Summarized Causal Explanations For Aggregate Views

BRIT YOUNGMANN, Technion - Israel Institute of Technology, Israel MICHAEL CAFARELLA, CSAIL MIT, USA AMIR GILAD, Hebrew University, Israel SUDEEPA ROY, Duke University, USA

SQL queries with group-by and average are frequently used and plotted as bar charts in several data analysis applications. Understanding the reasons behind the results in such an aggregate view may be a highly non-trivial and time-consuming task, especially for large datasets with multiple attributes. Hence, generating automated explanations for aggregate views can allow users to gain better insights into the results while saving time in data analysis. When providing explanations for such views, it is paramount to ensure that they are succinct yet comprehensive, reveal different types of insights that hold for different aggregate answers in the view, and, most importantly, they reflect reality and arm users to make informed data-driven decisions, i.e., the explanations do not only consider correlations but are *causal*. In this paper, we present CAUSUMX, a framework for generating summarized causal explanations for the entire aggregate view. Using background knowledge captured in a causal DAG, CAUSUMX finds the most effective causal treatments for different groups in the view. We formally define the framework and the optimization problem, study its complexity, and devise an efficient algorithm using the Apriori algorithm, LP rounding, and several optimizations. We experimentally show that our system generates useful summarized causal explanations compared to prior work and scales well for large high-dimensional data.

 $\label{eq:ccs} CCS \ Concepts: \bullet \ Information \ systems \rightarrow \ Database \ views; Relational \ database \ query \ languages; \ Data \ analytics.$

Additional Key Words and Phrases: Group-by SQL Queries, Query Results Explanation, Causal Inference.

ACM Reference Format:

Brit Youngmann, Michael Cafarella, Amir Gilad, and Sudeepa Roy. 2024. Summarized Causal Explanations For Aggregate Views. *Proc. ACM Manag. Data* 2, 1 (SIGMOD), Article 71 (February 2024), 27 pages. https://doi.org/10.1145/3639328

1 INTRODUCTION

As database interactions grow in popularity and their user base broadens to data analysts and decision-makers with varied backgrounds, it becomes important to generate insightful and automated explanations for results of the queries users run on the data. One simple yet important class of queries used in data analysis is the class of *SQL queries with group-by and average*, which are frequently used and plotted as barcharts in data analysis applications. These queries show how the average varies in different sub-populations in the data by creating an *aggregate view* over the input database (e.g., average salary per country, occupation, race, or gender; average severity of car accidents per major city in the USA, etc.). Understanding the *causal reasons* behind the high/low values of the average in different groups for such queries can enable sound data-driven

Authors' addresses: Brit Youngmann, brity@technion.ac.il, Technion - Israel Institute of Technology, Israel; Michael Cafarella, CSAIL MIT, USA, michjc@csail.mit.edu; Amir Gilad, Hebrew University, Israel, amirg@cs.huji.ac.il; Sudeepa Roy, Duke University, USA, sudeepa@cs.duke.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

https://doi.org/10.1145/3639328

ACM 2836-6573/2024/2-ART71

decision-making to address unwarranted situations. For instance, if a policymaker knows the possible causal reasons behind lower average salary of a certain race or gender in a certain region in the USA, they can try to improve the situation with corrective measures, which may not be possible with insights that are based on non-causal associational factors. Here we give a running example that we will frequently use in the paper.

EXAMPLE 1.1. Consider the Stack Overflow annual developer survey [1], where respondents from around the world answer questions about their job. We consider a subset of the data with 38090 tuples from 20 countries and 5 continents appearing the most in the dataset and augmented the data with additional attributes that describe the economy of each country: HDI (Human Development Index, higher values mean more human development), Gini (measures income inequity, higher values imply more inequity), GDP (Gross Domestic Product per capita, a measure of country's economic health, higher is better). Table 1 shows a few sample tuples with a subset of the attributes. The other attributes are SexualOrientation, EducationParents, Dependents, Student, Hobby, HoursComputer, and Exercise. Now consider the following group-by SQL query measuring the average salary in different countries:

SELECT Country, AVG(Salary) FROM Stack-Overflow GROUP BY Country

The results are plotted as a barchart in Figure 1 (the colors will be explained later). There is a huge variation in average salary (converted to USD) in different countries. The user may wonder (i) what are the main factors for this variation across countries, and also (ii) within each country, what is **causing** developers to earn more or less. However, the dataset is too big, both in terms of the number of tuples and attributes, to look for a succinct yet informative explanation by manual inspection. While tools like Tableau give highly sophisticated visualizations by slicing and dicing the data across several dimensions, they return the aggregates for these dimensions and do not differentiate between causal and non-causal reasons behind Figure 1.

Understanding the importance of generating insightful explanations for aggregated query results, several approaches have been proposed in database research on *explanations for aggregated query answers*. A simple form is given by the *provenance* for aggregate query answers that show how the output was computed using the input tuples [9]. However, an aggregate answer over a large dataset uses many input tuples, hence several approaches have focused on providing high-level explanations as *predicates* on input tuples that are responsible for producing query answers of interest [38, 56, 74] or provide other types of insights explaining them (e.g., the counterbalance approach in [45]). While these approaches provide predicates as explanations, which are easy to comprehend, they aim to explain certain answers in the view (e.g., outliers [74], high/low values of an answer [38, 45], or comparisons of a set of answers [56]), and do not provide a summarized explanation for the entire view. Further, although the explanations returned by these approaches reveal many interesting insights, they are not causal.

Causal inference, nevertheless, has been studied for several decades in Artificial Intelligence (AI) by *Pearl's Graphical Causal Model* [50], and in Statistics by *Rubin's Potential Outcome Framework* [58]. Causal analysis is a vital tool in determining the effect of a treatment on an outcome, and has been used in decision-making in medicine [53], economics [12], biology [66], and in critical applications like understanding the efficacy of a new vaccine using *randomized controlled trials*. While randomized trials cannot be performed in many applications due to ethical or feasibility issues, fortunately, the above causal models provide ways to do sound causal analysis on *observed* datasets under some assumptions (ref. Section 3).

ID	Country	Continent	Gender	Age	Role	Education	Major	Salary
1	US	N. America	Male	26	Data Scientist	PhD	C.S	180k
2	US	N. America	Non-binary	32	QA developer	B.Sc.	Mech. Eng.	83k
3	India	Asia	Male	29	C-suite executive	B.Sc.	C.S	24k
4	India	Asia	Female	25	Back-end developer	M.S.	Math.	7.5k
5	China	Asia	Male	21	Back-end developer	B.Sc.	C.S	19k

 Table 1. A subset of the Stack Overflow dataset.



Fig. 1. A visualization of the Stack Overflow query results.

Recent works have introduced causality to the field of database research [29, 42, 60, 61, 80], allowing users to benefit from this well-founded approach and infer solid causal conclusions from their data and queries. In particular, there has been prior work on extending Pearl's causal model for relational databases [61], providing causal hypothetical reasoning for what-if and how-to queries [29], and providing explanations for aggregate queries using causal analysis that focused on revealing unobserved factors influencing the results [60, 80]; however, [60, 80] provide a single explanation of the entire view, and do not offer fine-grained explanations for individual groups. A recent work [42] introduces a framework that searches for predicates that explain the difference in two average outcomes, and marks the patterns as either causal or not, by proposing a new causal discovery algorithm that extends the PC algorithm [67]. However, they do not search for important treatments affecting the outcomes and do not give causal explanations summarizing the entire view. On the other hand, *summarization* techniques for data and query answers form another active topic in database research [13, 17, 26, 35, 36, 62, 73, 77], often with *diversity* and *coverage* factors, however these summaries are also not causal.

Our contributions. In this work, we present a novel framework called CAUSUMX (Causal Summarized EXplanations) to explain the entire aggregate view from a query with group-by-average. Given a database D, causal background knowledge in the form of a causal DAG by Pearl's graphical causal model [50], a group-by-average query Q, and parameters k and θ , CAUSUMX generates a set of k explanation patterns (predicates) that explain at least θ fraction of groups in Q(D). An explanation pattern contains a grouping pattern capturing a subset of output groups covered by the explanation, and a treatment pattern with a high or low value of conditional average treatment effect (CATE) (ref. Section 3) on the average attribute A_{avg} as the outcome. In standard causal analysis, the goal is to estimate the causal effect of a given treatment on a given outcome, whereas in CAUSUMX we search for treatments with high and low causal effects for different subsets of groups defined by the grouping pattern. CAUSUMX combines the features of (i) causal inference, (ii) explanation, and (iii) summarization to provide succinct yet comprehensive and causal explanations for group-by-average queries, helping save time and effort in data analysis.

• For countries in Europe (blue bars with '/'), the most substantial effect on high salaries (effect size of 36K, p < 1e-3) is observed for individuals under 35 with a Master's degree. Conversely, being a student has the greatest adverse impact on annual income (effect size: -39K, p < 1e-3).

• For countries with a high GDP level (yellow bars with '*'), the most substantial effect on high salaries (effect size of 41K, p < 1e-3) is observed for <u>C-level executives</u>. Conversely, being over 55 with a bachelor's degree has the greatest adverse impact on annual income (effect size: -35K, p < 1e-4).

• For countries with a high Gini coefficient (pink bars with '-'), the most substantial effect on high salaries (effect size of 29K, p < 1e-4) is observed for white individuals under 45. Conversely, being having no formal degree has the greatest adverse impact on annual income (effect size: -28K, p < 1e-3).

Fig. 2. Causal explanation summary by CAUSUMX.

EXAMPLE 1.2. Reconsider the dataset and query from Example 1.1. The user runs CAUSUMX to search for an explanation for her query with no more than three insights while covering all groups, and receives the answers shown in Figure 2. The mapping between countries and insights is visualized using the bars' color and texture in Figure 1. Each country can be mapped to more than one insight, but for simplicity, only one color/texture is visualized. CAUSUMX uses a causal DAG (a partial DAG is shown in Figure 3), explores multiple patterns, and evaluates their causal effect on the salary across different countries. There are three parts in each insight: (a) A grouping pattern (first underlined text in magenta), illustrates a property or predicate on the groups or countries (group-by attribute in the query) for which this insight holds. (b) A positive treatment pattern (second underlined text in blue) is a predicate on the individuals from the above groups with a high positive treatment effect. (c) A negative treatment pattern (third underlined text in red) is a predicate on the individuals from the above groups with a high negative treatment effect.

Without having to manually explore this large dataset by running many subsequent queries, the user learns the main reasons for high and low salaries in different countries. These reasons are not just predicates summarizing tuples in the dataset, they have high and low causal effects as determined by the causal model. Moreover, the user knows that these explanations not only hold for one country but hold for several countries that share the same grouping pattern. The user can continue the exploration by varying parameters in CAUSUMX.

Our main contributions are as follows.

(1) We develop a framework called CAUSUMX that generates a summarized causal explanation to explain an aggregate view Q(D) for Q with group-by and average. We define explanation patterns that comprise a grouping pattern and a treatment pattern, define an optimization problem to maximize the causal explainability of these explanations subject to a size constraint on the number of explanations and a coverage constraint on the number of output groups covered by them, and show its NP-hardness.

(2) We design a three-step algorithm named CAUSUMX. The first step mines frequent grouping patterns using the seminal Apriori algorithm [7]. The second step uses a greedy lattice-based algorithm for mining promising treatment patterns for each grouping pattern from the previous step. In the third step, we model the optimization problem as an Integer Linear Program (ILP) and solve it by randomized rounding of its LP relaxation using the grouping and treatment patterns from previous steps.

(3) We provide a thorough **experimental analysis** and **multiple case study** that include five datasets, six baselines, and two variations of our solution as additional comparison points. We show that the explanations generated by CAUSUMX are of high quality compared to existing approaches and may provide different (and more justifiable) explanations from those given by associational and

Related Work			Entire	Supports
Kelated Work			View	Groups
	[15, 43-46, 55, 56, 74]	X	X	X
Ou any Basult Fundametian	[38]	×	×	\checkmark
Query Result Explanation	[42]	\checkmark	×	\checkmark
	[60, 80]	\checkmark	×	×
Interpretable Prediction Models	[19, 34]	×	1	×
Data Summanization	[26]	X	\checkmark	X
Data Summarization	[62, 73, 77]	×	✓	×
CAUSUMX			\checkmark	\checkmark

Table 2. Positioning of CAUSUMX w.r.t. Query Result Explanation, Interpretable Prediction Models, and Data Summarization.

causal approaches. Additionally, we analyze the runtime and accuracy of the proposed algorithms, and show that CAUSUMX is both efficient and useful in providing explanations.

2 RELATED WORK

Table 2 summarizes the differences between CAUSUMX and previous work. Columns in bold highlight our novelty, namely: CAUSUMX generates a **summarized** and **causal** explanation to the **entire aggregated view** generated by a SQL query while **accounting for variations among the groups.** In contrast, alternative approaches either provide non-causal explanations (as indicated in the Causal column), focus solely on elucidating specific portions of query results (as evident in the Entire View column), or offer a single explanation for query results, neglecting differences among groups (as represented in the Support Groups column).

Query Result Explanation. A substantial body of research has been dedicated to query result explanations. Multiple works used *data provenance* to obtain explanations for (possibly missing) query results [16, 18, 23, 37, 38, 43, 44, 46, 69]. Other forms of explanations include (non-causal) interventions [25, 55, 56, 68, 74], entropy [27], Shapley values [40, 52], and counterbalancing patterns [45]. Those works are orthogonal to our work, as we aim to explain an entire aggregated view via a small set of causal explanations (as explained above). Recent works [60, 79, 80] propose using causal inference to explain query results. In particular, [60] detects bias in a query in the form of confounder variables as explanations for Simpson's Paradox, while [80] finds confounders that explain the correlation between the grouping attribute and AVG queries. In both, the treatment is set as the grouping attribute and is fixed, and the same explanation is provided for all groups. Herein, we mine attribute-value pairs as treatments and search for ones with high causal effects for sets of groups. Another work [42] introduced a framework that identifies causal and non-causal patterns to explain the differences between two groups of tuples. In Section 6.2 we empirically demonstrate that this framework aims to solve a different task, and hence is unsuitable to solve the problem studied in this work.

Causal Inference. There is an extensive body of literature on causal inference over observational data in AI and Statistics [30, 50, 58, 70]. We employ standard techniques from this literature to compute causal effects. In related research, estimating heterogeneous treatment effects has been explored [72, 75]. This refers to variations in treatment effects across different population subgroups. However, this research differs from our framework. They assume known treatment and outcome variables and focus on identifying subpopulations with varying treatment effects. In contrast, we assume only the outcome variable is given and aim to identify treatments that influence the outcome for each subpopulation, potentially leading to different treatments for each subgroup.

Interpretable Prediction Models. Previous work developed models that offer both high predictive accuracy and interoperability [33, 41, 59, 63]. Rule-based interpretable prediction models [19, 34, 76] often utilize association rule mining processes to produce predictive rules. For instance, [34] introduced IDS, a model for binary classification that aims to optimize both the accuracy and interpretability of the chosen rules. This model generates a short, non-overlapping rule set encompassing the entire feature space and classes. Similarly, [19] devised FRL, an ordered rule list that includes probabilistic if-then rules. We compare against both [19, 34] in Section 6.

Data Summarization. Data summarization is the process of condensing an input dataset into interpretable and representative subsets [32, 73, 82]. A broad spectrum of approaches have been proposed for data and view summarization [13, 17, 35, 36, 62, 73, 77]. Unlike our work, none of these methods specifically aim to uncover causal explanations for aggregated views. In [26], the authors use *explanation tables*, which is one of the baselines in Section 6.

3 BACKGROUND ON CAUSAL INFERENCE

We use Pearl's model for *observational causal analysis* on collected datasets [50] and present the following concepts according to it.

Causal inference, Treatment, ATE, and CATE. The broad goal of *causal inference* is to estimate the effect of a *treatment variable T* on an outcome variable *Y* (e.g., what is the effect of higher Education on Salary). The gold standard of causal inference is by doing *randomized controlled experiments*, where the population is randomly divided into a *treated* group that receives the treatment (denoted by do(T = 1) for a binary treatment) and the *control* group (do(T = 0)). One popular measure of causal estimate is *Average Treatment Effect* (ATE). In a randomized experiment, ATE is the difference in the average outcomes of the treated and control groups [50, 58]

$$ATE(T, Y) = \mathbb{E}[Y \mid do(T = 1)] - \mathbb{E}[Y \mid do(T = 0)]$$
⁽¹⁾

The above definition assumes that the treatment assigned to one unit does not affect the outcome of another unit (called the Stable Unit Treatment Value Assumption (SUTVA)) [58]¹.

In our work on causal explanations for SQL group-by-average queries, where the treatment with maximum effect may vary among different tuples in the query answer, we are interested in computing the *Conditional Average Treatment Effect* (CATE), which measures the effect of a treatment on an outcome on *a subset of input units* [31, 57]. Given a subset of units defined by (a vector of) attributes *B* and their values *b*, we can compute CATE(T, Y | B = b) as:

$$\mathbb{E}[Y \mid do(T = 1), B = b] - \mathbb{E}[Y \mid do(T = 0), B = b]$$
(2)

However, randomized experiments where treatments are assigned at random cannot be done in many practical scenarios due to ethical or feasibility issues (e.g., effect of higher education on salary). In these scenarios, *Observational Causal Analysis* still allows sound causal inference under additional assumptions. Randomization in controlled trials mitigates the effect of *confounding factors* or *covariates*, i.e., attributes that can affect the treatment assignment and outcome. Suppose we want to understand the causal effect of Education on Salary from the SO dataset. We no longer apply Eq. (1) since the values of Education were not assigned at random in this data, and obtaining higher education largely depends on other attributes like Gender, EducationParents, and Country. Pearl's model provides ways to account for these confounding attributes *Z* to get an unbiased causal

¹This assumption does not hold for causal inference on multiple tables and even on a single table where tuples depend on each other, which we discuss in Section 7.

estimate from observational data under the following assumptions (*L* denotes independence):

Unconfoundedness:
$$Y \perp T | Z = z$$
 (3)

Overlap:
$$0 < Pr(T=1|Z=z) < 1$$
 (4)

The unconfoundedness assumption, Eq. (3), states that if we condition on Z, then treatment T in the dataset and the outcome Y are independent. In SO, assuming that only Z ={Gender, EducationParents, Country} affects T = Education, if we condition on a fixed set of values of Z, i.e., consider people of a given gender, from a given country, and with a given education level of parents, then T = Education and Y = Salary are independent. For such confounding factors Z, Eq. (1) and (2) respectively reduce to the following form (Eq. (4) gives feasibility of the expectation difference):

$$ATE(T, Y) = \mathbb{E}_{Z} \left[\mathbb{E}[Y \mid T = 1, Z = z] - \mathbb{E}[Y \mid T = 0, Z = z] \right]$$

$$CATE(T, Y \mid B = b) = \mathbb{E}_{Z} \left[\mathbb{E}[Y \mid T = 1, B = b, Z = z] - \mathbb{E}[Y \mid T = 0, B = b, Z = z] \right]$$
(5)

The above equations no longer have the do(T = b), and can be estimated from an observed dataset. Pearl's model gives a systematic way to find such a *Z* when a causal DAG is available.

Causal DAG. Pearl's Probabilistic Graphical Causal Model model [50] can be written as a tuple $(\mathcal{E}, \mathcal{N}, Pr_{\mathcal{E}}, \psi)$, where \mathcal{E} is a set of *unobserved exogenous (noise)* variables, $Pr_{\mathcal{E}}$ is the joint distribution of \mathcal{E} , and \mathcal{N} is a set of *observed endogenous variables*. Here ψ is a set of structural equations that encode dependencies among variables. The equation for $A \in \mathcal{N}$ takes the following form:

$$\psi_A : \operatorname{dom}(Pa_{\mathcal{E}}(A)) \times \operatorname{dom}(Pa_{\mathcal{N}}(A)) \to \operatorname{dom}(A)$$

Here $Pa_{\mathcal{E}}(A) \subseteq \mathcal{E}$ and $Pa_{\mathcal{N}}(A) \subseteq \mathcal{N} \setminus \{A\}$ respectively denote the exogenous and endogenous parents of *A*. A causal relational model is associated with a *causal DAG*, *G*, whose nodes are the endogenous variables \mathcal{N} and whose edges are all pairs (X, Y) (directed edges from *X* to *Y*) such that $Y \in \mathcal{N}$ and $X \in Pa_{\mathcal{N}}(Y)$. The causal DAG obfuscates exogenous variables as they are unobserved. Any given set of values for the exogenous variables completely determine the values of the endogenous variables by the structural equations (we do not need any known closed-form expressions of the structural equations in this work). The probability distribution $Pr_{\mathcal{E}}$ on exogenous variables \mathcal{E} induces a probability distribution on the endogenous variables \mathcal{N} by the structural equations ψ .

Figure 3 depicts a causal DAG for the SO dataset over the attributes in Table 1 as endogenous variables (we use a larger causal DAG with all 20 attributes for the SO dataset in our experiments). Given this causal DAG, we can observe that the Role that a coder has in their company depends on the values of their Education, Age, Major, and YearsCoding attributes.

A causal DAG can be constructed by a domain expert as in the above example, or using existing *causal discovery* [50] algorithms, which we study further in our experiments (Section 6).

In Pearl's model, a treatment T = t (on one or more variables) is considered as an *intervention* to a causal DAG by mechanically changing the DAG such that the values of node(s) for T in G are set to the value(s) in t, which is denoted by do(T = t). Following this operation, the probability distribution of the nodes in the graph changes as the treatment nodes no longer depend on the values of their parents. Pearl's model gives an approach to estimate the new probability distribution by identifying the confounding factors Z described earlier using conditions such as *d*-separation and *backdoor criteria* [50], which we do not discuss in this paper.



Fig. 3. Example causal DAG.

4 FRAMEWORK FOR SUMMARIZED CAUSAL EXPLANATIONS FOR AGGREGATE QUERIES

Databases and Queries. We consider a single-relation database² over a schema A. The schema is a vector of attribute names, i.e., $\mathbb{A} = (A_1, \ldots, A_s)$, where each A_i is associated with a domain dom (A_i) , which can be categorical or continuous. A database instance D, populates the schema with a set of tuples $t=(a_1, \ldots, a_s)$ where $a_i \in \text{dom}(A_i)$. We use $t[A_i]$ to denote the value of attribute A_i of tuple t. In this paper, we will only consider the *active domain* of every A_i as dom (A_i) , i.e., the set of values of A_i in the given D.

We consider an important class of SQL queries for data analysis, with group-by and average as the aggregate function:

 $Q = \text{SELECT } \mathcal{A}_{qb}, \text{ AVG } (A_{avq}) \text{ FROM D WHERE } \phi \text{ GROUP BY } \mathcal{A}_{qb};$

Here, $\mathcal{A}_{gb} \subseteq \mathbb{A}$ is a set of categorical group-by attributes, $A_{avg} \in \mathbb{A}$ is the average attribute, and ϕ is a predicate. The result of evaluating Q over D is denoted by Q(D). We denote |Q(D)|, i.e., the number of groups in Q(D), by m. The WHERE condition ϕ simply reduces the table D to tuples satisfying ϕ before the techniques in this section can be applied, so we do not discuss ϕ further in this section.

Consider again the query presented in Example 1.1. Here \mathcal{A}_{gb} = Country, A_{avg} = Salary, and ϕ is the empty predicate. The query results Q(D) are shown in Figure 1, where m = |Q(D)| = 20.

4.1 Explanation Patterns

The *patterns* in our framework of summarized explanations are *conjunctive predicates* on attribute values that are prevalent in previous work on explanation, e.g., [27, 56, 74].

DEFINITION 4.1 (PATTERN). Given a database instance D with schema \mathbb{A} , a simple predicate is an expression of the form $\varphi = A_i$ op a_i , where $A_i \in \mathbb{A}$, $a_i \in \text{dom}(A_i)$, and $\text{op} \in \{=, <, >, \le, \ge\}$. A pattern is a conjunction of simple predicates $\mathcal{P} = \varphi_1 \land \ldots \land \varphi_k$.

Grouping and treatment patterns. Our explanations consist of pairs of patterns: (i) A **grouping pattern** \mathcal{P}_g captures a subset of groups in Q(D) and must be well-defined over Q(D), i.e., each query answer $s \in Q(D)$ is either *covered* or not by \mathcal{P}_g . Therefore, \mathcal{P}_g can only contain attributes W such that the Functional Dependency (FD) from the grouping attributes, $\mathcal{A}_{gb} \to W$ holds for all W in \mathcal{P}_g . (ii) A **treatment pattern**, \mathcal{P}_t , is defined over the dataset D (as opposed to \mathcal{P}_g that is defined over Q(D)) and partitions the input tuples into treated (T = 1 if \mathcal{P}_t evaluates to true for a tuple) and control groups (T = 0 if \mathcal{P}_t evaluates to false). This partition is then used to assess the causal effects of the treatment pattern on the outcome $Y = A_{avg}$, the attribute for average in the query Q.

A pair of grouping and treatment pattern ($\mathcal{P}_g, \mathcal{P}_t$) together define an **explanation pattern**. Intuitively, \mathcal{P}_g specifies the subpopulation of interest (equivalent to the condition B = b in Eq.

Proc. ACM Manag. Data, Vol. 2, No. 1 (SIGMOD), Article 71. Publication date: February 2024.

²We discuss adjustments for supporting multi-dimensional datasets in Section 7.

(2) for CATE), while \mathcal{P}_t (equivalent to treatment *T*) explains the observed outcome *Y* within that subpopulation as per the CATE value.

DEFINITION 4.2 (EXPLANATION PATTERN). Given a database instance D with schema \mathbb{A} and a query Q with group-by attributes \mathcal{A}_{gb} and attribute A_{avg} for average, an explanation pattern $(\mathcal{P}_g, \mathcal{P}_t)$ where \mathcal{P}_g is a grouping pattern on Q(D), i.e., the FD $\mathcal{A}_{gb} \to W$ holds in D for all attributes W in \mathcal{P}_g , and \mathcal{P}_t is a pattern defined over D.

EXAMPLE 4.1. In the first insight in Figure 2, one explanation pattern is $(\mathcal{P}_g, \mathcal{P}_t)$ with \mathcal{P}_g : (Continent = Europe) and \mathcal{P}_t : (Age < 35) \wedge (Education = Master's degree). Note that the FD from the group-by attribute Country \rightarrow Continent holds.

Partitioning attributes for grouping and treatment patterns. We partition the attributes in \mathbb{A} into two disjoint sets. All attributes $W \subseteq \mathbb{A}$ s.t the FD $\mathcal{A}_{gb} \to W$ holds in D are considered for grouping patterns. All other attributes U are considered for treatment patterns. The necessity for FD for grouping patterns is explained above. Further, any attribute W where $\mathcal{A}_{gb} \to W$ holds, and in general, any grouping pattern, cannot be a valid treatment pattern. By the overlap condition in Eq. (4), conditioned on a grouping pattern \mathcal{P}_g , there should be at least one unit with T=1 and at least one unit with T=0. If we pick a pattern W=w with $\mathcal{A}_{gb} \to W$ or a pattern with multiple such attributes as the treatment, for all tuples in D contributing to a query answer in Q(D), either it evaluates to true or to false, so we do not get both T=1 and T=0 to estimate the CATE value. **Explainability.** To evaluate the effectiveness of an explanation pattern ($\mathcal{P}_g, \mathcal{P}_t$), we define its explainability.

DEFINITION 4.3 (EXPLAINABILITY). Given a database instance D with schema \mathbb{A} , a query Q groupby attributes \mathcal{A}_{gb} and attribute A_{avg} , and a causal model $\mathcal{M}_{\mathbb{A}}$ on \mathbb{A} associated with a causal DAG, the explainability of an explanation pattern ($\mathcal{P}_q, \mathcal{P}_t$) is defined as:

$$Explainability(\mathcal{P}_q, \mathcal{P}_t) := CATE_{\mathcal{M}_{\mathbb{A}}}(\mathcal{P}_t, A_{avg} \mid \mathcal{P}_q)$$

using the definition of CATE as in Eq. (2) with treatment $T=\mathcal{P}_t$, outcome $Y=A_{avg}$, and subpopultion defined by \mathcal{P}_g . The subscript $\mathcal{M}_{\mathbb{A}}$ denotes that the CATE is estimated using the causal model $\mathcal{M}_{\mathbb{A}}$ as explained in Section 3.

In particular, CATE given by Eq. (2) in the above definition is reduced to Eq. (5) using confounding variables Z obtained from the causal DAG of $\mathcal{M}_{\mathbb{A}}$, which then can be estimated from the data D.

Our focus is on computing CATE (Eq. 2) rather than ATE (Eq. 1) to understand the causal factors that influence outcomes for groups in Q(D). For instance, the effect of having a Master's degree on Salary in countries in Europe can be different from that in countries in Asia. Therefore, computing the ATE while considering all individuals from across the globe may not provide meaningful insights. For \mathcal{P}_g : (Continent = Europe) and \mathcal{P}_t : (Education = MA degree), we define the treatment group as individuals with a master's degree from European countries and the control group as individuals without a Masters degree from European countries. This enables us to draw relevant conclusions.

EXAMPLE 4.2. In Example 1.2 and Figure 2, there are two explanation patterns using the same grouping pattern. Here \mathcal{P}_g : (Continent = Europe), \mathcal{P}_{t1} : (Age < 35) \wedge (Education = Master's degree), and \mathcal{P}_{t2} : (Student = yes). The explainability of ($\mathcal{P}_g, \mathcal{P}_{t1}$) = 36K and that of ($\mathcal{P}_g, \mathcal{P}_{t2}$) = -39K, indicating that age below 35 and having a Master's degree has a high positive causal effect for individuals from European countries while being a student has a high negative effect.

4.2 Problem Definition and Hardness

Our goal is to obtain a succinct yet comprehensive set of explanation patterns for the groups in Q(D) from the huge search space of possible explanation patterns (e.g., the number of explanation patterns for Example 1.1 was 22350). To achieve this, we frame a constrained optimization problem. We apply three constraints: (1) the number of explanation patterns should not exceed a specified threshold, (2) the number of groups in Q(D) explained by the patterns must be at least a specified θ -fraction of all the groups in Q(D), and (3) an explanation pattern should not explain the same set of groups explained by another explanation pattern. Finally, our goal is to find the set of explanation patterns that abide by these constraints and whose overall explainability is maximized. First, we define the *coverage* of a grouping pattern \mathcal{P}_q .

DEFINITION 4.4 (COVERAGE). Given a database instance D and a query Q with group-by attributes \mathcal{A}_{gb} , a grouping pattern \mathcal{P}_g , and a group $s \in Q(D)$, \mathcal{P}_g is said to **cover** s if for any tuple $t \in D$ such that $t[\mathcal{A}_{gb}] = s[\mathcal{A}_{gb}]$, it holds that $t \models \mathcal{P}_g$, i.e., t satisfies the predicate \mathcal{P}_g . The set of groups in Q(D) covered by \mathcal{P}_g is denoted by $Cov(\mathcal{P}_g)$.

Next, we define the problem of Summarized Causal Explanations.

DEFINITION 4.5 (SUMMARIZED CAUSAL EXPLANATIONS). Given a database D, a causal model $\mathcal{M}_{\mathbb{A}}$, a query Q with group-by attributes \mathcal{A}_{gb} and attribute A_{avg} for average, a collection of explanation patterns $\{\mathcal{P}_i\}_{i=1}^l$, an integer $k \in [1, m]$ where m = |Q(D)|, and a threshold $\theta \in [0, 1]$, we aim to find a set $\Phi \subseteq \{\mathcal{P}_i\}_{i=1}^l$ of explanation patterns such that the following conditions hold:

- (Size constraint) $|\Phi| \leq k$.
- (Coverage constraint) at least θ -m groups from Q(D) are covered by Φ , i.e., $\bigcup_{\mathcal{P}_g \in \Phi} Cov(\mathcal{P}_g) \ge \theta \times m$.
- (Incomparability constraint) There are no pairs of explanation patterns, $(\mathcal{P}_g, \mathcal{P}_t)$ and $(\mathcal{P}'_g, \mathcal{P}'_t)$ in Φ such that $Cov(\mathcal{P}_g) = Cov(\mathcal{P}'_g)$.

(Objective) The objective is to maximize the total explainability of Φ under the above constraints, i.e., maximize $\sum_{\mathcal{P} \in \Phi} Explainability(\mathcal{P})$.

The size and incomparability constraints ensure that the size of the explanation will be small with no redundancy and therefore it will be more easily grasped by the user. The coverage constraint ensures that the explanation will be extensive and capture a significant part of the view, and the objective aims to maximize the validity of the explanation as the cause of the trends in the view.

EXAMPLE 4.3. In Example 1.2 and Figure 2, the parameters k=3 and $\theta=1$ are used, i.e., we aim to find a set of at most 3 explanation patterns that reveal the causes of the outcome for all groups in Q(D). As Proposition 4.1 shows, even deciding whether there is a set Φ of explanation patterns for a given size constraint k and coverage constraint θ simultaneously is NP-Hard, although for this example, we find a solution covering all 20 countries in Q(D) as shown in Figure 1.

Positive and negative explanation patterns. Given an outcome variable (e.g., income), positive explanations correspond to treatments that have an impact on making its value higher (what increases income), and negative explanations are treatments that make its value lower (what reduces income). These positive/negative treatments can vary for different groups (as illustrated in Figure 2). In an application, both or one of them may be valuable. Consequently, our proposed framework supports both. In a prototype with a UI, analysts have the flexibility to choose whether they want to view one or both types of explanations and even top-k positive/negative treatments for a grouping pattern. This helps understand the cause of both high and low values of outcomes for different groups without explicitly asking for explanations for high and low values as done in previous work [38, 45, 56, 74].

Summarized Causal Explanations For Aggregate Views

71:11

To generate positive and negative patterns for a grouping pattern, we slightly vary the optimization objective in Definition 4.5. For a grouping pattern \mathcal{P}_g , we find a treatment pattern $\mathcal{P}_{g,tp}^+$ with the highest explainability value of $(\mathcal{P}_g, \mathcal{P}_{g,t}^+)$ $(\mathcal{P}_{g,t}^+$ is called a *positive treatment pattern* for \mathcal{P}_g). A negative treatment pattern $\mathcal{P}_{g,t}^-$ is defined similarly using the lowest explainability value. In our system, for each grouping pattern \mathcal{P}_g we compute the sum of absolute values of two explainabilities: $|Explainability(\mathcal{P}_g, \mathcal{P}_{g,t}^+)| + |Explainability(\mathcal{P}_g, \mathcal{P}_{g,t}^-)|$. We treat this sum as the weight of the explanation pattern combination $(\mathcal{P}_g, \mathcal{P}_{g,t}^+, \mathcal{P}_{g,t}^-)$ and return top-k explanations that satisfy the constraints in Definition 4.5.

Hardness Result and enumeration of search space. Since in our optimization problem, we want to cover a certain fraction of answer tuples in Q(D) ('elements') with at most k patterns ('sets'), we can show the following NP-hardness result even if we ignore the optimization objective (proof in the full version [78]).

PROPOSITION 4.1. It is NP-hard to decide whether the Summarized Causal Explanations problem is feasible (i.e., has any solution satisfying the constraints) for a given k and θ .

Definition 4.5 assumes that the search space of explanation patterns is given, while in practice it is not efficient to enumerate all explanation patterns ahead of time. In Section 5 we give an efficient algorithm that give good solutions for the optimization problem without explicitly enumerating all explanation patterns upfront.

5 THE CAUSUMX ALGORITHM

Proposition 4.1 shows that even deciding the feasibility of the Summarized Causal Explanations problem for a given k and θ is NP-hard. Further, it is not practical to enumerate all explanation patterns and compute their explainability upfront. Given four-dimensional desiderata in Definition 4.5 (unlike the standard set-cover or max-cover problems that have two), it is non-trivial to design a good approximation algorithm or heuristics for this problem. In this section, we present the CAUSUMX algorithm (for <u>Causal Summarized EXplanations</u>) that aims to address these challenges.

A brute-force approach considers all grouping and treatment patterns and results in long runtimes (as we demonstrate in Section 6.5). Instead, CAUSUMX employs the Apriori algorithm [7] to mine frequent grouping patterns (the coverage of such patterns is monotone) as a heuristic for finding promising grouping patterns that are short (we ignore the rest). To mine treatment patterns (which are non-monotone for CATE) we adopt a lattice traversal approach and use a greedy heuristic materializing only promising treatment patterns. Both steps can lose optimality. In Section Section 6 we compare CAUSUMX against Brute-Force and show that while CAUSUMX is much faster, the difference in results (in terms of explainability, coverage, and accuracy) is modest.

Overview. The pseudo-code in Algorithm 1 outlines the operation of CAUSUMX. It takes a database D, an integer k, a query Q with an outcome A_{avg} and grouping attributes \mathcal{A}_{gb} , and a coverage threshold θ as input. The output is a set of explanation patterns Φ . The algorithm consists of three steps: (1) extracting candidate grouping patterns (line 2). (2) Focusing on promising treatment patterns for each extracted grouping pattern, *materializing and evaluating only them* (lines 3–4). (3) Utilizing Linear Programming (LP) to obtain a set of explanation patterns (line 5).

5.1 Mining Grouping Patterns

Considering every possible grouping pattern (Definition 4.2) is infeasible as their number is exponential $(O(agrmax_{A_i \in \mathbb{A}} | dom(A_i) |^{|\mathbb{A}|})$. Instead, our approach utilizes the Apriori algorithm [7] to generate candidate patterns. The Apriori algorithm gets a threshold τ , and ensures that the mined patterns are present in at least τ tuples of D. Formally, given a set of attribute $W \subseteq \mathbb{A}$ s.t the FD

Algorithm 1: The CAUSUMX algorithm	
------------------------------------	--

	nput : A database relation D , an integer k , a query Q with an aggregate attribute A_{avg} and grouping attributes \mathcal{A}_{gb} , and a coverage threshold θ output : A set Φ of explanation patterns.		
1	$\Phi \leftarrow \emptyset;$		
2	$Candidates \leftarrow GetGroupingPatterns(D, A_{avg}, \mathcal{A}_{gb});$	// Section 5.1	
3	for $\mathcal{P}_q \in Candidates$ do		
4	$ \bar{\mathcal{P}}_{g}.\mathcal{P}_{t} \leftarrow \texttt{GetTopTreatment}(\mathcal{P}_{g},k,A_{avg},D); $	// Section 5.2	
5	$\Phi \leftarrow SolveLP(Candidates, k, \theta);$	// Section 5.3	
6	return Φ		

 $\mathcal{A}_{gb} \rightarrow W$ holds in *D*, we apply the Apriori to extract frequent patterns defined solely by these attributes. The algorithm guarantees that each mined pattern covers at least τ tuples from *D* and is well-defined over Q(D), making them promising candidates for covering the necessary number of groups (see item (2) in Definition 4.5).

Post-Processing. Certain extracted grouping patterns may be superfluous, i.e., the set of groups they define from Q(D) could be indistinguishable. Redundant grouping patterns can emerge even in the absence of FDs among attributes. To illustrate, consider the patterns {HDI = High} and GDP = High, which both identify the same set of countries. However, it is possible that two countries with medium HDI values exhibit different GDP values. Thus, there is no functional dependency between HDI and GDP. Nevertheless, two grouping patterns defined by these attributes can still define the same set of tuples, rendering the consideration of both redundant. Therefore, following the mining stage, we remove redundant grouping patterns to ensure the obtained solution will satisfy the incomparability constraint (item (3) in Definition 4.5). Additionally, we favor more succinct patterns as they are easier to comprehend. To achieve this, we utilize a hash table that records the groups from Q(D) associated with each mined grouping pattern. In each group set, we retain only the shortest grouping pattern.

5.2 Mining Treatment Patterns

As opposed to a standard causal analysis setting where a causal question of the form "What is the effect of treatment T on outcome Y?" is posed, in our setting, we aim to find T that yields the highest effect on Y. Our subsequent goal, as discussed at the end of Section 4.2, is to identify a positive treatment pattern \mathcal{P}_{tg}^+ and a negative treatment pattern \mathcal{P}_{tg}^- for each mined grouping pattern \mathcal{P}_{g} . Which respectively explain the cause for high and low outcomes of tuples in D that satisfy \mathcal{P}_{g} . However, for simplicity and without loss of generality, we present the approach for finding positive explanation patterns only denoted by \mathcal{P}_{t} (i.e., treatments that have the highest positive CATE for a grouping pattern \mathcal{P}_{g}).

Since the number of potential treatment patterns for \mathcal{P}_g can be large (exponential in $|\mathbb{A}|$), we propose a greedy heuristic approach to materialize and assess the CATE only for promising treatment patterns. This is done by leveraging the notion of lattice traversal [10, 24]. In particular, the set of all treatment patterns can be represented as a lattice where nodes correspond to patterns and there is an edge between \mathcal{P}_t^1 and \mathcal{P}_t^2 if \mathcal{P}_t^2 can be obtained from \mathcal{P}_t^1 by adding a single predicate. This lattice can be traversed in a top-down fashion while generating each node at most once. Since not all nodes correspond to treatments that have a positive CATE, we only materialize nodes if all their parents have a positive CATE. The primary distinction from existing solutions (e.g., [10]) lies in the non-monotonic nature of CATE (meaning that adding a predicate to a treatment pattern can either increase or decrease its CATE value). Consequently, exclusively materializing nodes where all of their parents exhibit a positive CATE might lead the algorithm to overlook certain relevant treatment patterns. For example, we observe that for the pattern $\mathcal{P}_{t1}=(role = QA)$, by adding

Algorithm 2: Top treatment pattern for a grouping pattern				
input : A grouping pattern \mathcal{P}_g , the outcome attribute A_{avg} , the dataset D , a causal DAG G , and direction				
$\sigma \in \{+, -\}.$				
output : A treatment pattern \mathcal{P}_t .				
1 /* Get all single-predicate patterns. *	/			
2 $C \leftarrow \text{GenChildren}(D, A_{avg}, \mathcal{P}_g);$				
3 $C \leftarrow \text{ComputeCATEnFilter}(C, D, A_{avg}, \mathcal{P}_g, G);$				
4 $\mathcal{P}_t^{max} \leftarrow GetTopTreatment(C);$				
5 while True do				
6 /* Get patterns in the next level. *	/			
$\tau C \leftarrow \text{GenChildrenNextLevel}(C, \sigma);$				
8 $C \leftarrow \text{ComputeCATEnFilter}(C, D, A_{avg}, \mathcal{P}_g, G);$				
9 $\mathcal{P}_t \leftarrow \text{GetTopTreatment}(C);$				
10 if $\mathcal{P}_t.CATE > \mathcal{P}_t^{max}.CATE$ then				
11 $\left\lfloor \mathcal{P}_t^{max} \leftarrow \mathcal{P}_t \right\rangle$				
12 else				
13 Break;				

```
14 return \mathcal{P}_t^{max};
```



Fig. 4. Partial treatment-patterns lattice (Example 5.1).

the predicate (education = MA), its CATE value increases, while the CATE decreases by adding the predicate (education = no degree). However, according to our experiments (Section 6.2), combining treatment patterns that exhibit a positive CATE is highly likely to result in a treatment with a positive CATE as well. As a result, we account for most treatments with a positive CATE, and the accuracy of this algorithm is relatively high.

Algorithm 2 describes the search for the treatment pattern \mathcal{P}_t^{max} with the highest positive CATE value for a given grouping pattern \mathcal{P}_g . It traverses the pattern lattice in a top-down manner, starting from patterns with a single atomic predicate in line 2. In particular, the function GenChildren generates all atomic predicates of the form A_i op a_j . For each treatment pattern, it evaluates its CATE, and discards the ones with CATE values that do not have the same sign σ , either + or – (line 3). The algorithm stores the pattern with the highest positive CATE identified thus far (line 4). It then proceeds to traverse the pattern lattice only for nodes that satisfy the condition of having all their parents with a positive CATE (line 7-8), and updates \mathcal{P}_t^{max} if a pattern with a higher CATE was found (lines 9-12). It terminates at the first level, which does not include the maximum value recorded (lines 13-14).

EXAMPLE 5.1. We illustrate the operation of Algorithm 2 using Figure 4. Initially, it considers all treatment patterns with a single predicate (line 2). Assuming σ =+ (searching for the treatment pattern with the highest CATE), it moves to the next level (line 5). At this level, patterns with two predicates are considered only if both of their parents have a positive CATE. Thus, the pattern {Role = QA, Gender = Male} is excluded as the CATE of Role = QA is negative (line 8). The treatment pattern with the maximum CATE (\mathcal{P}_t^{max}) is found at the second level (marked in yellow) (line 9). After another iteration and generating the third level of the lattice, the algorithm does not proceed to the fourth level because \mathcal{P}_t^{max} is not present in the third level (line 10).

By changing σ , Algorithm 2 can also find the treatment patterns with the lowest negative CATE. It can also be generalized to output the highest explainability (i.e., the highest absolute CATE value), or both the patterns with the highest and lowest CATE values.

Optimizations. We implement more optimizations for Algorithm 2.

(a) **Pruning attributes:** We eliminate attributes that do not have a causal relationship with the outcome attribute (lines 2, 7). Since these attributes have no impact on CATE values, they can be disregarded. We can detect such attributes by utilizing the input causal DAG or by removing attributes with low correlation to the outcome.

(b) Pruning treatments: In constructing the lattice (lines 2, 7), we exclude patterns with a nearzero CATE. We observe that combining patterns with a CATE close to zero value often yields a similar result. Consequently, when advancing to the next lattice level, we only consider the top 50% patterns with the highest or lowest CATE.

(c) Parallelism: The process of extracting treatment patterns for each grouping pattern (lines 2, 7) can be performed in parallel since this procedure is dependent only on the grouping pattern.

(d) Estimating CATE Values by Sampling: To decrease runtime, we use a fixed-size random sample of tuples to estimate CATE values (lines 3, 8) and investigate the impact of sample size on runtime and accuracy (Section 6.6). Random sampling of tuples is done only for evaluating CATE values. Our study shows that using a random sample size of 1 million tuples yields CATE estimates that closely match those obtained from the entire dataset.

5.3 Linear Program Formulation and Rounding

Proposition 4.1 shows that finding any feasible solution to the optimization problem in Definition 4.5 is NP-hard. Our study shows that intuitive combinatorial greedy algorithms targeting the size and coverage constraints and maximizing explainability are unable to find good solutions; hence, we use an LP-rounding algorithm.

Given a collection of explanation patterns $\{\mathcal{P}_j\}_{j=1}^l$ with weights w_j corresponding to their explainability, an integer k, and a threshold θ , we construct the Integer Linear Program (ILP) shown in Figure 5 (which extends the ILP for the *max-k-cover* problem). Here g_j are the variables for patterns \mathcal{P}_j . t_i , i = 1 to m are the variables for m groups in Q(D). $s_i \in \mathcal{P}_j$ denotes that s_i satisfies \mathcal{P}_j .

We consider the LP-relaxation of this ILP with variables in [0, 1] instead of $\{0, 1\}$, then use an LP-solver to find solutions. If no solution is returned, we know that the original ILP had no feasible solution either. If any solution is returned by the LP, the original ILP may or may not have a feasible solution. We use the standard randomized rounding algorithm for max-k-cover [51] (sample *k* patterns with probability $\frac{g_j}{k}$), which guarantees at most *k* sets, an $(1 - \frac{1}{e})$ -approximation to the coverage constraint $\sum_{i=1}^{m} t_i \ge \theta \cdot m$, and a $\frac{1}{k}$ -approximation to the maximization objective $\sum_{j=1}^{l} g_j \cdot w_j$ in the LP as well as the ILP since *OPT LP* $\ge OPT$ *ILP* (details in the full version [78]). However, the approximation guarantees of the proposed LP formulation are solely theoretical since the approximation to the objective holds when *all explanation patterns are considered*. In practice, we use this LP-rounding algorithm in conjunction with the grouping and treatment pattern mining procedures described in Sections 5.1 and 5.2, therefore incur a trade-off between value and efficiency and lose the theoretical guarantees.

$$\max \sum_{j=1}^{l} g_j \cdot w_j \quad \text{s.t.} \quad (1) \sum_{j=1}^{l} g_j \le k, \quad (2) \ t_i \le \sum_{j:T_i \in \mathcal{P}_j} g_j \ \forall i = 1 \text{ to } m,$$
$$(3) \sum_{i=1}^{m} t_i \ge \theta \cdot m, \quad (4) \ t_i, g_j \in \{0, 1\} \ \forall i = 1 \text{ to } m, \ \forall j = 1 \text{ to } l$$

Fig. 5. ILP for optimization problem (line 5 in Algorithm 1).

Dataset	tuples	atts	max values per att	grouping patterns	
German [11]	1000	20	53	10	
Adult [2]	32.5K	13	94	13	
SO [1]	38K	20	20	75	
IMPUS-CPS [28]	1.1M	10	67	9	
Accidents [47]	2.8M	40	127	15	

Time complexity analysis for Algorithm 1. The maximum number of explanation patterns in a database D with attributes \mathbb{A} is bounded by $|D|^{|\mathbb{A}|}$ (considering both grouping and treatment patterns and active domain of attributes), which is polynomial in *data complexity* assuming a fixed schema [71]. The number of patterns dominates the number of variables in the ILP in Figure 5. Hence even if all patterns are enumerated and evaluated explicitly, the LP relaxation of the ILP can be solved (CAUSUMX uses z3 [22]) and rounded to an integral solution in polynomial time in |D|. The additional operations in this section (e.g., computation of CATE in Algorithm 2) are polynomial in D, giving a worst-case polynomial data complexity of CAUSUMX. However, $|D|^{|\mathbb{A}|}$ can have a large value for large |D| and \mathbb{A} . The suite of optimizations developed in this section reduces the running time of CAUSUMX (ref. Section 6).

6 EXPERIMENTAL EVALUATION

We present experiments that evaluate the effectiveness and efficiency of our proposed framework. We aim to address the following research questions. Q1: What is the quality of our explanations, and how does it compare to that of existing methods? Q2 : How does each phase of CAUSUMX contributes to its ability to find an explanation summary that satisfies our optimization goal and constraints? Q3 : What is the efficiency of the CAUSUMX algorithm? Q4: How sensitive are the explanations to various parameters?

Prototype implementation. CAUSUMX was written in Python, and is publicly available in [78]. CATE values computation was performed using the DoWhy library [64], utilizing their linear regression approach. CAUSUMX generates a solution in natural language using predefined templates, as shown in Figure 2. Those templates were generated via prompt questions to ChatGPT [4], asking it to transform predicates into a human-readable text.

6.1 Experimental Setup

Datasets. We examine multiple commonly used datasets:

German: This dataset contains details of bank account holders, including demographic and financial information, along with their credit risk. The causal graph was used from [21].

Adult: This dataset comprises demographic information of individuals along with their education, occupation, annual income, etc. We used the causal graph from [21].

SO: This is discussed in Example 1.1. The causal DAG was constructed using [81].

IMPUS-CPS: This dataset is derived from the Current Population Survey conducted by the U.S. Census Bureau, which includes demographic details for individuals, e.g., education, occupation, and annual income. We adopted the causal dag from [21].

Accidents: This dataset provides comprehensive coverage of car accidents across the USA. It includes numerous environmental stimuli features that describe the conditions surrounding the accidents, such as visibility, precipitation, and traffic signals. To construct a causal DAG, we followed the methodology outlined in [81].

Synthetic: We constructed a schema comprising the attributes $G, G_1, \ldots, G_i, T_1, \ldots, T_j, O$, where: G serves as the grouping attribute in the query, with each tuple taking a unique value ranging from 1 to n (where n is the number of tuples). Attributes G_1, \ldots, G_i are used to establish grouping patterns. Each attribute divides the values of G into varying numbers of buckets. Attributes T_1, \ldots, T_j are employed to define treatment patterns. Each tuple is assigned a random value between 1 and 5 for each attribute independently. The outcome O is defined as: $T_1 - T_2 + T_3 - \ldots + T_j$. There is a large number of grouping and treatment patterns to be considered. For each grouping pattern, the treatment patterns that exhibit the highest causal effects encompass all the attributes T_1, T_2, \ldots, T_j . For instance, treatment patterns with high positive causal effect on O are characterized by high values for odd T attributes and low values for even T attributes (e.g., $T_1 = 5, T_2 = 1, T_3 = 5, \ldots$).

Baselines. We compare CAUSUMX with the following baselines:

Brute-Force: The optimal solution according to Definition 4.5. This algorithm implements an exhaustive search over all possible grouping and treatment pattern combinations.

IDS: The authors of [34] have proposed a framework for generating Interpretable Decision Sets (IDS) for prediction tasks. The framework incorporates parameters restricting the percentage of uncovered data tuples and the number of rules. These parameters were assigned the same values used in our system.

FRL: The authors of [19] introduce a framework for producing Falling Rule Lists (FRL) as a probabilistic classification model. FRLs consist of a sequence of if-then rules, with the if-clauses containing antecedents and the then-clauses containing probabilities of the desired outcome. The order of rules in a falling rule list reflects the order of the probabilities of the outcome.

Explanation-Table: The authors of [26] introduced an efficient method to generate *explanation tables* for multi-dimensional datasets. The proposed algorithm employs an information-theoretic approach to select patterns that provide the most information gain about the distribution of the outcome attribute. Since this algorithm does not consider an input SQL query, we have included another variant, denoted as Explanation-Table-G, to consider the query. This variant considers the grouping patterns found by CAUSUMX and reports the pattern found by Explanation-Table separately for each group.

XInsight: The authors of [42] introduced a framework that employs causal discovery and identifies causal patterns to explain the differences between two groups in a SQL query result. We generate an explanation by comparing all $\binom{m}{2}$ pairs of groups in Q(D). For a fair comparison, we skipped the causal discovery phase and gave XInsight the causal DAG CAUSUMX uses.

Since IDS, FRL, and Explanation-Table assume a binary outcome attribute, we binned the outcome variable in each examined scenario using the average outcome values. We gave FRL, IDS, and Explanation-Table the input table, as they do not consider a SQL query, obtaining a single set of rules for each dataset.

We also considered ChatGPT [4] as a baseline. We provided it with multiple prompts comprising a task description, SQL query, causal DAG, and dataset link. However, ChatGPT's responses consistently indicated its lack of direct access to external datasets. Consequently, it couldn't provide Summarized Causal Explanations For Aggregate Views

• For countries in Europe (e.g., Spain, Italy), the most substantial effect on high salaries (effect size of 34K, p < 1e-3) is observed for individuals under 45. Conversely, being under 25 has the greatest adverse impact on salary (effect size of -32K, p < 1e-3).

• For countries with a high GDP level (e.g., Sweden, Spain), the most substantial effect on high salaries (effect size of 40K, p < 1e-3) is observed for white individuals under 35. Conversely, being over 55 has the greatest adverse impact on salary (effect size of -34K,p < 1e-4). • For countries with a high Gini coefficient (e.g., Turkey, Brazil), the most substantial effect on high salaries (effect size of 29K, p < 1e-4) is observed for white individuals under 45. Conversely, being

being over 55 has the greatest adverse impact on salary (effect size of -28K, p < 1e-3).

Fig. 6. Causal explanation summary by CauSumX for SO use-case (sensitive attributes only).

specific insights on the datasets, offering only general domain-related insights instead. Although these insights included identifying attributes with strong causal effects on the outcome, they did not provide specific patterns. Consequently, we excluded this baseline from presentation.

Variations of CAUSUMX. We consider the following variations: **Brute-Force-LP**: As in Brute-Force, all grouping and treatment patterns are examined. In the final step, an approximated solution is obtained by employing the LP formulation described in Section 5.3.

Greedy-Last-Step: This baseline utilizes the approaches described in Section 5 to generate promising grouping and treatment patterns. In the final step, it employs a greedy strategy instead of solving an LP. The strategy involves iteratively selecting explanation patterns based on their explainability and the increase in coverage they offer.

Unless otherwise specified, the size constraint is set to 5, and the coverage threshold is set to 0.75. The threshold of the Apriori algorithm is set to 0.1. IDS, FRL, and Explanation-Table use their default parameters. The time cutoff is set to 3 hours. The experiments were executed on a PC with a 4.8GHz CPU, and 16GB memory.

6.2 Quality Evaluation (Q_1)

We assess the quality of our explanations relative to the baselines when examining a range of queries on diverse real-world datasets. Since no ground truth is available, we examine the consistency of our findings with insights from prior research (as was done in [29, 60, 80]). For each dataset, we present the result for a representative aggregated SQL query. Our queries are inspired by real-life sources, such as the Stack Overflow annual reports [1], media websites (e.g., The 19th Newsletter [3]), and academic papers (e.g., [8, 60]). More details and additional use cases are provided in [78].

SO. As in our running example, we consider a SQL query that computes the average salary of developers across the 20 most commonly mentioned countries among respondents (accounting for more than 85% of the SO dataset). To define grouping patterns, we considered attributes having FDs with COUNTRY: CONTINENT, HDI, GINI, and GDP. Here, for brevity, we set the solution size to 3. As mentioned, the explanation summary generated by CAUSUMX(shown in Figure 2), reveals key insights regarding the factors influencing salary across different countries. It highlights that job role, age, and education level are the primary determinants of income in the examined countries. Notably, individuals in C-level positions tend to earn significantly more than students. These results align with previous research [1, 5], which emphasized the significant influence of education level and job responsibilities on salaries in high-tech. Additionally, our results indicate that being below the age of 35 positively impacts salary, while being over 55 has a negative impact on income. Age discrimination toward people in the IT industry was identified in the literature [6, 39].

XInsight is designed to identify patterns that account for variations in the average salary across pairs of countries $\binom{20}{2}$, and thus, it results in an extensive explanation (of size exceeding 500KB). As, CAUSUMX, it focuses on patterns involving attributes that causally influence salary, such as

• For cities in the Northeast region (e.g., Boston, Albany), the combination of overcast weather conditions and low visibility has been found to have a substantial positive effect on severity (effect size of 0.55, p < 1e-3). The presence of traffic signals has been identified as the factor with the largest adverse impact on severity (effect size of -0.42, p < 1e-4).

• For cities in the Midwest region (e.g., Chicago, Detroit), the combination of cold temperatures and snow has been found to have a substantial positive effect on severity (effect size of 0.61, p < 1e-3). Clear weather has been identified as the factor with the largest adverse impact on severity (effect size of -0.31, p < 1e-3).

For cities in the South region (e.g., Huston, Miami), rain has been found to have a substantial positive effect on severity (effect size of 0.3, p < 1e-3). The presence of traffic calming measures has been identified as the factor with the largest adverse impact on severity (effect size of -0.44, p < 1e-3).
For cities in the West region (e.g., Phoenix, Los Angeles), the absence of traffic signals and traffic calming measures has been found to have a substantial positive effect on severity (effect size of 0.53, p < 1e-4). City roads (as opposed to highways) has been identified as the factor with the largest adverse impact on severity (effect size of -0.25, p < 1e-4).



Fig. 10. Analysis of precision and recall

Formal Education and Gender. However, these explanations offer distinct perspectives on the query results. For instance, XInsight's explanation sheds light on the substantial variance in average

Proc. ACM Manag. Data, Vol. 2, No. 1 (SIGMOD), Article 71. Publication date: February 2024.

salaries between the US and Poland, highlighting differences in the distribution of role types. In the US, there is a prevalence of machine learning specialists and C-level executives, while Poland has a higher concentration of Back-end developers and DevOps specialists. Conversely, CAUSUMX explanation uncovers commonalities among various groups, such as the positive impact of being under the age of 35 on salaries across all European countries. Therefore, these two solutions can be viewed as complementary efforts, each revealing different causal insights behind aggregate views. Nevertheless, if the query returns only two groups, CAUSUMX can find treatments (patterns) that explain the difference between their outcomes. In the US and Poland case, the returned treatments by CAUSUMX were similar to that obtained from XInsight(e.g., CAUSUMX also recognized that having a C-level executive position is the treatment with the highest positive effect on respondents from the US). In contrast, in the presence of multiple groups in the query result, XInsight lacks a straightforward extension for generating a summarized explanation for the entire aggregate view.

Both Explanation-Table and Explanation-Table-G aim to discover patterns related to high or low salaries. Notably, YearsCoding emerged as a significant factor in these patterns, indicating that individuals with over 30 years or less than 2 years of coding experience tend to have lower salaries. This aligns with our observation related to being a student or being over 55, which also leads to reduced salaries. However, Age and Role have a more substantial causal effect on salary than YearsCoding. One limitation of Explanation-Table is its inability to capture variations among different groups (countries). It either identifies country-specific patterns (e.g., low income for individuals from India) or universal patterns (e.g., YearsCoding). In contrast, Explanation-Table-G is designed to capture variations among groups but similarly highlights YearsCoding as a highly informative factor. This underscores the fundamental difference between Explanation-Table, which prioritizes patterns with high information gain, and our approach, which focuses on identifying patterns with strong causal effects.

Focusing on Sensitive Attributes. To identify potential biases, we focused exclusively on sensitive attributes (such as ethnicity, gender, and age) when examining treatment patterns. The generated explanation is depicted in Figure 6. Our results indicate that demographic factors significantly influence salary in all countries examined. Specifically, being under 35 positively impacts income, while being over 55 has a negative effect. Similarly, being a white male correlates with higher salary. These findings align with previous research on demographic impact on income, such as the gender wage gap [3] and disparities based on ethnicity [14, 20]. This showcases CAUSUMX's versatility in identifying causal explanations across different attributes. It also emphasizes its potential in uncovering disparities among demographic groups, supporting efforts to combat discrimination, and promoting equality.

Accidents. We investigated the average severity of car accidents across cities in the US. Our generated explanation is shown in Figure 7. It indicates that adverse weather conditions, such as cold temperatures and snow, tend to escalate the severity of car accidents. Conversely, the presence of traffic signals and calming measures appears to mitigate severity. Our results highlight variations in weather conditions across different regions. For instance, in the Midwest, cold temperatures and snow commonly contribute to severe accidents, whereas in the South, rainy weather emerges as a more significant factor for severe car accidents as snow and cold temperatures are less common. Previous studies [48, 49] have provided evidence supporting the effectiveness of traffic signals and calming measures in reducing accident severity. This analysis underscores the potential value of our system in extracting insights that go beyond common patterns and correlations. Such insights can be valuable for decision-makers, enabling a deeper understanding of causal factors and aiding in implementing effective road safety measures. The vast number of cities (over 50,000) made it unfeasible to generate an explanation using XInsight within our time constraints. The IDS, FRL, Explanation-Table and Explanation-Table-G baselines also exceeded our time cutoff.



6.3 Accuracy Evaluation

To evaluate the accuracy of the CAUSUMX algorithm, we conducted experiments over syntactic data where the ground truth is known. Here, we set n=1k and used the default system parameters. **Metrics of evaluation.** We report the precision and recall of the grouping and treatment pattern mining algorithms compared to the exhaustive Brute-Force baseline. To assess the precision and recall of the grouping mining algorithm (Section 5.1), we compare the tuples covered by the grouping patterns chosen by CAUSUMX against those covered Brute-Force. To assess the performance of the treatment mining algorithm (Section 5.2), we consider a fixed count of 20 grouping patterns, where both CAUSUMX and Brute-Forceselected the same grouping patterns. To evaluate the precision and recall of the treatment mining algorithm in comparison to the treatments chosen by Brute-Force, we compared the tuples identified as the treated group by CAUSUMX with those defined as the treated group according to Brute-Force. We report average precision and recall across all grouping-treatment pattern combinations.

Grouping patterns. We manipulate the count of attributes utilized for grouping patterns, thereby controlling the number of grouping patterns. The results are shown in Figure 10(a). Our results show that when the count of potential grouping patterns is relatively low (up to 20), our grouping patterns mining algorithm selects patterns that cover the same tuples as Brute-Force. As the number of grouping patterns to consider increases, the algorithm prunes certain patterns, leading to a decrease in precision and recall. Nevertheless, these scores consistently remain high (above 0.78), signifying that CAUSUMX and Brute-Force cover nearly the same tuples.

Treatment patterns. The results are shown in Figure 10(b). Observe that the recall remains consistently high, irrespective of the number of treatments. This suggests that the treated group according to Brute-Force is always encompassed by the treated group identified by CAUSUMX. In contrast, precision decreases as the count of treatment patterns increases. This signifies that the treated group according to CAUSUMX may contain irrelevant tuples. This is due to the pruning optimizations employed by CAUSUMX, which prevent it from materializing "long" patterns. However, even with many treatment patterns, the precision consistently exceeds 0.75.

6.4 Ablation Study (Q_2)

A breakdown analysis by step of CAUSUMX (Figure 14) shows that in all cases, mining the treatment pattern phase (Algorithm 2) consumes most of the time. The first and last steps are relatively fast. This aligns with our time complexity analysis. We next analyze the impact of our grouping and treatment pattern mining algorithms, as well as the LP formulation, compared to the optimal solution determined by Brute-Force based on Definition 4.5.

Runtime. Consider Figure 8(a). As expected, the algorithms that utilize our grouping and treatment patterns algorithms (CAUSUMX, Greedy-Last-Step) exhibit significantly faster performance



Fig. 14. Runtime by-step of the CAUSUMX algorithm

Stack Overflow

IMPLIS-CPS

20

compared to Brute-Forcevariants. Both Brute-Force and Brute-Force-LP could not process any dataset other than German, as their runtimes exceeded our time cutoff. *This clearly demonstrates the efficiency gained by our algorithms*. The disparity in running times between Greedy-Last-Step and CAUSUMX is negligible. This can be attributed to the fact that the treatment pattern detection phase consumes most of the execution time. Additionally, the relatively low number of grouping patterns explains the minimal difference in execution times during the last phase. Despite Greedy-Last-Step being faster than solving the LP, there are not too many explanation patterns to consider.

Coverage & Explainability. Figures 8(b) and (c) display the explainability and coverage of each baseline. In German, all baselines achieve the same coverage determined by k, but the approaches considering all patterns have higher explainability. This minimal difference highlights the effectiveness of our pattern mining and pruning techniques, improving runtime without compromising explainability significantly. In Accidents, CAUSUMX and Greedy-Last-Step yield identical solutions. However, in Adult, Greedy-Last-Step achieves higher explainability but falls short in coverage compared to CAUSUMX. This emphasizes the better balance achieved by our LP formulation in satisfying coverage and maximizing the objective compared to the greedy approach.

In-depth Analysis. In a comprehensive analysis comparing CAUSUMX and Greedy-Last-Step, we explored their performance by varying the solution size k on the SO dataset. The findings

are presented in Figure 9. As k increases, both algorithms exhibit improved overall explainability, with similar performance in this aspect. However, their behavior diverges in terms of coverage. CAUSUMX demonstrates a faster ability to satisfy the coverage constraint (shown by the dashed horizontal line). This is because CAUSUMX treats coverage as a constraint, while Greedy-Last-Step has no guarantees for coverage. As a result, Greedy-Last-Step only satisfies the coverage constraint for k=6. Based on our findings, CAUSUMX surpasses Greedy-Last-Step, as both achieve comparable objective values, but CAUSUMX has a higher likelihood of satisfying constraints.

6.5 Efficiency Evaluation (Q_3)

We showcase the scalability of CAUSUMX. For brevity, we excluded certain figures while discussing the trends they reveal within the text. We omit the results for IDS and FRL from the presentation, as their response times exceed 10 minutes. We also exclude the results of Explanation-Table-G since they exhibited similar trends to those of Explanation-Table. As we are not making a direct comparison with XInsight but instead with a variant that generates explanations for all pairs of groups, we have chosen to omit this baseline from presentation, as its runtime should be evaluated for the task of generating an explanation for a single pair.

Data Size. We analyze the impact of dataset size on runtime through random sampling of tuples. The results for the Adult and IMPUS-CPS datasets are shown in Figure 11. CAUSUMX and Brute-Force demonstrate a nearly linear increase in runtime for Adult (and SO - omitted from presentation) due to their full utilization of data for CATE value computation. However, CAUSUMX employed a sampling optimization for the larger IMPUS-CPS (and Accidents) dataset, resulting in a more consistent runtime. Explanation-Table's runtime is unaffected by dataset size due to sampling, but it is unable to handle datasets with more than 10 attributes (e.g., SO).

Attributes. We examine the impact of attribute quantity on runtime, by randomly excluding attributes from consideration. The results for the SO and Accidents datasets are shown in Figure 12. Brute-Force and Explanation-Table show exponential runtime increases with attribute number due to the growing number of patterns to consider. In contrast, CAUSUMX exhibits linear growth in runtime, thanks to pruning techniques mentioned in Section 5.2 that eliminate non-promising treatment patterns.

Treatment Patterns. We analyze the impact of treatment pattern quantity on runtime. We vary the number of bins for ordinal attributes and randomly exclude values for non-ordinal attributes. The results for Adult and IMPUS-CPS are displayed in Figure 13. Runtime increases linearly for all algorithms, which is expected due to the increased solution space.

Grouping Patterns. We examine the impact of grouping pattern quantity on runtime. By adjusting the threshold of the Apriori algorithm, we explore different numbers of grouping patterns. For CAUSUMX, the runtime remains relatively unchanged across all scenarios due to its simultaneous exploration of promising treatment patterns for each grouping pattern. However, Brute-Force's runtime increases linearly with the number of grouping patterns.

Solution Size. Lastly, we explore the impact of varying k. As this parameter affects only the final phase of CAUSUMX and Brute-Force, we observe negligible changes in their runtimes.

6.6 Explanations Sensitivity (Q₄)

We evaluate the impact of various parameters on the quality of the explanation summaries. The measures we focus on are overall explainability and coverage. Full details are provided in [78].

Apriori Threshold. We investigate the effect of varying the threshold parameter τ in the Apriori algorithm. Increasing τ leads to a reduction in the number of grouping patterns considered. Our findings indicate that higher τ values lead to a decrease in both explainability and coverage. Based on our findings, we recommend using a default threshold of 0.1, which provides satisfactory results



(a) Explainability



in terms of runtime, explainability, and coverage. However, it can be adjusted according to specific coverage requirements.

CATE Values Estimation. Recall that we use fixed-size random sampling for estimating CATE values. We investigate the impact of sample size on CATE value estimation. Figure 15 illustrates the results for the Accidents dataset. In Figure 15(a), we present the estimated CATE values for 5 random treatments using various sample sizes. In Figure 15(b), we evaluate the agreement between rankings using Kendall's τ correlation coefficient. We randomly selected 20 treatments and ranked them based on their CATE values, comparing this ranking with rankings obtained using different sample sizes. Notably, for a sample size of 1m tuples, CATE values exhibit an error of no more than 5%, and Kendall's τ reaches a high and stable value of 0.95. Similar trends were observed for the IMPUS-CPS dataset. Consequently, we conclude that a sample size of 1*m* tuples is suitable for accurate estimation of CATE values.

Causal DAG. We depart from the assumption of a given causal DAG and instead use existing solutions to construct DAGs. We conducted tests with multiple widely used causal discovery algorithms (the PC [67], FCI [67], and LiNGAM [65] algorithms), as well as with a simple, straightforward solution. In particular, we considered a causal DAG (referred to as No-DAG) wherein all attributes are directly linked to the outcome variable via edges, and no other edges exist within the DAG, similar to the approach taken in [29]. Our findings reveal that even the employment of basic causal discovery algorithms produces superior results when compared to the absence of any assumed causal DAG (Figure 16). We examined the effects on overall explainability and the ranking of treatment patterns when using different causal DAGs (we report the Kendall tau values, comparing the ranking of top-20 treatments based on their CATE values with a ranking obtained using the ground truth causal DAG). Notably, no single causal discovery algorithm outperforms all others, however, all of them outperform the No-DAG baseline. We note that causal DAGs can originate from various sources, including domain knowledge, GPT, or existing causal discovery methods.

This experiment illustrates that while our results do rely on the input causal DAG, the currently available methods for deriving such a DAG are capable of producing meaningful results.

7 LIMITATIONS AND FUTURE WORK

Our framework offers explanations to group-by-avg queries. Restricting to AVG is fundamental for causal explanations unlike non-causal methods. Our framework aims to find summarized *causal explanations* for the query answers in Q(D) using the concept of CATE discussed in Section 3. The causal estimates by CATE (as shown in Eq. (2) and (5)), inherently use *expectations, i.e., weighted average*. Thus, CATEs can be used when considering the aggregate average on the outcome column. Aggregate functions like SUM or COUNT depend on the *number of units* satisfying the grouping and treatment patterns, which does not have a correspondence with the estimate of causal effect (except that larger groups might reduce variance in the estimate). While other non-causal work on explanations or summarizations for query answers [36, 38, 45, 56, 73, 74] can support other aggregate functions, methods based on causal estimates focus on average as well [42, 60, 80].

Our framework currently supports a single-relation database with no dependencies among the tuples. The rationale behind this assumption is to ensure the SUTVA assumption [58] (as mentioned in Section Section 3) holds. Even for a single-table database with dependencies among tuples, this assumption no longer holds. For example, in a Flights dataset for flight delay, the delay of one flight has an impact on subsequent flights using the same aircraft, and is also dependent on flights leaving and arriving the same airport. Consequently, attempting to compute causal effects on such datasets would result in invalid results. When dealing with a single table, treatment patterns are well-defined, and grouping patterns are defined with attributes with FDs to the grouping attributes. However, extending these concepts to multi-table scenarios, where grouping attributes, treatments, and outcomes may originate from distinct tables, poses a challenging formalization task. This complexity arises from the need to address many-to-many relationships and patterns spanning multiple tables. Additionally, in multi-relational databases, intricate dependencies among tuples can exist, potentially violating the SUTVA assumption. While previous work [61] and [29] have extended causal models to accommodate multi-table data and employed them for hypothetical query answering, they have not specifically addressed explanations for groupby-avg queries. The extension of our framework to support multi-relational databases with complex dependencies is an important future work. It is worth noting that prior work on causal explanations [42, 60, 80] has primarily focused on single tables as well.

Lastly, note that attributes used for defining the grouping patterns must be categorical, as they need to exhibit a functional dependency with the grouping attribute. However, attributes used to define treatment patterns can take either continuous or categorical values. Handling continuous treatment variables poses specific challenges, but there are standard approaches in causal inference to address them, such as propensity weighting [54] or variable discretization. We note that with continuous treatment variables, the search space for potential treatment patterns significantly expands, and thus, while the operation of the treatment mining algorithm remains consistent, the execution times increase. Here, variable discretization may also be helpful, although this should be done cautiously to preserve the integrity of causal analysis. In our implementation, we did not discretize continuous variables, and we leave this optimization for future work.

ACKNOWLEDGMENTS

This work was partially supported by NSF awards IIS-1552538, IIS-1703431, IIS-2008107, IIS-2147061, and the NSF Convergence Accelerator Program award number 2132318.

REFERENCES

- [1] 2021. 2021 Stackoverflow Developer Survey. https://insights.stackoverflow.com/survey/2021.
- [2] 2021. Adult Census Income Dataset. https://www.kaggle.com/datasets/uciml/adult-census-income.
- [3] 2023. The 19th*. https://19thnews.org/2023/03/parenthood-stereotypes-gender-pay-gap/. Accessed: 2023-05-18.
- [4] 2023. OpenAI Introducing ChatGPT. https://openai.com/blog/chatgpt.
- [5] 2023. Viza jobs. https://vizajobs.com/what-do-technology-jobs-paysalary-insights-and-compensation-factors/.
- [6] 2029. Tech Talks. https://bdtechtalks.com/2019/03/29/ageism-in-tech-age-limit-software-developers-face/.
- [7] Rakesh Agrawal, Ramakrishnan Srikant, et al. 1994. Fast algorithms for mining association rules. In Proc. 20th int. conf. very large data bases, VLDB, Vol. 1215. Santiago, Chile, 487–499.
- [8] Mohamed Aljaban. 2021. Analysis of car accidents causes in the usa. (2021).
- [9] Yael Amsterdamer, Daniel Deutch, and Val Tannen. 2011. Provenance for aggregate queries. In Proceedings of the 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2011, June 12-16, 2011, Athens, Greece, Maurizio Lenzerini and Thomas Schwentick (Eds.). ACM, 153–164. https://doi.org/10.1145/1989284.1989302
- [10] Abolfazl Asudeh, Zhongjun Jin, and HV Jagadish. 2019. Assessing and remedying coverage for a given dataset. In 2019 IEEE 35th International Conference on Data Engineering (ICDE). IEEE, 554–565.
- [11] Arthur Asuncion and David Newman. 2007. UCI machine learning repository.
- [12] Abhijit V Banerjee, Abhijit Banerjee, and Esther Duflo. 2011. Poor economics: A radical rethinking of the way to fight global poverty. Public Affairs.
- [13] Senjuti Basu Roy, Sihem Amer-Yahia, Ashish Chawla, Gautam Das, and Cong Yu. 2010. Constructing and exploring composite items. In Proceedings of the 2010 ACM SIGMOD International Conference on Management of data. 843–854.
- [14] Marianne Bertrand and Sendhil Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. American economic review 94, 4 (2004), 991–1013.
- [15] Aline Bessa, Juliana Freire, Tamraparni Dasu, and Divesh Srivastava. 2020. Effective Discovery of Meaningful Outlier Relationships. ACM Transactions on Data Science 1, 2 (2020), 1–33.
- [16] Nicole Bidoit, Melanie Herschel, and Katerina Tzompanaki. 2014. Query-based why-not provenance with nedexplain. In Extending database technology (EDBT).
- [17] Shaofeng Bu, Laks VS Lakshmanan, and Raymond T Ng. 2005. Mdl summarization with holes. In Proceedings of the 31st international conference on Very large data bases. Citeseer, 433–444.
- [18] Adriane Chapman and HV Jagadish. 2009. Why not?. In Proceedings of the 2009 ACM SIGMOD International Conference on Management of data. 523–534.
- [19] Chaofan Chen and Cynthia Rudin. 2018. An optimization approach to learning falling rule lists. In International conference on artificial intelligence and statistics. PMLR, 604–612.
- [20] Raj Chetty, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter. 2020. Race and economic opportunity in the United States: An intergenerational perspective. *The Quarterly Journal of Economics* 135, 2 (2020), 711–783.
- [21] Silvia Chiappa. 2019. Path-specific counterfactual fairness. In Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 33. 7801–7808.
- [22] Leonardo Mendonça de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In TACAS. 337-340.
- [23] Daniel Deutch, Nave Frost, and Amir Gilad. 2020. Explaining Natural Language query results. VLDB J. 29, 1 (2020), 485–508.
- [24] Daniel Deutch and Amir Gilad. 2019. Reverse-Engineering Conjunctive Queries from Provenance Examples. In Advances in Database Technology - 22nd International Conference on Extending Database Technology, EDBT 2019, Lisbon, Portugal, March 26-29, 2019, Melanie Herschel, Helena Galhardas, Berthold Reinwald, Irini Fundulaki, Carsten Binnig, and Zoi Kaoudi (Eds.). OpenProceedings.org, 277–288. https://doi.org/10.5441/002/edbt.2019.25
- [25] Daniel Deutch, Amir Gilad, Tova Milo, Amit Mualem, and Amit Somech. 2022. FEDEX: An Explainability Framework for Data Exploration Steps. Proc. VLDB Endow. 15, 13 (2022), 3854–3868. https://www.vldb.org/pvldb/vol15/p3854-gilad.pdf
- [26] Kareem El Gebaly, Parag Agrawal, Lukasz Golab, Flip Korn, and Divesh Srivastava. 2014. Interpretable and informative explanations of outcomes. *Proceedings of the VLDB Endowment* 8, 1 (2014), 61–72.
- [27] Kareem El Gebaly, Parag Agrawal, Lukasz Golab, Flip Korn, and Divesh Srivastava. 2014. Interpretable and Informative Explanations of Outcomes. Proc. VLDB Endow. 8, 1 (sep 2014), 61–72. https://doi.org/10.14778/2735461.2735467
- [28] Sarah Flood, Miriam King, Steven Ruggles, and J Robert Warren. 2015. Integrated public use microdata series, current population survey: Version 9.0.[Machine-readable database]. *Minneapolis: University of Minnesota* 1 (2015).
- [29] Sainyam Galhotra, Amir Gilad, Sudeepa Roy, and Babak Salimi. 2022. Hyper: Hypothetical reasoning with what-if and how-to queries using a probabilistic causal approach. In Proceedings of the 2022 International Conference on Management of Data. 1598–1611.
- [30] Sander Greenland and James M Robins. 1999. Epidemiology, justice, and the probability of causation. Jurimetrics 40 (1999), 321.

- [31] Paul W Holland. 1986. Statistics and causal inference. *Journal of the American statistical Association* 81, 396 (1986), 945–960.
- [32] Alexandra Kim, Laks VS Lakshmanan, and Divesh Srivastava. 2020. Summarizing hierarchical multidimensional data. In 2020 IEEE 36th International Conference on Data Engineering (ICDE). IEEE, 877–888.
- [33] Been Kim, Cynthia Rudin, and Julie A Shah. 2014. The bayesian case model: A generative approach for case-based reasoning and prototype classification. *Advances in neural information processing systems* 27 (2014).
- [34] Himabindu Lakkaraju, Stephen H Bach, and Jure Leskovec. 2016. Interpretable decision sets: A joint framework for description and prediction. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. 1675–1684.
- [35] Laks VS Lakshmanan, Raymond T Ng, Christine Xing Wang, Xiaodong Zhou, and Theodore J Johnson. 2002. The generalized MDL approach for summarization. In VLDB'02: Proceedings of the 28th International Conference on Very Large Databases. Elsevier, 766–777.
- [36] Laks VS Lakshmanan, Jian Pei, and Jiawei Han. 2002. Quotient cube: How to summarize the semantics of a data cube. In VLDB'02: Proceedings of the 28th International Conference on Very Large Databases. Elsevier, 778–789.
- [37] Seokki Lee, Bertram Ludäscher, and Boris Glavic. 2020. Approximate summaries for why and why-not provenance (extended version). arXiv preprint arXiv:2002.00084 (2020).
- [38] Chenjie Li, Zhengjie Miao, Qitian Zeng, Boris Glavic, and Sudeepa Roy. 2021. Putting Things into Context: Rich Explanations for Query Answers using Join Graphs. In Proceedings of the 2021 International Conference on Management of Data. 1051–1063.
- [39] Yin Lin, Brit Youngmann, Yuval Moskovitch, HV Jagadish, and Tova Milo. 2021. On detecting cherry-picked generalizations. Proceedings of the VLDB Endowment 15, 1 (2021), 59–71.
- [40] Ester Livshits, Leopoldo E. Bertossi, Benny Kimelfeld, and Moshe Sebag. 2020. The Shapley Value of Tuples in Query Answering. In *ICDT*, Vol. 155. 20:1–20:19.
- [41] Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. 2013. Accurate intelligible models with pairwise interactions. In Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining. 623–631.
- [42] Pingchuan Ma, Rui Ding, Shuai Wang, Shi Han, and Dongmei Zhang. 2023. XInsight: EXplainable Data Analysis Through The Lens of Causality. Proc. ACM Manag. Data, Article 156 (jun 2023), 27 pages.
- [43] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F Moore, and Dan Suciu. 2009. Why so? or why no? functional causality for explaining query answers. arXiv preprint arXiv:0912.5340 (2009).
- [44] Alexandra Meliou, Wolfgang Gatterbauer, Katherine F Moore, and Dan Suciu. 2010. The complexity of causality and responsibility for query answers and non-answers. *arXiv preprint arXiv:1009.2021* (2010).
- [45] Zhengjie Miao, Qitian Zeng, Boris Glavic, and Sudeepa Roy. 2019. Going beyond provenance: Explaining query answers with pattern-based counterbalances. In Proceedings of the 2019 International Conference on Management of Data. 485–502.
- [46] Tova Milo, Yuval Moskovitch, and Brit Youngmann. 2020. Contribution Maximization in Probabilistic Datalog. In 2020 IEEE 36th International Conference on Data Engineering (ICDE). IEEE, 817–828.
- [47] Sobhan Moosavi, Mohammad Hossein Samavatian, Srinivasan Parthasarathy, and Rajiv Ramnath. 2019. A countrywide traffic accident dataset. *arXiv preprint arXiv:1906.05409* (2019).
- [48] Göran Nilsson. 1982. Effects of speed limits on traffic accidents in Sweden.
- [49] Ndhlovu Pardon and Chigwenya Average. 2013. The effectiveness of traffic calming measures in reducing road carnage in masvingo urban. *International Journal* 3, 2 (2013), 2305–1493.
- [50] Judea Pearl. 2009. Causal inference in statistics: An overview. (2009).
- [51] Prabhakar Raghavan and Clark D Tompson. 1987. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica* 7, 4 (1987), 365–374.
- [52] Alon Reshef, Benny Kimelfeld, and Ester Livshits. 2020. The Impact of Negation on the Complexity of the Shapley Value in Conjunctive Queries. In PODS, Dan Suciu, Yufei Tao, and Zhewei Wei (Eds.). 285–297.
- [53] James M Robins, Miguel Angel Hernan, and Babette Brumback. 2000. Marginal structural models and causal inference in epidemiology.
- [54] Paul R Rosenbaum and Donald B Rubin. 1983. The central role of the propensity score in observational studies for causal effects. *Biometrika* 70, 1 (1983), 41–55.
- [55] Sudeepa Roy, Laurel Orr, and Dan Suciu. 2015. Explaining query answers with explanation-ready databases. Proceedings of the VLDB Endowment 9, 4 (2015), 348–359.
- [56] Sudeepa Roy and Dan Suciu. 2014. A formal approach to finding explanations for database queries. In Proceedings of the 2014 ACM SIGMOD international conference on Management of data. 1579–1590.
- [57] Donald Bruce Rubin. 1971. The use of matched sampling and regression adjustment in observational studies. Ph.D. Dissertation. Harvard University.

- [58] Donald B Rubin. 2005. Causal inference using potential outcomes: Design, modeling, decisions. J. Amer. Statist. Assoc. 100, 469 (2005), 322–331.
- [59] Omer Sagi and Lior Rokach. 2021. Approximating XGBoost with an interpretable decision tree. Information Sciences 572 (2021), 522–542.
- [60] Babak Salimi, Johannes Gehrke, and Dan Suciu. 2018. Bias in olap queries: Detection, explanation, and removal. In Proceedings of the 2018 International Conference on Management of Data. 1021–1035.
- [61] Babak Salimi, Harsh Parikh, Moe Kayali, Lise Getoor, Sudeepa Roy, and Dan Suciu. 2020. Causal Relational Learning. In Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020, David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo (Eds.). ACM, 241–256. https://doi.org/10.1145/3318464.3389759
- [62] Gayatri Sathe and Sunita Sarawagi. 2001. Intelligent rollups in multidimensional OLAP data. In VLDB. 307-316.
- [63] Holger Schielzeth. 2010. Simple means to improve the interpretability of regression coefficients. Methods in Ecology and Evolution 1, 2 (2010), 103–113.
- [64] Amit Sharma and Emre Kiciman. 2020. DoWhy: An End-to-End Library for Causal Inference. arXiv preprint arXiv:2011.04216 (2020).
- [65] Shohei Shimizu, Patrik O Hoyer, Aapo Hyvärinen, Antti Kerminen, and Michael Jordan. 2006. A linear non-Gaussian acyclic model for causal discovery. *Journal of Machine Learning Research* 7, 10 (2006).
- [66] Bill Shipley. 2016. Cause and correlation in biology: a user's guide to path analysis, structural equations and causal inference with R. Cambridge university press.
- [67] P. Spirtes et al. 2000. Causation, prediction, and search. MIT press.
- [68] Yuchao Tao, Amir Gilad, Ashwin Machanavajjhala, and Sudeepa Roy. 2022. DPXPlain: Privately Explaining Aggregate Query Answers. Proc. VLDB Endow. 16, 1 (2022), 113–126. https://www.vldb.org/pvldb/vol16/p113-tao.pdf
- [69] Balder ten Cate, Cristina Civili, Evgeny Sherkhonov, and Wang-Chiew Tan. 2015. High-level why-not explanations using ontologies. In Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems. 31–43.
- [70] Jin Tian and Judea Pearl. 2000. Probabilities of causation: Bounds and identification. Annals of Mathematics and Artificial Intelligence 28, 1-4 (2000), 287–313.
- [71] Moshe Y. Vardi. 1982. The Complexity of Relational Query Languages (Extended Abstract). In Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing (San Francisco, California, USA) (STOC '82). ACM, New York, NY, USA, 137–146. https://doi.org/10.1145/800070.802186
- [72] Stefan Wager and Susan Athey. 2018. Estimation and inference of heterogeneous treatment effects using random forests. J. Amer. Statist. Assoc. 113, 523 (2018), 1228–1242.
- [73] Yuhao Wen, Xiaodan Zhu, Sudeepa Roy, and Jun Yang. 2018. Interactive summarization and exploration of top aggregate query answers. In Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases, Vol. 11. NIH Public Access, 2196.
- [74] Eugene Wu and Samuel Madden. 2013. Scorpion: Explaining away outliers in aggregate queries. (2013).
- [75] Yu Xie, Jennie E Brand, and Ben Jann. 2012. Estimating heterogeneous treatment effects with observational data. Sociological methodology 42, 1 (2012), 314–347.
- [76] Hongyu Yang, Cynthia Rudin, and Margo Seltzer. 2017. Scalable Bayesian rule lists. In International conference on machine learning. PMLR, 3921–3930.
- [77] Brit Youngmann, Sihem Amer-Yahia, and Aurélien Personnaz. 2022. Guided Exploration of Data Summaries. Proc. VLDB Endow. 15, 9 (2022).
- [78] Brit Youngmann, Michael Cafarella, Amir Gilad, and Sudeepa Roy. 2023. Techinical Report. https://anonymous.4open. science/r/Explanation_Summarization-F736
- [79] Brit Youngmann, Michael Cafarella, Yuval Moskovitch, and Babak Salimi. 2023. NEXUS: On Explaining Confounding Bias. In Companion of the 2023 International Conference on Management of Data. 171–174.
- [80] Brit Youngmann, Michael Cafarella, Yuval Moskovitch, and Babak Salimi. 2023. On Explaining Confounding Bias. 2023 IEEE 39th International Conference on Data Engineering (ICDE) (2023).
- [81] Brit Youngmann, Michael J. Cafarella, Babak Salimi, and Anna Zeng. 2023. Causal Data Integration. Proc. VLDB Endow. 16, 10 (2023), 2659–2665.
- [82] Cong Yu, Laks Lakshmanan, and Sihem Amer-Yahia. 2009. It takes variety to make a world: diversification in recommender systems. In Proceedings of the 12th international conference on extending database technology: Advances in database technology. 368–378.

Received July 2023; revised November 2023; accepted December 2023