

# CausalPrism: A Visual Analytics Approach for Subgroup-based Causal Heterogeneity Exploration


Jiehui Zhou , Xumeng Wang, Wong Kam-Kwai, Wei Zhang, Xingyu Liu, Juntian Zhang, Minfeng Zhu and Wei Chen



Fig. 1: *CausalPrism* helps identify, explore, rank, and interpret causal subgroups in observational data. (A) The Causal Subgroup View includes a tailored tabular visualization of subgroup descriptions, a subgroup editing window, and a ranking visualization of multiple evaluation metrics to support subgroup overview, modification, and ranking. (B) The Covariate Projection View reduces units with high-dimensional covariates to two dimensions, allowing users to analyze similarities between subgroups and assisting in merging and splitting subgroups. (C) The Treatment Effect Validation View consists of propensity score histograms, individual treatment effect dot plots, and detailed information of matched units, which helps users interpret effect strength and uncertainty, thereby increasing trust.

**Abstract**—In causal inference, estimating Heterogeneous Treatment Effects (HTEs) from observational data is critical for understanding how different subgroups respond to treatments, with broad applications such as precision medicine and targeted advertising. However, existing work on HTE, subgroup discovery, and causal visualization is insufficient to address two challenges: first, the sheer number of potential subgroups and the necessity to balance multiple objectives (e.g., high effects and low variances) pose a considerable analytical challenge. Second, effective subgroup analysis has to follow the analysis goal specified by users and provide causal results with verification. To this end, we propose a visual analytics approach for subgroup-based causal heterogeneity exploration. Specifically, we first formulate causal subgroup discovery as a constrained multi-objective optimization problem and adopt a heuristic genetic algorithm to learn the Pareto front of optimal subgroups described by interpretable rules. Combining with this model, we develop a prototype system, *CausalPrism*, that incorporates tabular visualization, multi-attribute rankings, and uncertainty plots to support users in interactively exploring and sorting subgroups and explaining treatment effects. Quantitative experiments validate that the proposed model can efficiently mine causal subgroups that outperform state-of-the-art HTE and subgroup discovery methods, and case studies and expert interviews demonstrate the effectiveness and usability of the system. Code is available at [OSF](https://github.com/OSF).

**Index Terms**—Causal inference, data heterogeneity, subgroup discovery, optimization, interpretability, visual analytics

## 1 INTRODUCTION

Causal inference is a data analysis process aiming at conclusions about whether and to what extent treatments affect outcomes [48]. Data heterogeneity must be considered when estimating treatment effects, as the effect of the same treatment may vary across subgroups. As shown in Fig. 2, subgroups within the population responded differently to the treatment. The treatment exerts larger effects on Subgroup 1 and Subgroup 3 than Subgroup 2. Nevertheless, the high variance of Subgroup

- J. Zhou, W. Zhang, X. Liu, J. Zhang, M. Zhu and W. Chen are with the State Key Lab of CAD&CG, Zhejiang University. E-mail: {zhoujehui, 22151190, liu\_xingyu, 3200105799, minfeng\_zhu, chenwis}@zju.edu.cn.
- X. Wang is with TMCC, CS, Nankai University. E-mail: wangxumeng@nankai.edu.cn.
- W. Kam-Kwai is with Hong Kong University of Science and Technology and Georgia Institute of Technology. E-mail: kkwongar@connect.ust.hk.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: [reprints@ieee.org](mailto:reprints@ieee.org).

Digital Object Identifier: [xx.xxxx/TVCG.201x.xxxxxx](https://doi.org/10.1109/TVCG.201x.xxxxxx)




Fig. 2: An illustrative toy example. There is only one covariate, and the change in the outcome between the treatment and control group can be informally thought of as the treatment effect. Subgroup 3 has a high effect and low variance, which is better than Subgroup 1 and 2.

1 indicates that individual differences (uncertainty of the outcome) exist. Discovering those subgroups with strong treatment effect and low outcome variance (hereinafter referred to as significant treatment effect) compared to the overall population is widely used in domains such as healthcare [50], marketing [56], and public administration [23]. For example, marketers want to find customer groups where advertising more effectively drives purchases. Since Randomized Controlled Trials (RCTs), known as the gold standard for causal inference [31], are not always feasible due to cost or ethical concerns, there is a strong need to uncover those causal subgroups from observational data effectively.

In practical applications, analyzing causal heterogeneity faces two challenges. First, it is nontrivial and challenging to identify important subgroups from a large number of subgroup candidates. Subgroups can be described by different combinations of variables, which could lead to a combinatorial explosion of candidates. Selecting optimal subgroups requires trade-offs among various targets in the objective space, such as effect strength, outcome variance, and subgroup coverage, further complicating the subgroup discovery process. Coordinating the above factors has to follow users’ analysis requirements. Nevertheless, the identification model can hardly communicate with humans due to the lack of interpretability, which causes the second challenge. The subgroups obtained by the black-box model may hardly be interpreted by or fail to support analysis tasks. Users still need to tediously analyze and compare multiple subgroups to determine which one they prefer. In addition, without a treatment effect explanation, numerical causal conclusions alone are difficult to convince users, especially in high-stakes safety and life-critical fields.

Existing work is insufficient to address these challenges. Automated heterogeneous treatment effect (HTE) estimation methods, such as causal trees [2] and causal forests [57], mainly construct hierarchical structures for individuals in datasets and identify leaf nodes in the hierarchies as subgroups. However, not all leaf nodes can reflect a significant treatment effect, leading to less useful results. Therefore, users still need to go through tedious review and analysis to find subgroups that meet their requirements. Subgroup discovery methods [6, 22, 55] can directly optimize correlation objectives but lack attention to the more complex causal effects that require statistical inference and confounding bias correction. Some researchers use visualization to assist causal analysis, but they either focus on the representation of causal graph structure [37, 58, 59, 66], manual selection of variables to divide subgroups [29], or clustering to obtain subgroups whose meaning is difficult to describe [28]. Therefore, how to support the cooperation between human intelligence and computing power in the analysis of HTE is still underexplored.

In this work, we propose a visual analytics approach for subgroup-based causal heterogeneity exploration that supports users in effectively identifying optimal subgroups from observational data, comparing and

ranking different subgroups, and verifying treatment effects. First, we propose a causal subgroup discovery model based on constrained multi-objective optimization (MOO). Subgroups are described by interpretable rules, where the rule antecedents are conjunctions containing covariates and corresponding values, and the consequents correspond to subgroup evaluation metrics, such as effects and variances. Coverage and antecedent length are used as constraints to ensure interpretability. Due to multiple objectives and constraints, the optimal Pareto front of the subgroup is learned by a heuristic searching algorithm for user analysis. Second, we design and develop an interactive prototype system, *CausalPrism*, which incorporates intuitive visualizations of subgroups, evaluation metrics, and explanations of treatment effects, thereby facilitating users’ understanding, comparison, and verification of causal subgroups. Quantitative experiments, case studies, and expert interviews demonstrate the effectiveness and usability of the proposed model and system. In summary, our contributions are as follows:

- We propose a subgroup discovery model based on constrained multi-objective optimization, which can mine rule-explained subgroups with significant treatment effects from a large amount of high-dimensional observational data that outperforms state-of-the-art methods.
- We designed and implemented an interactive visual analytics prototype system, *CausalPrism*, which includes table-based subgroup visualization, multi-attribute ranking, and matching unit-based explanation of treatment effects. The system supports users in understanding, comparing, and validating causal subgroups. Its utility has been proved through case studies, and positive feedback has been received during expert interviews.

## 2 RELATED WORK

### 2.1 Heterogeneous Treatment Effect Estimation

Treatment effects vary across the whole population. Conditional average treatment effect (CATE), individual average treatment effect (ITE) and causal rules comprise current HTE research. Reviews [30, 68] offer in-depth analyses of causal inference.

*CATE* evaluates treatment effects on specific subgroups of the population, given similar covariates like demographics. To optimize the heterogeneity of treatment effects, tree-based methods [2, 3, 57] are commonly employed to partition the covariate space into subspaces. For instance, Causal Tree [2] constructs the tree and estimates treatment effects in each subspace using separate data, avoiding overfitting by cross-validation. Wager *et al.* [57] suggested Causal Forest, combining causal tree ensemble results for more robust and smooth estimation. Root-to-leaf node routes naturally define subgroups of heterogeneous CATEs, making the tree model interpretable. However, tree-based methods may have limited performance due to greedy tree building process and does not necessarily return the “optimal” structure.

*ITE* compares outcomes with and without treatment. Since only one outcome is visible, another must be estimated. Existing techniques are single- or multi-model-based depending on whether treatment and control groups are estimated independently. The former uses regression to fit treatment effects. For example, Hill *et al.* [33] employs Bayesian additive regression trees to fit the outcome surface. The latter fits the treated and control groups separately, achieving better performance for significant differences between groups’ outcomes. The base model uses off-the-shelf estimators like linear regression [12] or neural networks [38]. With well-tuned parameters, these models can accurately estimate effects but are uninterpretable.

Several researchers have tried to find *causal rules* in data. For example, CRE [9] is a two-stage method that first produces rules using methods such as random forest or Gradient Boosting Machines, then picks robust ones using stability selection regularization. After mining association rules from data, Li *et al.* [46] conducted a cohort study to test whether the association rules were causal.

Many methods have presented estimators that can accurately estimate HTE despite confounding biases. We leverage existing estimators but focus on finding subgroups with significant treatment effects among many candidates. We present a constrained MOO-based causal

subgroup discovery model that outperforms tree-based and black-box approaches in significance and interpretability.

## 2.2 Visual Causality Analysis

Automated causal detection algorithms are built upon assumptions and complex causal mechanisms that are hardly fulfilled in real life, causing accuracy and interpretability issues. Visualizations have been used to explore and verify causality interactively, which can be divided into homogeneous and heterogeneous causality investigations.

*Homogeneous* causality assumes that data's causal mechanism is static and stable. Visualization helps users grasp complex causal relationships and make decisions. Graph-based visualizations [7, 35, 58, 66] have been widely used to demonstrate causality in multi-attribute datasets, emphasizing the use of statistics to locate and manipulate improper relationships for what-if analysis. They employ advanced layout designs to highlight attribute distributions and enhance graph readability. In addition, design factors like node size [39], edge shape [8], and crowd beliefs [70] affect users' understanding of causal relationships. Bar charts and scatter plots are also used to infer causality [40, 67, 69]. However, these approaches lack generalizability due to population variances such as demographics and environmental factors.

*Heterogeneous* causality examines causal relationships or effects that vary over time or data subgroups. Most work examines causal structural heterogeneity. In Causal Structure Investigator [59], users can acquire data subdivisions through manual filtering and k-means clustering. Then, these subdivisions are mapped to causal graphs for detailed analysis of causal paths. Jin *et al.* [37] focus on subsets in event sequences. Overlapping adjacency matrices with inner and outer sections lets users easily identify the differences in causal relationships between subsets. Deng *et al.* [19] created causal graph bands with compass glyphs for spatio-temporal sequences to show dynamic causal relationships in period-based time windows. This helps users understand influence transmission and identify spurious causalities. DOMINO [60] applies time delays and event constraints to temporal causality analysis, facilitating hypotheses formulation and validation.

Other research examines causal inference heterogeneity. The Absolute Standardized Mean Difference (ASMD) plot is used to assess covariate balance in groups after weighting and propensity score matching [25, 52]. Guo *et al.* [28] created VAINE to enable users to find statistical phenomena like Simpson's paradox by manually selecting clusters in covariate projections and observing their impact. Causalvis [29], a later proposal, enables the visualization of a whole causal analysis workflow. The raincloud and beeswarm plots in the Treatment Effect Explorer module let users manually pick subgroups faceted by covariates and analyze ITE distribution to examine heterogeneity.

However, the present HTE visualization work involves time-consuming manual participation to locate subgroups, and the subgroups obtained through clustering lack explicit interpretable descriptions. Therefore, we propose *CausalPrism* to automatically obtain rule-described subgroups with significant treatment effects through optimization and design visualizations for subgroup exploration, comparison, and treatment effect validation.

## 2.3 Subgroup Discovery and Visualization

Subgroup discovery (SD) is a descriptive data mining method that finds data subgroups with intriguing patterns on certain goals, as summarized in comprehensive surveys [4, 32]. Data subgroups can be represented using description languages like attribute-value pairs and logical forms (*e.g.*, conjunctions, inequalities, and fuzzy logic). Subgroup interestingness can be measured using binary, nominal, or numerical targets. Post-processing methods have been applied to select diverse and less redundant subgroups. Search methodologies like exhaustive and heuristic search have been used due to the large number of candidate subgroups.

Using the exhaustive techniques [6, 26, 27, 62], all possible subgroups are searched. Since viable subgroups are exponentially large, a naive exhaustive search is time-consuming. Minimum support, optimistic estimate pruning, and generalization-aware pruning can reduce the hypothesis space. SD-Map [6] is an exhaustive SD approach that uses depth-first search to produce candidates, extending the Frequent Pattern

(FP) Growth-based association rule mining method. The SD-Map\* [5] is extended with binary, categorical, and continuous target variables.

Further studies [18, 22, 43, 55, 72] employed efficient heuristic methods. For example, DSSD [55] uses beam search, which starts with an initial solution and subsequently spreads to several candidates. Top performers are kept for the next iteration until a stopping condition is reached. SDIGA [18] is an evolutionary fuzzy rule induction method that facilitates the discovery of general rules by allowing variables to take multiple values. Subgroups can be evaluated in terms of confidence, support, and unusualness.

Visualization techniques have also been proposed in order to support subgroup-level analysis tasks, such as subgroup multi-feature visualization [21, 24], model diagnosis on data subsets [14, 20, 42, 54, 71], and high-dimensional data subspace exploration [47, 64, 65]. For example, Taggle [21] employs a tabular visualization design that allows for hierarchical grouping and sorting of massive amounts of data. The icicle plot [20] and the map-based metaphor [47] provide help for comparisons between subgroups.

However, most SD methods only focus on correlations, involving just covariates and outcomes. It is unsuitable for SD in causal scenarios (treatment, covariates, and outcomes must be considered). To this end, we formulate causal SD as a constrained MOO problem that can be efficiently solved using heuristic search. A range of subgroup visualization techniques, such as multi-attribute ranking [24], are incorporated into the *CausalPrism* system to help users explore and compare subgroups.

## 3 BACKGROUND

### 3.1 Preliminaries

We introduce the basis of causal inference under the potential outcome (PO) framework [51] and give examples based on medical scenarios.

A **unit** is an individual or object under study. A medical study unit may be a patient. The subscript  $i$  denotes the  $i$ -th unit.

A **treatment** is an intervention or exposure that subjects to a unit. A new medicine or therapy could be used as a treatment in a medical study. Let a binary  $T$  indicate whether a unit has received a treatment. Units satisfying  $T = 1$  belong to the treatment group, while those  $T = 0$  belong to the control group.

**Outcomes** are what would happen to units under different treatments. Each unit has two potential outcomes: factual outcome and counterfactual outcome. For instance, patient survival time is an outcome in a medical study. The potential outcome with treatment is  $Y(T = 1)$ , also abbreviated as  $Y(1)$ , and without treatment, it is  $Y(0)$ .

**Covariates** are background variables that affect treatment assignment and outcome. For example, patient demographic information such as age may influence medication use (treatment assignment) and blood pressure (outcome). Observational studies often control for covariates to mitigate confounding and provide more unbiased effects estimates. Covariates are represented as a vector  $\mathbf{X}_i = (x_{i,1}, \dots, x_{i,d})$ , where  $d$  is the number of covariates.

**Observational data** refers to data collected without the researcher manipulating the environment or the subjects being studied. It differs from RCTs, which randomly assign treatment to each unit. The observational data containing  $n$  units is denoted by  $\mathcal{D} = \{(T_i, \mathbf{X}_i, Y_i)\}_{i=1}^n$ .

**Treatment effect** refers to the impact of a treatment on an outcome. It can be obtained by quantitatively comparing the potential outcomes in the treatment and control conditions at different levels, such as populations, subgroups, and units. For unit  $i$ , its individual treatment effect (ITE) is defined as:

$$\tau_i = Y_i(1) - Y_i(0). \quad (1)$$

Unfortunately, for any unit, only one of the two potential outcomes can be observed, so ITE is not identifiable. One way to address this lack of counterfactual outcomes is to estimate the average treatment effect (ATE) on the population, defined as follows:

$$\tau = \mathbb{E}[Y(1) - Y(0)]. \quad (2)$$

ATE may fail to accurately reflect treatment effects due to the heterogeneity of units. This is overcome by conditional average treatment effect (CATE) on subgroups, which is defined as follows:

$$\tau(\mathbf{x}) = \mathbb{E}[Y(1) - Y(0) \mid \mathbf{X} = \mathbf{x}]. \quad (3)$$

**P propensity score** is a balancing score  $e(\mathbf{x}) = P(T = 1 \mid \mathbf{X} = \mathbf{x})$ , defined as the conditional probability of getting a treatment given the covariates. In observational data, a biased treatment effect would be obtained by directly using the difference between the average outcome of the treatment and control groups because the treatment assignment is correlated with covariates. Rubin *et al.* [10] proved that  $\{Y(0), Y(1)\} \perp T \mid e(\mathbf{X})$  under the assumption of unconfoundedness. For binary treatment, the Logistic regression model is commonly used to estimate propensity scores [13].

**Outcome Variance** refers to the outcome variability among the units in the treatment and control group. A lower variance means that the treatment leads to a more consistent outcome among units. Let  $\sigma^2$  be the variance.

**Causal Subgroup** refers to specific subgroups within the population that exhibit significant treatment effects. For example, the preventive effect of influenza vaccine is more significant in the elderly and immunocompromised people.  $\mathcal{S}$  is used to represent subgroups.

### 3.2 Design Requirements

We distilled the design requirements from interviews with three experts (E1-3) and a literature review. Data analysts E1 and E2 have three years of work experience in a technology firm. Their daily tasks include evaluating KPI anomalies and guiding advertising placement using causal analysis on observational data. E3, a university associate professor, has written multiple causal inference studies. They noted that causal inference is plagued by data heterogeneity, and massive amounts of observational data lack appropriate exploration tools. Causality interpretability is also crucial since users cannot make decisions if they don't trust the result. Finally, the requirements are listed.

- R1 Descriptive subgroup identification.** Observational data usually contains a large potential exploration space with many variables. Although traditional data clustering methods can be used to discover clusters, they do not give a corresponding interpretation. Experts mentioned that "Although high-value groups can be manually segmented based on domain knowledge, it often requires multiple attempts of different filtering conditions." Users can take advantage of automatically identified subgroups to further discover interesting causal patterns.
- R2 Subgroup understanding and valuation.** Subgroups involve rich information such as variables used in subgroup descriptions, value distribution of covariates, treatment effect, and variance scores. Users should be able to browse such information to understand the characteristics of a subgroup. It is necessary for the system to provide a clear and intuitive visualization for subgroups.
- R3 Subgroup adjustment and hypothesis.** The subgroups automatically discovered by the model may not satisfy users. Experts say that for advertising scenarios, they aim to boost user spending and meet total profit goals. Therefore, when necessary, they will relax the filtering conditions or merge small groups to enlarge the target subgroup. Our approach needs to support users to adjust subgroups. The basis of adjustment could be the understanding of the target subgroup, domain knowledge from users, or the analysis results of what-if tests on new subgroups.
- R4 Subgroup comparison and ranking.** Users have diverse preferences for causal subgroups. For example, when determining the target audience for advertising, conservative users are willing to choose subgroups that are generally effective and have smaller variances, while risk-takers try subgroups that have stronger effects but may also have greater outcome fluctuations. Therefore, we need to allow users to compare subgroups from multiple perspectives and rank them based on their preferences to select satisfactory subgroups.
- R5 Treatment effect validation.** Users need to understand why certain treatment effect is estimated and be provided with data

evidence to explain their rationality. In addition, the estimated treatment effect may be biased by the size and distribution of the data units. Seeking reliable conclusions, visualizations are needed to help users rule out suspicious causal effects.

## 4 WORKFLOW OVERVIEW

To address the requirements mentioned in Sec. 3.2, we developed *CausalPrism*, a visual analytics system for analyzing causal heterogeneity from a subgroup perspective, allowing users to identify, explore, rank, and validate causal subgroups. Figure 3 illustrates the workflow of *CausalPrism*. The input is observational data, including covariates (e.g., age, height, weight), a treatment (e.g., drug), and an outcome (e.g., survival time). Given the input data, the system works as follows:

(A) First, causal subgroups are automatically identified by the model from the observational data in which the treatment has a significant effect on the outcome. Since there are multiple objectives, including treatment effects and outcome variance, the discovered subgroups are usually an optimal solution set. Constraints such as subgroup coverage can also be imposed on the subgroup discovery process. (R1, Sec. 5.1)

(B) Second, to understand the subgroups, their interpretable descriptions (e.g., "age>60 AND sex=female") can be considered as a description of the older female group) and evaluation metrics (e.g., treatment effects, outcome variance) are presented. Users can get an overview of the discovered subgroups and view the details of a subgroup on demand. (R2, Sec. 6.2)

(C) Third, custom subgroups are allowed, where users can hypothesize subgroups based on domain knowledge for what-if analysis. By weighing evaluation metrics such as effect strength and subgroup size, users can make multi-criteria decisions and select a preferred subgroup. (R3, R4, Sec. 6.2, Sec. 6.3)

(D) Finally, after the subgroups of interest have been selected, explanations of the treatment effect within the subgroups can be examined. Analytical evidence, such as the uncertainty of individual treatment effects, is provided, thus helping users exclude the unreliable subgroup and enhance trust in the results. (R5, Sec. 5.2, Sec. 6.4)

## 5 MODELS

In this section, we introduce the models used in our system.

### 5.1 Causal Subgroup Discovery

We first frame causal subgroup mining as a constrained multi-objective optimization problem and solve the optimal subgroup set through an efficient heuristic algorithm. (R1)

#### 5.1.1 Problem formulation

Without loss of generality, we consider observational data whose covariates are binary, *i.e.*,  $\mathbf{X}_i = (x_{i,1}, \dots, x_{i,d}) \in \{0, 1\}^d$ . Categorical variables can be binarized by one-hot encoding (such as color=red, color=yellow, color=blue). As for numerical variables, we can convert the value intervals to binary by bucketing strategy (such as age≤10, age>10, ..., age≤100, age>100). Formally, our goal is to learn causal subgroups  $\mathcal{S}$  from the given observational data  $\mathcal{D}$ . We use interpretable rules (simple logical structures of the form "IF  $P$  THEN  $Q$ ") to describe the subgroup  $\mathcal{S} : \alpha \Rightarrow [\tau, \sigma^2(0), \sigma^2(1)]$ , which contain the antecedent  $\alpha$  and the consequent  $[\tau, \sigma^2(0), \sigma^2(1)]$ .

A **antecedent**  $\alpha$  is the condition of the subgroup, expressed as the conjunctive normal form (CNF) of a series of atoms  $\bigwedge_{j \in \Gamma} x_j$ , *e.g.*, "age > 25 AND job == teacher".  $\Gamma$  is the covariate indices used in the antecedent, which is a subset of the indices of all binary covariates, *i.e.*,  $\Gamma \in 2^{[d]}$ , where  $[d] = \{1, \dots, d\}$  and  $2^{[d]}$  means the power set of  $[d]$ . The atom  $x_j$  is the smallest interpretable element. The mapping from a  $\mathcal{S}$  to a CNF is given by  $\alpha_{\mathcal{S}}(\mathbf{X}_i) = \bigwedge_{j \in \Gamma_{\mathcal{S}}} x_{i,j}$ . For brevity, we call it as  $\bigwedge_{j \in \mathcal{S}} x_{i,j}$ . When  $\alpha_{\mathcal{S}}(\mathbf{X}_i)$  is true, the  $i$ -th unit is **covered** by the subgroup  $\mathcal{S}$ . We define the length of the antecedent  $|\alpha|$  as the number of different covariates. For example, the length of the antecedent "10 < age <= 25 AND sex = female" is 2.  $|\alpha|$  reflects the readability of the antecedent; shorter antecedents are easier for users to interpret.




Fig. 3: A four-step workflow for subgroup-based causal heterogeneity exploration. (A) The model automatically mines subgroups with significant treatment effects from observational data. (B) Subgroups can be explored through tabular and multi-attribute visualizations. (C) Users can interactively analyze new subgroup hypotheses and achieve multi-criteria decision-making based on their preferences. (D) Effect interpretation based on matched units simulates the user's familiar A/B testing, aiding in result validation.

The **consequent** is the evaluation results of the subgroup, consisting of the treatment effect  $\tau$ , the control group variance  $\sigma^2(0)$ , and the treatment group variance  $\sigma^2(1)$ .

To estimate the CATE  $\tau$  for units covered by the subgroup, we employ Inverse Probability Weighting (IPW) [34], which assigns appropriate weights  $w_i = \frac{T_i}{e_i} + \frac{1-T_i}{1-e_i}$  to each unit to balance the distribution of covariates in the treatment and control groups, thereby simulating RCTs. The normalized weighted average of the factual outcomes for the treatment and control groups can estimate treatment effects [36]:

$$\tau_S = \frac{\sum_{i \in \mathcal{D}_S^+} w_i Y_i}{\sum_{i \in \mathcal{D}_S^+} w_i} - \frac{\sum_{i \in \mathcal{D}_S^-} w_i Y_i}{\sum_{i \in \mathcal{D}_S^-} w_i}, \quad (4)$$

where  $\mathcal{D}_S$  denotes the covered data,  $\mathcal{D}^+$  denotes the data that received the treatment ( $T = 1$ ), and  $\mathcal{D}^-$  denotes the data that did not receive the treatment ( $T = 0$ ),  $\mathcal{D}_S^+ = \{i | i \in \mathcal{D}^+ \wedge \alpha_S(\mathbf{X}_i) = 1\}$  denotes the units in the treatment group that are covered by the subgroup  $S$ ,  $\mathcal{D}_S^- = \{i | i \in \mathcal{D}^- \wedge \alpha_S(\mathbf{X}_i) = 1\}$  denotes the units in the control group that are covered by the subgroup  $S$ .

Since obtaining treatment effects is a statistical estimation problem, it is important to consider the uncertainty of the treatment effect, which can be measured by the outcome variance, defined as:

$$\begin{aligned} \sigma_S^2(0) &= \frac{\sum_{i \in \mathcal{D}_S^-} w_i (Y_i - \bar{Y}_w)^2}{\sum_{i \in \mathcal{D}_S^-} w_i} \\ \sigma_S^2(1) &= \frac{\sum_{i \in \mathcal{D}_S^+} w_i (Y_i - \bar{Y}_w)^2}{\sum_{i \in \mathcal{D}_S^+} w_i}, \end{aligned} \quad (5)$$

where  $\bar{Y}_w$  is the weighted outcome mean.

Therefore, we formalize learning causal subgroups from observational data as a **constrained multi-objective optimization** problem:

$$\begin{aligned} &\max_S \tau_S \\ &\min_S \sigma_S^2(0), \sigma_S^2(1) \\ &\text{s.t. } |\mathcal{D}_S| \geq C \\ &\quad |\alpha_S| \leq L. \end{aligned} \quad (6)$$

To ensure interpretability and meet user personalized needs, two constraints are added.  $|\mathcal{D}_S| \leq C$  limits the unit covered by the subgroup to at least  $C$ , and  $|\alpha_S| \leq L$  limits the length of the antecedent by  $L$ .

### 5.1.2 Solving the problem

Solving the optimization problem of Eq. (6) is not easy because the decision variable  $S$  is a rule-described subgroup rather than a single continuous variable, and the problem contains multiple objectives and constraints that are difficult to differentiate. Analytic solution or gradient descent methods are thus not applicable to this problem. For multi-objective problems, it is often impossible to obtain a single ideal optimal solution because optimizing one objective is likely to be at the expense of another objective. Therefore, the solution most often




Fig. 4: Illustration of the proposed model to discover optimal causal subgroups. (A) Schematic diagram of the Pareto front, where circles represent feasible subgroups for which small objective values are preferred over large values. The subgroup in Front 1 is not dominated by other subgroups, and the subgroup in Front 2 is dominated only by those in Front 1. (B) Illustration of the iterative heuristic algorithm for solving the multi-objective optimization problem. In each iteration, offspring subgroups are generated from the parent subgroup and survived by first comparing the front level and then the crowding distance.

consists of a series of non-dominated subgroups, more formally called the Pareto front. As shown in Fig. 4-(A), assuming two minimizing objectives  $f_1$  and  $f_2$  (if there are both maximizing and minimizing objectives, we can convert them to minimizing by adding a negative sign to maximizing), the circle represent a feasible solution (*i.e.*, a subgroup in our problem). Front 1 denotes the set of non-dominated solutions because no other solution is better than them. In mathematical terms, one solution  $s_1$  (Pareto) dominates another solution  $s_2$ , if

$$\begin{aligned} &\forall i \in \{1, \dots, m\}, f_i(s_1) \leq f_i(s_2), \text{ and} \\ &\exists i \in \{1, \dots, m\}, f_i(s_1) < f_i(s_2), \end{aligned} \quad (7)$$

where  $m$  is the number of objectives.

In order to efficiently find the subgroups belonging to the Pareto front, we employ a heuristic genetic search algorithm [17]. As shown in Fig. 4-(B), the algorithm mainly consists of 4 steps.

1. **Subgroups initialization.** We use binary random sampling to generate binary vectors, where 0 represents not selecting the corresponding covariates, and 1 indicates that it is selected, resulting in different antecedents for describing the subgroups. The default initial number of subgroups is 100.

2. **Offspring generation.** Existing subgroups are used to generate new subgroups by crossover and randomly flipping binary vectors. For example, a new subgroup is described by splicing the first half of the antecedent in Subgroup  $S_1$  onto the second half of the antecedent in Subgroup  $S_2$ . This step helps to expand the search space as much as possible without falling into a local optimum. Together, the existing and new subgroups form a candidate population.

3. **Non-dominated sorting.** The dominance relationship between all subgroups is obtained through a pairwise check. The subgroups that are not dominated by any other subgroups form Pareto front 1; the subgroups dominated only by the front 1 are the front 2, and so on. Therefore, all subgroups are divided into several ordered levels.




Fig. 5: Illustration of propensity score matching. Colors represent different covariates, and drug icons indicate treatment. Matching reduces confounding bias by finding comparable treatment and control units.

**4. Crowding distance sorting.** The crowding distance is the Manhattan distance in the objective space, which measures the distribution density of subgroups. For instance, in Fig. 4-(A), the crowding distance of subgroup  $S_1$  is  $\frac{d_1}{f_1^{\max} - f_1^{\min}} + \frac{d_2}{f_2^{\max} - f_2^{\min}}$ . The boundary subgroups in front (such as  $S_3$  and  $S_4$ ) will be given an infinite crowding distance. The subgroup with a larger crowding distance is given priority, thereby ensuring the diversity of the solution space. After sorting with the front as the first priority and crowding distance as the second priority, the top-ranked subgroups are retained, while other subgroups are discarded.

Steps 1-4 are performed iteratively until a predetermined number of iterations is reached or the target values of the subgroups no longer improve. During this process, solutions that do not satisfy the constraints or are duplicated are eliminated. The subgroups belonging to front 1 obtained at the end of the iterative process are used as outputs.

## 5.2 Treatment Effect Explanation

Unlike prediction or classification problems, the ground truth of treatment effects is often unknown because we cannot simultaneously observe two contrasting outcomes. Existing causal inference methods usually estimate treatment effects through confounding balancing and statistical inference, which are naturally uncertain. Therefore, providing reasonable explanations for treatment effects is crucial to help users interpret the results and enhance trust. (R5)

We utilize propensity score matching to generate explanations of treatment effects, which mimic the A/B testing familiar to users. As shown in Fig. 5, for subgroups automatically discovered or manually added by users, the units covered may have different covariates. For example, the elderly may have different demographic characteristics such as height and weight. To balance these confounders, we identify pairs of units from treatment and control groups whose covariates are similar or even identical, *i.e.*,  $\text{Distance}(\mathbf{X}_i, \mathbf{X}_j) \leq \epsilon$ . The naive matching method is exactly matching, which requires that the matched unit covariates are identical; however, in high-dimensional settings, there are rarely exact matches. Therefore, we match units by the distance between propensity score, *i.e.*,  $\text{Distance}(\mathbf{X}_i, \mathbf{X}_j) = |\hat{\epsilon}(\mathbf{X}_i) - \hat{\epsilon}(\mathbf{X}_j)|$ , which transforms the problem of high-dimensional space matching into scalar propensity score matching and controls for confounding bias. In the implementation, we set the threshold  $\epsilon = 0.1$  and greedily find the nearest matching unit in the control group for the units in the treatment group, guaranteeing that each unit matches at most once. Based on the difference in the outcomes (individual treatment effects) for the matched units, users can learn about the concentration and distribution of effects for that subgroup and thus judge the reliability of the estimated effects.

## 6 INTERFACE DESIGN

We developed a prototype system, *CausalPrism*, to help users implement the workflow proposed in Sec. 4. This section presents an overview of the system and details of visual design and interaction.

### 6.1 System Overview

As depicted in Fig. 1, *CausalPrism* offers three views: Causal Subgroup, Covariate Projection, and Treatment Effect Validation. These views enable subgroup exploration, comparison, and explanation of treatment effects. We demonstrate an analysis flow that uses these views to analyze causal heterogeneity in observational data. A data

analyst wants to see if the treatment affects the outcome differently across subgroups. She enters the data into the system and sets treatment, outcome, maximum antecedent length, and minimum coverage. The model (Sec. 5.1) automatically identifies subgroups with significant treatment effects (R1). The Causal Subgroup View (Fig. 1-A) offers an overview of causal subgroups, including their description and evaluation metrics (R2). Based on domain knowledge, she adds a new subgroup to analyze together (R3). In the Covariate Projection View (Fig. 1-B), she integrates small, close subgroups into a large one (R3). Since there are many possible subgroups, the multi-attribute ranking function is resorted to identify preferred ones (R4). To better understand the treatment effect in the subgroup, she consults the Treatment Effect Validation View (Fig. 1-C). The propensity score histogram shows the proportion of matched units in treatment and control groups. She then interprets the effect strength and uncertainty using dot plots of matched pairs and detailed information (Sec. 5.2, R5).

### 6.2 Causal Subgroup View

The Causal Subgroup View contains three parts: (1) a table-based antecedent visualization to intuitively convey the meaning of the subgroup (R2); (2) a subgroup editing window for additions and modifications to improve the identified subgroups (R3); and (3) an evaluation metrics visualization that supports multi-attribute ranking to facilitate personalized comparison of subgroups (R4).

Users need antecedents and consequents to grasp causal subgroups. Antecedents include covariates and values in the form of CNF, and consequents include evaluation metrics. Thus, various one-to-many relationships (subgroup  $\rightarrow$  covariates, covariate  $\rightarrow$  values, subgroup  $\rightarrow$  metrics) exist. Tabular, matrix, and multi-attribute ranking visualizations inspire us to intuitively visualize this information. Tabular/matrix forms like UpSet [45] and Taggle [21] are useful for analyzing set relationships. Multi-attribute ranking visualizations like LineUp [24] and SRVis [61] aid in multi-criteria decision-making. Our bespoke table-based view (Fig. 1-A1) shows subgroups as rows and covariates and evaluation metrics as columns. Expanding a covariate column displays its distribution. Bar charts display distinct value counts for category covariates. Smooth line charts show the count over numerical covariates' value domain. The cell displays subgroup covariate values as circles or rectangles, reflecting discrete values and continuous intervals. Clicking the "Add Subgroup" or "Edit Subgroup" button opens the "Subgroup Edit Box" (Fig. 1-A2), allowing users to alter covariate values and identify the subgroup with relevant domain semantics.

The right side of the table (Fig. 1-A3) displays a multi-attribute depiction of subgroup evaluation metrics, with each column representing a measure. Horizontal bars represent metric values. Users can apply an inverted mapping to measures like variance, where more minor is better. Users can drag to combine metrics and add weights to construct stacked bar charts to convey preferences. Sorting and filtering lets users find relevant subgroups faster.

### 6.3 Covariate Projection View

The Covariates Projection View (Fig. 1-B) displays each unit's low-dimensional projection. Users have trouble intuitively comparing units because raw observations include many covariates. To downscale high-dimensional units to the 2D plane, we use non-metric multidimensional scaling [41], where similar units are close. The scatter plot highlights units from the selected subgroup in the Causal Subgroup View, showing subgroup size and connectivity. To adjust subgroup size, users can click "Merge subgroup" or "split subgroup" (R3). Analysis noise can be reduced by using a switch to hide non-subgroup units.

### 6.4 Treatment Effect Validation View

The Validation View (Fig. 1-C) has three parts: (1) The histogram depicts the treatment and control group units' propensity score distribution. (2) The treatment effect dot plot shows the distribution of ITE of sampled matched pairs in the treatment and control groups. (3) The unit information table gives detailed treatment, covariates, and outcomes of matching units, enhancing users' trust (R5).

Table 1: Dataset statistics for quantitative experiments.

Dataset	#Units	#Categorical	#Numerical
Syn-1	3000	5	5
Syn-2	3000	5	15
Syn-3	4000	5	25
Syn-4	4000	5	45
Syn-5	4000	5	75
Syn-6	4000	5	95
Twins	23968	3	46
IHDP	747	19	6

The histogram (Fig. 1-C1) of unit distribution with varied propensity scores indicates covariate similarity across units. The propensity score is on the horizontal axis, and the number of units is on the vertical. The treatment and control groups had yellow and pink distributions. The overlapping distribution is brown. More overlap between these two groups' distributions means better covariate balancing, reducing effect estimation bias. The top of the histogram shows the number of units in the two groups and matched pairs.

The treatment effect dot plot (Fig. 1-C2) helps to understand how effect is calculated using the propensity score matching. Each dot represents a matched pair of treatment and control units. A dot's horizontal position indicates the ITE. The average of ITE is represented as vertical lines, while confidence intervals are represented as gray rectangles at the bottom of the axis. The unit information table (Fig. 1-C3) provides details on matched units. The left columns represent unit ID, propensity score, treatment, and outcome. Other columns are covariates. Clicking on the dot in the dot plot highlights corresponding rows. Unit-based data can help users understand how treatment affects outcomes through particular examples.

## 7 EVALUATION

In this section, we implement quantitative experiments to evaluate the causal subgroup discovery model. Two case studies and expert interviews further validate the usefulness of *CausalPrism*.

### 7.1 Quantitative Experiments

The quantitative experiment aims to assess the efficacy of the proposed model in identifying subgroups with significant treatment effects. Based on multiple synthetic datasets and real datasets, we compare the model with various baselines on different metrics.

**Datasets.** We employed synthetic datasets and real-world datasets. Following the settings in [2, 63], we sampled units under the assumption of unconfoundedness. Some covariates are categorical, and others are numerical with a normal distribution. We simulated non-random treatment assignment in observational data by creating a treatment variable  $T$  determined by a Bernoulli distribution. We also produced the treatment effect TE and the outcome  $Y$ , calculated from covariates and parameter vectors. Categorical covariates are converted to one-hot encoding for computation. For detailed descriptions of the generation process, please see OSF. We also collected real-world dataset including Twins<sup>1</sup> and IHDP<sup>2</sup>. The details of the datasets are shown in Tab. 1.

**Baselines.** We compare the proposed model with two groups of algorithms. The first group is the popular HTE estimation algorithms: (1) Causal Tree (CT) [2]; (2) Causal Forest (CF) [57]; and (3) Causal Rule Ensemble (CRE) [9]. The second group is the rule learning and subgroup discovery algorithms: (1) BRCA [16]; (2) Decision Tree (DT) [11]; (3) Pysubgroup (PYS) [44]. In the first group, CRE can explicitly obtain the antecedent and treatment effect of the subgroup. For CT and CF, it can be considered that the path from the root to the leaf nodes in the tree structure is the antecedent of the causal subgroup. The second group of methods can only get the correlation subgroups.

<sup>1</sup><https://github.com/AMLab-Amsterdam/CEVAE/tree/master/datasets/TWINS>

<sup>2</sup><https://search.r-project.org/CRAN/refmans/bartcs/html/ihdp.html>

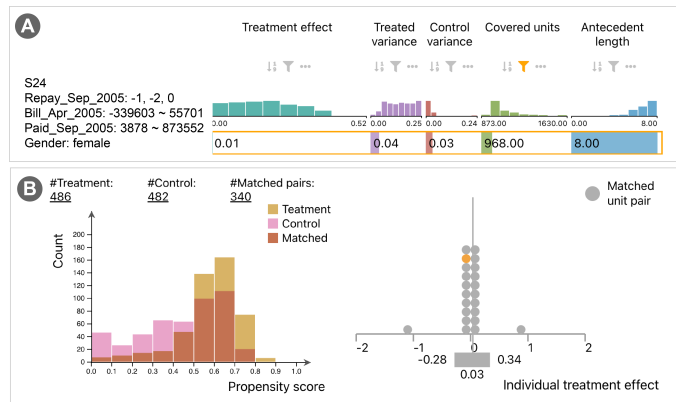


Fig. 6: Descriptions of a subgroup with good credit in Case 1. (A) Explanation of treatment effects shows that the subgroup has low effect and variance. (B) A balanced histogram of propensity scores and most matched pairs have a zero ITE.

In order to adapt to the causality setting, we add a post-processing step. CATE and variance are calculated on the data covered by each subgroup via Eq. (4) and Eq. (5).

**Metrics.** We evaluate the quality of causal subgroups obtained from different perspectives. First, in order to evaluate the multi-objective optimization of treatment effect and outcome variance, it is proposed that (1) Precision(P) = (the true number of dominating subgroups)/(the number of subgroups in the front discovered by the method). Due to the lack of ground truth for subgroups belonging to the Pareto front, we collected subgroups in the front obtained by all methods and assumed that a subgroup is considered a true dominating subgroup if it is not dominated by any other subgroup. We also considered the interpretability of subgroups, including the metrics (2) #Subgroups(S) = number of subgroups in front, (3) Avg\_len(L) = average length of antecedent (*i.e.*, number of covariates) used to describe the subgroups and (4) Coverage(C%) = The average percent of units in a subgroup to the total number of units.

**Result analysis.** The experimental results are reported in Tab. 2. Our model has near-perfect precision (the bigger, the better), indicating that for “Pareto fronts” in other methods, our model always finds a dominant subgroup that is better in at least one objective. This is mainly due to the fact that we directly formalize and solve the constrained multi-objective optimization problem. Some methods use a two-stage approach to subgroup generation and selection, such as CT and CF, which partition the covariate space through trees and select the best subgroups. DT, on the other hand, only considers covariates and outcomes, ignoring treatment changes. Missing important subgroups in the initial stage may result in suboptimal results and reduced precision. We also uncover more diverse subgroups, averaging 15.4 subgroups distributed over the Pareto front, which is about 3 times that of other techniques. The small antecedent length and coverage of our method indicate fine-grained subgroups. Adjusting the minimal coverage limit can yield larger subgroups. As the number of covariates increases (10 to 100), our model remains stable, demonstrating that the efficient heuristic genetic algorithm can find satisfactory subgroups in high-dimensional covariate spaces, whereas other approaches or manual selection fail.

### 7.2 Case Studies

We use two cases to demonstrate the analysis process of *CausalPrism* to explore causal heterogeneity from real datasets.

#### 7.2.1 Case 1: Default of Credit Card Clients

This dataset contains behavioral data of about 30,000 credit card customers. The treatment is the credit limit (1 for greater than 200,000 and 0 for less), the outcome is default status (1 is defaulted, 0 is not), and the covariates include gender, education, marital status, age, and historical bill amounts *etc.* An account manager plans to raise credit

Table 2: Quantitative metrics for different causal subgroups discovery methods. **Black**: best. Underline: second best. (The experimental results of CRE and BRCG on Twins are missing due to code execution errors.)

Methods	Ours				CT				CF				CRE				DT				PYS				BRCG							
Metric	P	S	L	C	P	S	L	C	P	S	L	C	P	S	L	C	P	S	L	C	P	S	L	C	P	S	L	C	P	S	L	C
Syn-1	<b>1.0</b>	17	4.0	3.8	0.0	6	6.0	7.0	0.3	8	5.7	3.2	<u>0.5</u>	2	2.3	30.2	0.0	8	3.8	9.1	0.0	3	1.4	13.7	0.0	2	3.0	12.5	0.0	2	3.0	12.5
Syn-2	<b>0.9</b>	14	3.0	3.5	<u>0.3</u>	3	4.3	12.8	0.0	7	4.5	3.2	0.1	8	2.3	28.7	0.1	7	4.0	8.6	0.0	3	1.2	16.9	0.0	3	3.0	11.8	0.0	3	3.0	11.8
Syn-3	<b>1.0</b>	17	3.8	3.7	<b>1.0</b>	3	6.4	5.8	0.0	6	5.2	3.1	0.0	3	2.1	32.7	0.0	4	4.0	9.9	0.0	5	1.0	20.1	0.0	3	3.2	9.1	0.0	3	3.2	9.1
Syn-4	<b>1.0</b>	18	3.6	3.5	0.4	5	6.2	7.6	<u>0.5</u>	2	4.4	2.7	0.2	5	1.7	32.2	0.2	5	4.0	8.3	0.0	8	1.0	20.1	0.0	1	10.0	35.4	0.0	1	10.0	35.4
Syn-5	<u>0.9</u>	19	4.0	3.4	0.0	7	8.5	5.2	0.0	2	5.5	2.5	<b>1.0</b>	1	1.7	35.7	0.0	4	4.0	12.5	0.0	3	1.0	20.0	0.0	1	8.0	33.5	0.0	1	8.0	33.5
Syn-6	<b>1.0</b>	13	3.1	3.2	0.0	6	5.4	10.6	0.0	1	4.4	2.6	0.0	3	2.6	31.3	0.0	3	4.0	9.6	0.0	11	1.0	20.1	0.0	1	1.0	40.0	0.0	1	1.0	40.0
Twins	<u>0.8</u>	15	4.0	4.0	0.2	5	5.1	9.7	<b>1.0</b>	2	5.7	4.4	/	/	/	/	0.0	2	2.9	13.9	0.0	7	2.6	21.1	/	/	/	/	/	/	/	/
IHDP	<b>1.0</b>	10	4.0	3.1	<u>0.1</u>	7	4.1	7.5	0.0	11	3.1	11.9	0.0	3	1.5	28.8	0.0	3	2.6	18.1	0.0	2	2.8	19.4	0.0	2	2.5	24.2	0.0	2	2.5	24.2
Average	<b>1.0</b>	15.4	3.7	3.5	<u>0.3</u>	5.3	5.8	8.3	0.2	4.9	4.8	4.2	<u>0.3</u>	3.6	2.0	31.4	0.0	4.5	3.7	11.0	0.0	5.3	1.5	19.0	0.0	1.9	4.4	23.8	0.0	1.9	4.4	23.8

limits to boost interest profits and transaction volume but avoids default risk by using *CausalPrism* to identify acceptable credit recipients.

He first loads the observational data on client behavior and sets the maximum subgroup antecedent length to 7 and the minimum coverage to 5%. Afterward, the Causal Subgroup View (Fig. 1-A1) displays the causal subgroups that the model automatically mines from the data (R1). He scrolls through the table to see descriptions of the covariates involved in the subgroup antecedents and clicks to see the details (R2). He found that two subgroups had similar distributions of covariates and were both of small size, so he clicked the “Merge Subgroups” button to generate a large one (R3). He wanted to study risky subgroups and creditworthy subgroups, but checking each subgroup one by one was tedious, so he used the multi-attributes ranking function. To find risky subgroups, he first set an inverted mapping from a variance value to the length of a bar, that is, a longer bar represents less variance. Then, he combined the treatment effects, the outcome variances of the treatment and the control group as objectives of subgroup identification and assigned weights of 6:2:2. As shown in Fig. 1-A3, after sorting in descending order of the combined column and filtering by the number of covered units greater than 950, S98 rose to the first place (R4), which includes past bill amounts, repayment records, etc. S98 has an effect of 0.4, a treatment group outcome variance of 0.24, and a control group outcome variance of 0.03. He explained that these customers often overdue their repayments (the repay status indicates that the number of overdue months is up to 8), and their recent bill amounts are relatively high. Therefore, increasing their credit limit makes them more prone to defaults and a priority group to monitor.

To further verify the causal subgroup (R5), he went to the Treatment Effect Validation View (Fig. 1-C). In S98, the treatment group contains 722 units, the control group contains 238 units, and the number of matched unit pairs reaches 208. The propensity score histogram proves that for each unit in the control group, a similar matching unit can be found in the treatment group almost all the time. In the dot plot of the ITE, the dots are mostly distributed in the middle or on the right side, indicating that increasing the credit limit has a positive effect on default. The detailed information in the table suggests that the outcome goes from 0 to 1 (i.e., from non-default to default) in matched pairs of similar units (e.g., units with id 2184 and 6246) after the treatment, which enhances his trust in the causal conclusion.

He next sought higher-credit subgroups for comparison. According to experience, these subgroups should have lower effects and variances, so he canceled the invert mapping of variances and sorted the combined columns in ascending order (R4). As shown in Fig. 6, the first subgroup, S24, caught his attention. It had an effect of 0.01, a treatment group outcome variance of 0.04, a control group outcome variance of 0.03, and a coverage of 968 people. Based on the antecedent description, he believed that this group had good repayment records in the past (the repay status is -1, -2, and 0) and had the habit of automatic repayment

(the presence of a negative value of bill represents an automatic deposits into the credit card every month). The dot plot with dots concentrated at 0 also proves that increasing the limit hardly causes defaults.

## 7.2.2 Case 2: Bank Marketing

The dataset comes from a bank’s marketing campaign, in which the treatment is the number of phone calls to customers (1 for more than two times, 0 otherwise), and the outcome is whether the customer makes a deposit (1 yes, 0 no), and the other 15 covariates consist of the customer’s age, occupation, marital status, education level, recent contact data, and other socioeconomic indicators. An analyst wanted to identify the groups where increased contact could lead to deposits.

With the Causal Subgroup View, she found that many subgroups have covariates of job, month, and duration, aligning with her domain knowledge, as there are differences in the financial capacity of clients across occupations and deposits have low and peak times of the year (R1, R2). The length of the last contact also reflects the customer’s wishes. Since telemarketing failure will not bring serious consequences, she placed greater emphasis on the effect strength and customer coverage (R4). Then, she combined the effect and coverage metrics, set the weights to 70% and 30% respectively, and sorted to get the subgroup S70 that suits her preference (Fig. 7-A). The subgroup is described as job: blue-collar, entrepreneur, management, retired, etc.; duration: 204-3881, age: 31-95; contact day: 1-13. The corresponding effect is 0.31 with 394 covered units. She guessed that these middle-aged and elderly people had a sound financial foundation, and previous contacts also showed that they were interested. In addition, the beginning of the month is suitable for promoting deposits because wages are usually just paid.

To verify the reliability of the effect, she turned to the Treatment Effect Validation View. She recognized that almost all units in the treatment group had matching counterparts in the control group, ensuring the validity of the effect estimates. The dot plot revealed a high mean and a wide spread of dots, implying greater uncertainty (R5). Leveraging the detailed information provided by the table (Fig. 7-B), she found that an increase in the number of contacts indeed promotes deposits, so she decided to start by selecting those customers (e.g., the user with id 10888) with fewer contacts in the matched pair with an individual treatment effect of 1 since they are more likely to purchase deposits after several more contacts.

## 7.3 Expert Interviews

We conducted informal interviews with four experts, two of whom (E1, E2) are data analysts at technology companies, E3 is a Ph.D. student studying causality, and E4 is a researcher at a university medical school who studies big data mining for medical data. They had not been involved in the previous design process.



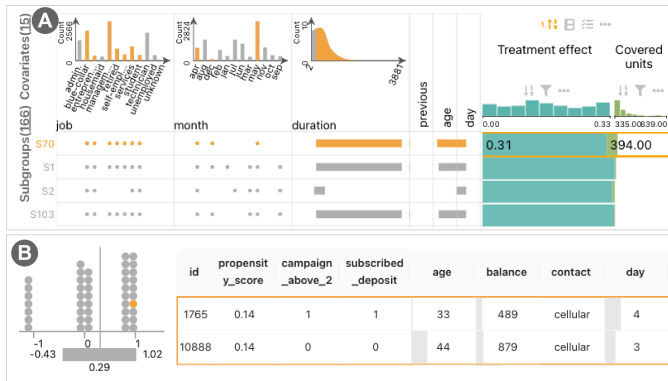


Fig. 7: The subgroup identified in Case 2 as suitable for multiple contact to market deposits. (A) This subgroup is mostly people who have good jobs and showed interest in the last communication. They also have a large treatment effect and coverage rate. (B) The dot plot of ITE is skewed towards the strong side of the effect. Although it has a large variance, it is within the acceptable range of the user.

**Procedure.** Each interview consisted of four steps, which are an introduction of background (5min), a demo of usage cases (15min), a think-aloud exploration (15min), and a talk for feedback collection (10min). Finally, we summarized comments, which includes reviews on both the proposed model and the prototype system, and suggestions.

**Reviews on the proposed model.** Experts believe the causal subgroup discovery model can find subgroups with significant treatment effects compared to the population average. E1 said, “When facing multi-dimensional observation data, I am usually left with time-consuming manual subgroup analysis, whereas this model gives me quick guidance on subgroups and helps to uncover unexpected conclusions.” E3 praised that rules make subgroups easier to understand than clustering. Also, causal tree or forest-based approaches include numerous repeated covariate descriptions since many leaves share ancestor routes. Solving subgroups with multi-objective optimization is straightforward and yields better results. E4 stated that it would be more effective if they could freely combine different objectives and constraints.

**Reviews on the prototype system.** Experts confirmed that *CausalPrism*'s visual design is easy to grasp. E1 and E2 thought the Causal Subgroup View was intuitive because they utilized similar Excel spreadsheets for work. E4 recommended dot plots for their clear explanation of the trade-off between bias and variation in treatment effects estimation. E3 said the multi-attribute ranking visualization is useful for decision-making situations where users have diverse preferences for several objectives, such as weighing revenues and costs when choosing an advertisement audience. E4 highlighted that adding and modifying subgroups helps incorporate domain knowledge and improves analytical freedom beyond model outputs.

**Suggestions.** Experts offered insightful advice on *CausalPrism*. E2's advice is to the function of merging subgroups with neighboring attribute values on the covariates, which can reduce the number of fragmented subgroups. Taking the advice, we enabled subgroup editing options, consisting of merging or splitting subgroups based on similarity and size, in the Covariate Projection View. E4 told us that manual editing of subgroups would be tedious if many predefined subgroups needed to be introduced. It would be nice if the Causal Subgroup View could automatically generate subgroups that have been recorded by experts by connecting with external knowledge bases. We plan to explore this feature by taking knowledge graphs as input.

## 8 DISCUSSION

This section discusses the scalability, lessons learned, and the limitations and future work of *CausalPrism*.

**Scalability.** (1) The running efficiency of the proposed model is mainly affected by the number of subgroups  $P$ . The complexity in the subgroup generation stage is  $O(Pd)$ , where each subgroup could be

described by at most  $d$  covariates. The following three steps, checking the dominance relation in pairs, calculating the crowding distance of each subgroup, and sorting and selection, have complexities of  $O(P^2)$ ,  $O(P)$ , and  $O(P \log P)$ , respectively. The overall time complexity is  $O(Pd + P^2)$ . (2) For visualization, the covariate table may become crowded as the number of covariates increases; we moved the covariates that appear more frequently in the subgroup antecedents forward and added sliders and column folding functions to alleviate this. The dot plot in the Treatment Effect Validation View may appear to be overplotted due to an excessive number of matching units. We used sampling to reduce the number of dots and maintain the shape of the distribution. Other possible solutions are non-linear dot plots with adaptive dot sizes [49].

**Lessons learned.** We gained valuable insights by developing *CausalPrism*. (1) Whether to analyze the heterogeneity of the causal graph structure or the HTE depends on the task. Causal discovery can identify many causal relationships among variables, such as gene regulatory networks. In the absence of prior knowledge of the causal structure, causal inference is a more general method, such as studying the treatment effect of policies on economic indicators. (2) Causal heterogeneity analysis benefits from multi-objective optimization. Our initial optimization target was treatment effect strength, but experts stressed that in actual applications, they would also evaluate many objectives, such as outcome variance, cost per unit, and return on investment (ROI). We then reduce multi-objectives optimization to single-objectives optimization by weighting (e.g.,  $\tau + w_0\sigma^2(0) + w_1\sigma^2(1)$ ). However, experts said weights are sensitive and hard to modify, so the Pareto front of the MOO problem was eventually learned. We take effect and variances as objectives, but they can be expanded flexibly, which helps to meet diverse user preferences. (3) Explanations can support human-in-the-loop causal analysis, but too much information causes cognitive overload. Our initial prototype solution used a white-box causal subgroup discovery model based on decision trees, trading performance for interpretability. However, experts said they value the model's precision in real scenarios. If the model is bad, explaining it is pointless. Thus, we proposed an effective causal subgroup discovery model with a popular CATE estimator (IPW) and a heuristic search algorithm. We also provided post-hoc visualizations of subgroups and treatment effects to aid user interpretation.

**Limitations and future work.** Three limitations are observed in *CausalPrism*. First, subgroups support understanding data from the level of groups but may obscure individual uniqueness. For numerical attributes, units in the same subgroup may still have subtle covariate differences, and similar covariate values may belong to separate subgroups due to boundary divisions. Also, subgroup quality can be assessed based on external or internal similarity metrics. Second, our approach considers a single outcome. A treatment may have multiple outcomes. For example, a medicine may cause mortality, disease progression, or adverse events. A comprehensive evaluation of the treatment effectiveness should consider all possible outcomes. Future plans include combining causal inference approaches like mediation analysis [53] and multi-task learning [1] for multiple outcomes. Third, we assume enough confounding variables in observational data. However, unobserved variables may distort effect estimates in some cases. For example, patients' diets may also affect treatment effects. Combining RCTs and observational studies [15] is a promising approach.

## 9 CONCLUSION

In this paper, we propose a visual analytics approach to support users in identifying, comparing, ranking, and validating subgroups with significant treatment effects in observational data. We first introduce interpretable rules to describe subgroups and then formalize the causal subgroup discovery into a constrained multi-objective optimization problem, whose corresponding Pareto optimal subgroups are efficiently solved by a heuristic genetic algorithm. A visual analysis system, *CausalPrism*, is developed to implement the proposed approach, which features a series of visualizations such as subgroups—covariate tables, multi-attribute rankings, and treatment effect explanations. Quantitative experiments prove that our model outperforms state-of-the-art methods

in terms of precision and readability. Case studies and expert interviews demonstrate the usability of the system in helping users explore and interpret subgroups that satisfy different preferences.

#### ACKNOWLEDGMENTS

We would like to thank all the reviewers for their constructive comments. This work was supported by the National Natural Science Foundation of China (# 62132017) and Zhejiang Provincial Natural Science Foundation of China (# LD24F020011).

#### REFERENCES

- [1] A. M. Alaa and M. Van Der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *Advances in Neural Information Processing Systems*, 30, 2017. 9
- [2] S. Athey and G. Imbens. Recursive partitioning for heterogeneous causal effects. *Proc. NAS*, 113(27):7353–7360, 2016. doi: 10.1073/pnas.1510489113 2, 7
- [3] S. Athey and S. Wager. Estimating treatment effects with causal forests: An application. *Observational Studies*, 5(2):37–51, 2019. doi: 10.1353/obs.2019.0001 2
- [4] M. Atzmueller. Subgroup discovery. *WIREs Data Mining and Knowledge Discovery*, 5(1):35–49, 2015. doi: 10.1002/widm.1144 3
- [5] M. Atzmüller and F. Lemmerich. Fast subgroup discovery for continuous target concepts. In *Proc. ISMIS*, p. 35–44, 2009. doi: 10.1007/978-3-642-04125-9\_7 3
- [6] M. Atzmüller and F. Puppe. Sd-map - a fast algorithm for exhaustive subgroup discovery. In *Proc. ECML PKDD*, p. 6–17, 2006. doi: 10.1007/11871637\_6 2, 3
- [7] J. Bae, T. Helldin, and M. Riveiro. Understanding indirect causal relationships in node-link graphs. *Computer Graphics Forum*, 36(3):411–421, 2017. doi: 10.1111/cgf.13198 3
- [8] J. Bae, E. Ventocilla, M. Riveiro, T. Helldin, and G. Falkman. Evaluating multi-attributes on cause and effect relationship visualization. In *Proc. VISIGRAPP*, pp. 64–74, 2017. doi: 10.5220/0006102300640074 3
- [9] F. J. Bargagli-Stoffi, R. Cadei, K. Lee, and F. Dominici. Causal rule ensemble: Interpretable discovery and inference of heterogeneous causal effects. *arXiv preprint arXiv:2009.09036*, 2020. doi: 10.48550/arXiv.2009.09036 2, 7
- [10] C. J. Blumberg. Causal inference for statistics, social, and biomedical sciences: An introduction. *International Statistical Review*, 84(1):159–159, 2016. doi: 10.1111/insr.12170 4
- [11] L. Breiman and R. A. Olshen. Points of significance: Classification and regression trees. *Nature Methods*, 14:757–758, 2017. doi: 10.1038/nmeth.4370 7
- [12] T. Cai, L. Tian, P. H. Wong, and L. J. Wei. Analysis of randomized comparative clinical trial data for personalized treatment selections. *Biostatistics*, 12(2):270–282, 2011. doi: 10.1093/biostatistics/kxq060 2
- [13] M. Caliendo and S. Kopeinig. Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72, 2008. doi: 10.1111/j.1467-6419.2007.00527.x 4
- [14] F. Cheng, Y. Ming, and H. Qu. Dece: Decision explorer with counterfactual explanations for machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1438–1447, 2021. doi: 10.1109/TVCG.2020.3030342 3
- [15] B. Colnet et al. Causal inference methods for combining randomized trials and observational studies: a review. *arXiv preprint arXiv:2011.08047*, 2020. doi: 10.48550/arXiv.2011.08047 9
- [16] S. Dash, O. Günlük, and D. Wei. Boolean decision rules via column generation. In *Proc. NeurIPS*, 2018. doi: 10.48550/arXiv.1805.09901 7
- [17] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multi-objective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017 5
- [18] M. J. del Jesús, P. González, F. Herrera, and M. Mesonero. Evolutionary fuzzy rule induction process for subgroup discovery: A case study in marketing. *IEEE Transactions on Fuzzy Systems*, 15(4):578–592, 2007. doi: 10.1109/TFUZZ.2006.890662 3
- [19] Z. Deng, D. Weng, X. Xie, J. Bao, Y. Zheng, M. Xu, W. Chen, and Y. Wu. Compass: Towards better causal analysis of urban time series. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):1051–1061, 2022. doi: 10.1109/TVCG.2021.3114875 3
- [20] D. Dingen et al. Regressionexplorer: Interactive exploration of logistic regression models with subgroup analysis. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):246–255, 2019. doi: 10.1109/TVCG.2018.2865043 3
- [21] K. Furmanova et al. Taggle: Combining overview and details in tabular data visualizations. *Information Visualization*, 19(2):114–136, 2020. doi: 10.1177/1473871619878085 3, 6
- [22] D. Gamberger and N. Lavrac. Expert-guided subgroup discovery: Methodology and application. *Journal of Artificial Intelligence Research*, 17:501–527, 2002. doi: 10.1613/jair.1089 2, 3
- [23] M. Gangl. Causal inference in sociological research. *Annual review of sociology*, 36(1):21–47, 2010. doi: 10.1146/annurev.soc.012809.102702

2

- [24] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. Lineup: Visual analysis of multi-attribute rankings. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2277–2286, 2013. doi: [10.1109/TVCG.2013.173](https://doi.org/10.1109/TVCG.2013.173) 3, 6
- [25] N. Greifer. *cobalt: Covariate Balance Tables and Plots*, 2024. R package version 4.5.4, <https://github.com/ngreifer/cobalt>. 3
- [26] H. Grosskreutz and S. Rüping. On subgroup discovery in numerical domains. *Data Mining and Knowledge Discovery*, 19:210–226, 2009. doi: [10.1007/s10618-009-0136-3](https://doi.org/10.1007/s10618-009-0136-3) 3
- [27] H. Grosskreutz, S. Rüping, and S. Wrobel. Tight optimistic estimates for fast subgroup discovery. In *Proc. ECML PKDD*, pp. 440–456, 2008. doi: [10.1007/978-3-540-87479-9\\_47](https://doi.org/10.1007/978-3-540-87479-9_47) 3
- [28] G. Guo, M. Glenski, Z. Shaw, E. Saldanha, A. Endert, S. Volkova, and D. Arendt. Vaine: Visualization and ai for natural experiments. In *Proc. VIS*, pp. 21–25, 2021. doi: [10.1109/VIS49827.2021.9623285](https://doi.org/10.1109/VIS49827.2021.9623285) 2, 3
- [29] G. Guo, E. Karavani, A. Endert, and B. C. Kwon. Causalvis: Visualizations for causal inference. In *Proc. CHI, CHI '23*, article no. 462, 20 pages, 2023. doi: [10.1145/3544548.3581236](https://doi.org/10.1145/3544548.3581236) 2, 3
- [30] R. Guo, L. Cheng, J. Li, P. R. Hahn, and H. Liu. A survey of learning causality with data: Problems and methods. *ACM Computing Surveys*, 53(4), article no. 75, 37 pages, jul 2020. doi: [10.1145/3397269](https://doi.org/10.1145/3397269) 2
- [31] E. Hariton and J. J. Locascio. Randomised controlled trials - the gold standard for effectiveness research: Study design: randomised controlled trials. *BJOG*, 125(13):1716, June 2018. 2
- [32] F. Herrera, C. J. Carmona, P. González, and M. J. del Jesus. An overview on subgroup discovery: foundations and applications. *Knowledge and Information Systems*, 29:495–525, 2011. doi: [10.1007/s10115-010-0356-2](https://doi.org/10.1007/s10115-010-0356-2) 3
- [33] J. L. Hill. Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240, 2011. doi: [10.1198/jcgs.2010.08162](https://doi.org/10.1198/jcgs.2010.08162) 2
- [34] K. Hirano, G. Imbens, and G. Ridder. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica*, 71(4):1161–1189, 2003. doi: [10.1111/1468-0262.00442](https://doi.org/10.1111/1468-0262.00442) 5
- [35] 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, 2022. doi: [10.1109/TVCG.2021.3102051](https://doi.org/10.1109/TVCG.2021.3102051) 3
- [36] G. W. Imbens. Nonparametric Estimation of Average Treatment Effects Under Exogeneity: A Review. *The Review of Economics and Statistics*, 86(1):4–29, 2004. doi: [10.1162/003465304323023651](https://doi.org/10.1162/003465304323023651) 5
- [37] 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, 2021. doi: [10.1109/TVCG.2020.3030465](https://doi.org/10.1109/TVCG.2020.3030465) 2, 3
- [38] F. Johansson, U. Shalit, and D. Sontag. Learning representations for counterfactual inference. In *Proc. ICML*, vol. 48, pp. 3020–3029, 2016. doi: [10.48550/arXiv.1605.03661](https://doi.org/10.48550/arXiv.1605.03661) 2
- [39] N. R. Kadaba, P. Irani, and J. Leboe. Visualizing causal semantics using animations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1254–1261, 2007. doi: [10.1109/TVCG.2007.70528](https://doi.org/10.1109/TVCG.2007.70528) 3
- [40] 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, 2022. doi: [10.1109/TVCG.2021.3114824](https://doi.org/10.1109/TVCG.2021.3114824) 3
- [41] J. B. Kruskal. Nonmetric multidimensional scaling: a numerical method. *Psychometrika*, 29(2):115–129, 1964. 6
- [42] B. C. Kwon et al. Rmexplorer: A visual analytics approach to explore the performance and the fairness of disease risk models on population subgroups. In *Proc. VIS*, pp. 50–54, 2022. doi: [10.1109/VIS54862.2022.00019](https://doi.org/10.1109/VIS54862.2022.00019) 3
- [43] N. Lavrač, B. Kavšek, P. A. Flach, and L. Todorovski. Subgroup discovery with cn2-sd. *The Journal of Machine Learning Research*, 5:153–188, 2004. 3
- [44] F. Lemmerich and M. Becker. pysubgroup: Easy-to-use subgroup discovery in python. In *Proc. ECML PKDD*, pp. 658–662, 2018. doi: [10.1007/978-3-030-10997-4\\_46](https://doi.org/10.1007/978-3-030-10997-4_46) 7
- [45] A. Lex, N. Gehlenborg, H. Strobel, R. Vuillemot, and H. Pfister. Upset: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1983–1992, 2014. doi: [10.1109/TVCG.2014.2346248](https://doi.org/10.1109/TVCG.2014.2346248) 6
- [46] J. Li et al. From observational studies to causal rule mining. *ACM Transactions on Intelligent Systems and Technology*, 7(2):1–27, 2015. doi: [10.1145/2746410](https://doi.org/10.1145/2746410) 2
- [47] J. Li, C. Lai, and X. Yuan. Subspace-map: Interactive visual analysis for subspace data with a map metaphor. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–15, 2024. doi: [10.1109/TVCG.2024.3368094](https://doi.org/10.1109/TVCG.2024.3368094) 3
- [48] J. Peters, D. Janzing, and B. Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017. doi: [10.1080/00949655.2018.1505197](https://doi.org/10.1080/00949655.2018.1505197) 1
- [49] N. Rodrigues and D. Weiskopf. Nonlinear dot plots. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):616–625, 2018. doi: [10.1109/TVCG.2017.2744018](https://doi.org/10.1109/TVCG.2017.2744018) 9
- [50] K. J. Rothman and S. Greenland. Causation and causal inference in epidemiology. *American journal of public health*, 95(S1):S144–S150, 2005. doi: [10.2105/AJPH.2004.059204](https://doi.org/10.2105/AJPH.2004.059204) 2
- [51] J. Sekhon. 271 The Neyman—Rubin Model of Causal Inference and Estimation Via Matching Methods. In *The Oxford Handbook of Political Methodology*. Oxford University Press, 08 2008. doi: [10.1093/oxfordhb/9780199286546.003.0011](https://doi.org/10.1093/oxfordhb/9780199286546.003.0011) 3
- [52] Y. Shimoni et al. An evaluation toolkit to guide model selection and cohort definition in causal inference. *ArXiv*, abs/1906.00442, 2019. 3
- [53] T. R. Ten Have and M. M. Joffe. A review of causal estimation of effects in mediation analyses. *Statistical Methods in Medical Research*, 21(1):77–107, 2012. 9
- [54] X. Teng, Y. Ahn, and Y.-R. Lin. Vispur: Visual aids for identifying and interpreting spurious associations in data-driven decisions. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):219–229, 2024. doi: [10.1109/TVCG.2023.3326587](https://doi.org/10.1109/TVCG.2023.3326587) 3
- [55] M. van Leeuwen and A. J. Knobbe. Diverse subgroup set discovery. *Data Mining and Knowledge Discovery*, 25:208–242, 2012. doi: [10.1007/s10618-012-0273-y](https://doi.org/10.1007/s10618-012-0273-y) 2, 3
- [56] H. R. Varian. Causal inference in economics and marketing. *Proc. NAS*, 113(27):7310–7315, 2016. doi: [10.1073/pnas.1510479113](https://doi.org/10.1073/pnas.1510479113) 2
- [57] S. Wager and S. Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018. doi: [10.1080/01621459.2017.1319839](https://doi.org/10.1080/01621459.2017.1319839) 2, 7
- [58] 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, 2016. doi: [10.1109/TVCG.2015.2467931](https://doi.org/10.1109/TVCG.2015.2467931) 2, 3
- [59] J. Wang and K. Mueller. Visual causality analysis made practical. In *Proc. VAST*, pp. 151–161, 2017. doi: [10.1109/VAST.2017.8585647](https://doi.org/10.1109/VAST.2017.8585647) 2, 3
- [60] J. Wang and K. Mueller. Domino: Visual causal reasoning with time-dependent phenomena. *IEEE Transactions on Visualization and Computer Graphics*, 2022. doi: [10.1109/tvcg.2022.3207929](https://doi.org/10.1109/tvcg.2022.3207929) 3
- [61] D. Weng, R. Chen, Z. Deng, F. Wu, J. Chen, and Y. Wu. Srvis: Towards better spatial integration in ranking visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):459–469, 2019. doi: [10.1109/TVCG.2018.2865126](https://doi.org/10.1109/TVCG.2018.2865126) 6
- [62] S. Wrobel. An algorithm for multi-relational discovery of subgroups. In *Proc. ECML PKDD*, p. 78–87, 1997. doi: [10.1007/3-540-63223-9\\_108](https://doi.org/10.1007/3-540-63223-9_108) 3
- [63] A. Wu, K. Kuang, R. Xiong, B. Li, and F. Wu. Stable estimation of heterogeneous treatment effects. In *Proc. ICML*, 2023. 7
- [64] J. Xia, F. Ye, W. Chen, Y. Wang, W. Chen, Y. Ma, and A. K. Tung. Ldscanner: Exploratory analysis of low-dimensional structures in high-dimensional datasets. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):236–245, 2018. doi: [10.1109/TVCG.2017.2744098](https://doi.org/10.1109/TVCG.2017.2744098) 3
- [65] P. Xie, W. Tao, J. Li, W. Huang, and S. Chen. Exploring multi-dimensional data via subset embedding. *Computer Graphics Forum*, 40(3):75–86, 2021. doi: [10.1111/cgf.14290](https://doi.org/10.1111/cgf.14290) 3
- [66] 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, 2021. doi: [10.1109/TVCG.2020.3028957](https://doi.org/10.1109/TVCG.2020.3028957) 2, 3
- [67] 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, 2020. doi: [10.1109/TVCG.2019.2934399](https://doi.org/10.1109/TVCG.2019.2934399) 3
- [68] L. Yao, Z. Chu, S. Li, Y. Li, J. Gao, and A. Zhang. A survey on causal inference. *ACM Transactions on Knowledge Discovery from Data*, 15(5), article no. 74, 46 pages, may 2021. doi: [10.1145/3444944](https://doi.org/10.1145/3444944) 2
- [69] C. E. Yen, A. G. Parameswaran, and W. Fu. An exploratory user study of visual causality analysis. *Computer Graphics Forum*, 38(3):173–184,

2019. doi: [10.1111/cgf.13680](https://doi.org/10.1111/cgf.13680) 3
- [70] C.-H. Yen, H. Cheng, Y. Xia, and Y. Huang. Crowdidea: Blending crowd intelligence and data analytics to empower causal reasoning. In *Proc. CHI, CHI '23*, article no. 463, 17 pages, 2023. doi: [10.1145/3544548.3581021](https://doi.org/10.1145/3544548.3581021) 3
- [71] X. Zhang, J. P. Ono, H. Song, L. Gou, K.-L. Ma, and L. Ren. Sliceteller: A data slice-driven approach for machine learning model validation. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):842–852, 2023. doi: [10.1109/TVCG.2022.3209465](https://doi.org/10.1109/TVCG.2022.3209465) 3
- [72] E. Zitzler, M. Laumanns, and L. Thiele. Spea2: Improving the strength pareto evolutionary algorithm. *TIK report*, 103, 2001. doi: [10.3929/ethz-a-004284029](https://doi.org/10.3929/ethz-a-004284029) 3