG4**).

4 USAGE SCENARIOS

We demonstrate VAINE with two usage scenarios [12]. The first uses the Auto MPG dataset [8], often used to predict the miles per gallon of cars from variables such as weight. The second scenario looks at the Ames Housing dataset [6], which has been used to predict house sale prices from other variables. These datasets were selected because they are intuitive and contain well-established causal relationships, which allow us to validate VAINE.



(a) From the *Clusters* view, we identify 3 distinct clumps in the dataset. After setting the number of clusters to 3, the *ATE* of *weight* on *mpg* is estimated to be -0.01.



(b) From inspecting the *Covariates* view, we can see that the three clusters are clearly distinguished by their covariates. In particular, instances within the same cluster share similar values for the number of *cylinders*, *displacement* and *horsepower*.

Figure 2: After setting the number of clusters for a treatment and outcome pair, the clustering can be verified in the *Covariates* view.

4.1 Usage Scenario 1: Auto MPG Dataset

The Auto MPG dataset [8] contains 9 attributes. We exclude temporal and ordinal variables, and consider *mpg* and *horsepower* to be the outcome variables. All other variables are treatment variables.

It has often been claimed that the weight of a vehicle affects its fuel economy [1] (i.e., its miles per gallon (*mpg*) value). We can look at this proposed causal relationship in VAINE by selecting *weight* as the treatment variable of interest. In the *Clusters* view, we see three distinct clumps in the dataset. Setting the number of clusters to 3 identifies these (Fig. 2a). The *Covariates* view reveals that the clusters are clearly distinguished by their covariates. In particular, instances within the same cluster share similar values for the number of *cylinders*, *displacement* and *horsepower* (Fig. 2b). From the *Average Treatment Effect* view, we can see a small but significant estimated effect of *weight* on *mpg*. VAINE thus allows us to confirm from the data that when other variables are held nearly constant, increasing a car's weight decreases its *mpg*, a finding that agrees with both EPA guidelines and prior studies such as [19].

For comparison, we can also look at variables not known to affect *mpg*, such as *acceleration*. A simple linear regression of *acceleration* on *mpg* for the entire dataset yields a regression coefficient of 1.20. We can verify this relationship in VAINE by selecting *acceleration* as the treatment variable of interest. The *Clusters* view reveals multiple distinct clusters in the dataset; we thus adjust the number of clusters to 4 (Fig. 1). With this new clustering, the *ATE* is now 0.03. While still significant, this is a large decrease from the coefficient of 1.20 when only a simple linear regression was used without clustering. Additionally, the direction of the effect of *acceleration* on *mpg* is not constant across clusters. The overall regression line for the entire dataset suggests a strongly positive relationship of *acceleration* on *mpg* (Fig. 1). However, from examining the subgroups, we see that the blue and orange clusters show a relationship opposite to the


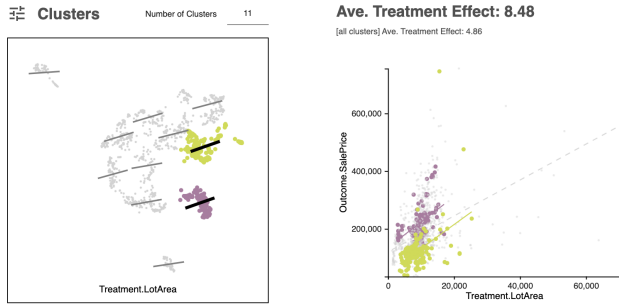



Figure 3: Outliers identified in the *Average Treatment Effect* view can be excluded from analysis, updating the axes and ATE value.



(a) The purple and lime clusters have the steepest regression lines. Selecting these two clusters increases the ATE to 8.48 (from an ATE of 4.86 when all clusters are selected).



(b) Unlike the rest of the dataset, data instances in the selected clusters (purple and lime) tend to have a *2ndFlrSF* of 0. For clarity, only a subset of covariates are shown.

Figure 4: By inspecting the purple and lime clusters in greater detail, we can conclude that *Lot Area* has a stronger effect on *Sale Price* for smaller properties without a second floor.

overall trend. This is a case of Simpson’s paradox, where trends in subgroups disappear or are reversed when groups are combined. By visualizing clusters in the dataset, VAINÉ thus supports the identification of inconsistent treatment effects in subgroups, allowing instances of Simpson’s paradox to be mitigated by expert judgement.

4.2 Usage Scenario 2: Ames Housing Dataset

The Ames Housing dataset [6] contains 81 attributes. We exclude all temporal and ordinal attributes, and consider *Sale Price* to be the outcome attribute of interest. All other attributes are treatments.

Of the available treatment variables, we would expect property area to affect property sale price. In order to estimate the change in sale price per unit change in property area, we set *Lot Area* as the treatment variable. From the *Average Treatment Effect* view, a clear outlier is identified. Excluding this outlier increases the estimated ATE of *Lot Area* on *Sale Price* to 4.86 (Fig. 3).

In the *Clusters* view, some clusters have a steeper regression line than others, suggesting that *Lot Area* has a stronger effect on *Sale Price* for those particular properties. We can investigate this further by deselecting all other clusters, highlighting only the purple and

lime clusters with the steepest regression lines. From this selection, we see that the ATE for these two clusters is 8.48, indicating that the effect of *Lot Area* on *Sale Price* for properties in these two clusters is almost double that of the entire dataset (Fig. 4a). From the *Average Treatment Effect* view, we see that the two clusters tend towards smaller *Lot Area* values. The steeper effect of *Lot Area* on *Sale Price* for smaller properties corresponds with findings from prior studies using the Ames housing dataset [11].

Using the *Covariates* view, we can further inspect the data instances in the two clusters. From Fig. 4b, it can be seen that the most distinct covariate differentiating these clusters from the rest of the data is the Second floor square feet (*2ndFlrSF*) variable. Instances in the two clusters tend to have a *2ndFlrSF* of 0. We thus conclude that per unit increases in *Lot Area* cause a greater increase in *Sale Price* for smaller properties without a second floor.

5 LIMITATIONS AND FUTURE WORK

VAINÉ was originally designed to operate on continuous treatments, outcomes, and covariates. Therefore, some decisions are less appropriate for ordinal or categorical data, such as using dimension reduction and linear regression. However, workarounds such as one-hot encodings are possible. We also note that VAINÉ assumes linearity when estimating the ATE. While the *Average Treatment Effect* view can be used to identify instances of non-linear relationships, VAINÉ does not currently take such factors into consideration. Moving forward, we would like to extend VAINÉ to account for non-linearity. Given the utility of time ordering to establish causal relationships, future work can also extend VAINÉ to account for a range of data types and experimental setups.

Finally, the design goals outlined in this paper were identified by the research team of domain experts conducting causal analysis. There may be further design goals not considered here. Preliminary applications of VAINÉ by the research team have also found limitations in terms of the scalability of the system to datasets with tens of thousands of instances. While computation time remains feasible, optimizing for scale represents a potential avenue of future work.

6 CONCLUSION

In this paper we presented VAINÉ, a visual analytics tool designed for research analysts familiar with using computational environments for causal analysis. VAINÉ helps users identify natural experiments in observational data by clustering data instances based on covariates. Analysts can thus control for confounding variables and isolate the effect of the treatment(s) on outcome variables. The visual representations in VAINÉ allow domain experts to explore potential causalities and contextualize results by adjusting clusters, identifying outliers, and inspecting covariates. We have also demonstrated how VAINÉ can validate causal relationships, estimate average treatment effects, and identify cases of Simpson’s paradox.

Despite the limitations addressed above, VAINÉ is a practical and novel technique for leveraging natural experiments in observational data for human-in-the-loop causal effect estimation. Having validated VAINÉ, the research team is currently applying it to understand various operational domains including information operations, human behavior, and causal explanations of model performance.

ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for their thoughtful and detailed feedback on this paper, and members of the Georgia Tech Visualization Lab for their helpful insights and comments.

This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions and/or findings expressed are those of the author and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Further development and evaluation of this work is supported by NSF IIS-1750474.

REFERENCES

- [1] Highlights of the automotive trends report, Jan 2021.
- [2] J. D. Angrist. Lifetime earnings and the vietnam era draft lottery: evidence from social security administrative records. *The American Economic Review*, pp. 313–336, 1990.
- [3] J. D. Angrist and V. Lavy. Using maimonides’ rule to estimate the effect of class size on scholastic achievement. *The Quarterly journal of economics*, 114(2):533–575, 1999.
- [4] D. L. Arendt, N. Nur, Z. Huang, G. Fair, and W. Dou. Parallel embeddings: a visualization technique for contrasting learned representations. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*, pp. 259–274, 2020.
- [5] J. Bae, T. Helldin, and M. Riveiro. Understanding indirect causal relationships in node-link graphs. In *Computer graphics forum*, vol. 36, pp. 411–421. Wiley Online Library, 2017.
- [6] D. De Cock. Ames, iowa: Alternative to the boston housing data as an end of semester regression project. *Journal of Statistics Education*, 19(3), 2011.
- [7] J. DiNardo. Natural experiments and quasi-natural experiments. In *Microeconometrics*, pp. 139–153. Springer, 2010.
- [8] D. Dua and C. Graff. UCI machine learning repository, 2017.
- [9] N. Elmqvist and P. Tsigas. Causality visualization using animated growing polygons. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No. 03TH8714)*, pp. 189–196. IEEE, 2003.
- [10] A. Finkelstein, S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group. The oregon health insurance experiment: evidence from the first year. *The Quarterly journal of economics*, 127(3):1057–1106, 2012.
- [11] F. Hohman, A. Head, R. Caruana, R. DeLine, and S. M. Drucker. Gamut: A design probe to understand how data scientists understand machine learning models. In *Proceedings of the 2019 CHI conference on human factors in computing systems*, pp. 1–13, 2019.
- [12] T. Isenberg, P. Isenberg, J. Chen, M. Sedlmair, and T. Möller. A systematic review on the practice of evaluating visualization. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2818–2827, 2013.
- [13] R. Jensen. The digital provide: Information (technology), market performance, and welfare in the south indian fisheries sector. *The quarterly journal of economics*, 122(3):879–924, 2007.
- [14] S. Mani and G. F. Cooper. Causal discovery using a bayesian local causal discovery algorithm. In *Medinfo*, pp. 731–735, 2004.
- [15] L. McInnes, J. Healy, and J. Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.
- [16] S. Nadkarni and P. P. Shenoy. A bayesian network approach to making inferences in causal maps. *European Journal of Operational Research*, 128(3):479–498, 2001.
- [17] J. Pearl. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688, 1995.
- [18] C. R. Shalizi and A. C. Thomas. Homophily and contagion are generically confounded in observational social network studies. *Sociological methods & research*, 40(2):211–239, 2011.
- [19] 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.
- [20] 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*, 2020.
- [21] H. Zhou, X. Yuan, H. Qu, W. Cui, and B. Chen. Visual clustering in parallel coordinates. In *Computer Graphics Forum*, vol. 27, pp. 1047–1054. Wiley Online Library, 2008.