

Suna: Scalable Causal Confounder Discovery over Relational Data

Jiaxiang Liu
jl6235@columbia.edu
Columbia University

Daniel Alabi
alabid@illinois.edu
University of Illinois Urbana-Champaign

Siyuan Xia
stevenxia@uchicago.edu
University of Chicago

Eugene Wu
ewu@cs.columbia.edu
Columbia University

ABSTRACT

Understanding the causal relationships between treatments and outcomes is fundamental in various areas. Causal inference aims to estimate the effect of one variable on another, and critically relies on access to those variables as well as the key confounders. Unfortunately, data analysts often start with datasets lacking these columns, leading to incorrect estimations. Relational data repositories hold significant potential to augment such datasets with an admissible set of confounders necessary for causal analysis. While recent work has advocated for this potential, these approaches face notable limitations. They either assume the existence of a complete causal diagram over all datasets in the repository, which is impractical; rely on computationally infeasible techniques that do not scale to large data repositories with many features; or can only detect confounders in the absence of causal relations, and are thus ineffective when a causal effect exists.

We observe that the asymmetry between causes and effects used in causal discovery can be exploited to directly identify confounders for causal queries. In this paper, we establish a connection between the existence of confounders and the presence of unconfounded ancestors of the treatment variable in the underlying causal diagram—without requiring access to the diagram. This makes it feasible to iteratively discover confounders until an admissible set is constructed. We propose Suna, a highly optimized, GPU-compatible system that implements a novel end-to-end algorithm for discovering confounders within large relational data repositories. Experiments on both real-world and synthetic datasets demonstrate that our system effectively discovers high-quality confounders. Furthermore, Suna employs algorithmic optimizations to accelerate confounder discovery without materializing joins. Our experiments show that Suna finds high-quality confounders while running $>100\times$ faster than existing confounder discovery systems.

PVLDB Reference Format:

Jiaxiang Liu, Siyuan Xia, Daniel Alabi, and Eugene Wu. . PVLDB, 18(11): 4158 - 4170, 2025.
doi:10.14778/3749646.3749684

PVLDB Availability Tag:

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 18, No. 11 ISSN 2150-8097.
doi:10.14778/3749646.3749684

The source code of this research paper has been made publicly available at <https://github.com/cudbg/Suna>.

1 INTRODUCTION

The goal of observational causal study is to understand the causal effect between a pair of treatment (T) and outcome (O). For example, medical researchers aim to understand to what extent smoking (T) elevates blood sugar levels (O), and advertisers try to determine whether a new ad campaign (T) leads to higher customer purchase rates (O). Understanding these causal relationships facilitates both knowledge discovery and decision-making. Due to the impossibility of observing the same subject under multiple treatment conditions simultaneously – known as the *fundamental problem of causal inference* – researchers must group data into comparable cohorts to effectively study these effects.

Despite the abundance of observational data, determining the causal effect from observational studies remains a challenging task. Researchers cannot actively form similar groups of individuals and assign treatments accordingly. A naive comparison of outcomes between individuals receiving different treatments is likely to yield a mixture of causal effect and spurious correlation. For example, smokers having a higher probability of blood sugar might be due to a specific type of gene that makes people more likely to smoke and have higher levels of blood sugar. Without controlling for these genetic factors, a naive causal analysis could mistakenly attribute the effect of these genetic factors to smoking, leading to a complete reversal of the true causal effect. To properly isolate spurious correlations from the causal effect, it is essential to identify key attributes of the data that forms an admissible adjustment set. By holding these attributes constant, different treatment groups become comparable in terms of their outcomes.

The identification of an admissible adjustment set presents significant practical challenges, because key attributes to control for confounding are often inaccessible. For example, data analysts struggle with picking an adjustment set without the aid of domain experts; more commonly, they start with datasets lacking critical attributes to form an admissible adjustment set for a treatment and outcome pair of interest [22]. Without access to these key attributes, the issue of *non-identifiability* arises, making it impossible to determine whether the correlation between the treatment and outcome is due to causal or non-causal effects. This can lead to incorrect conclusions regarding the causal query [36, 51].

Existing relational data repositories offer significant potential for enhancing data-oriented tasks, and prior work has applied this idea to machine learning tasks [14, 20]. In addition, past work [56]

has explored the overlap between causal inference and large data repositories by leveraging knowledge graphs to discover attributes that offset spurious correlation between a treatment and outcome pair. However, this work assumes that there is no causal relationship between the pair, making it inapplicable to non-zero causal effects.

In this paper, we present the first work that takes as input a table with treatment T and outcome O variables, and discovers an admissible adjustment set. We call this a *causal query* $Q = (T, O)$. We use relational data repositories as external knowledge sources to identify an admissible adjustment set. Since these repositories can contain many attributes, our approach needs to discover joinable datasets, integrate them, and identify a set of attributes as an admissible adjustment set from a substantial search space. While much existing work focuses on discovering and integrating [10, 26, 27], we tackle the challenge of efficiently searching for attributes to construct an admissible adjustment set for a causal query.

EXAMPLE 1. *Charlotte is deciding whether to pursue a master’s degree in computer science or enter the job market. She finds a dataset that contains worldwide education level and annual salary among software developers, and finds higher education correlated with much higher salaries. However, she did not expect such a high salary increase, and submits a causal query $Q = (\text{education}, \text{salary})$. Suna finds that high living and rent indices are confounders that drive higher education levels and higher salaries, and once conditioned on these attributes, the salary increase is much lower, which does not compensate the cost of pursuing an MS degree. As a result she decides to enter the job market.*

Finding confounders is straightforward given a complete *causal directed acyclic diagram* (causal DAG) over attributes in all datasets, where each edge is a causal relationship between two attributes. The *back-door criterion* [36, 37] directly finds an admissible adjustment set for a given $Q = (T, O)$ from the causal DAG. However, most datasets (and certainly not repositories) do not have causal DAGs. Thus, Charlotte’s causal query is not answerable.

The lack of a causal DAG is challenging for several reasons. First, it is difficult to even verify whether a candidate adjustment set is admissible for a causal query Q . Second, too many attributes are potentially relevant to the causal query, and there is an exponential number of attribute combinations that are potential candidate sets. Third, assessing each candidate adjustment set requires integrating of *all* participating datasets, which is impractically expensive.

In response, causal discovery methods aim to learn the underlying causal DAG from data. However, because these approaches target single tables, they do not scale to even modestly sized data repositories. For instance, Shimizu et al. [45] does not terminate within 10 hours on a data repository with 10K rows and 500 attributes. Some recent approaches rely on LLMs to construct the causal DAG [57], but it is not established whether LLMs have causal reasoning capabilities [11]. An additional consideration for practical confounder discovery systems is that the adjustment sets they find must remain a modest size so that analysts can interpret the output. While controlling for thousands of attributes might theoretically ensure validity, such expansive sets become computationally prohibitive and impossible to analyze.

To address the challenges above, we present a scalable confounder discovery system, Suna, that automatically searches for

a minimal, yet sufficient, set of relevant attributes over large data repositories. Given a causal query, Suna (1) identifies an admissible set of attributes as confounders with minimal size; (2) returns results within interactive timescales and (3) supports scaling to repositories with thousands or millions of attributes. The user submits a dataset and a causal query over a pair of attributes in the dataset, and Suna returns an admissible adjustment set (with links to the original datasets), an estimation of the average treatment effect, and the relevant tables. Internally, Suna iteratively finds attributes from datasets in the repository that are both unconfounded with the treatment attribute and a *causal confounder* of the causal query. We prove that this iterative approach can construct an admissible adjustment set without the need to build a causal DAG. To make Suna even faster and more scalable, we leverage factorized learning technique to avoid explicitly joining datasets and implement a GPU-friendly architecture that significantly accelerates execution.

Our key innovation is a novel framework that leverages *bivariate causal discovery*, a building block for parametric causal discovery, as an oracle to iteratively identify confounders. The main challenge is selecting a attribute that removes spurious correlations from a large pool of candidates; by doing this iteratively, we can find an admissible adjustment set. The novel insight from our main theorem shows that this step can be reduced to finding an unconfounded ancestor of the treatment attribute, that is not part of the adjustment set found so far. This insight also helps us determine when the current adjustment set becomes admissible, which avoids making the adjustment set unmanageably large. Experiments on real-world repositories show that Suna identifies meaningful confounders for causal queries that are consistent with prior studies. Synthetic experiments show that Suna scales to 1 million attributes, and answers the user’s causal query with an r^2 accuracy of >0.99 , compared to the ground truth. Our contributions include:

- We connect the existence of confounders to unconfounded ancestors of the treatment in an unknown causal diagram.
- We propose a novel framework that iteratively discovers confounders for causal queries based on bivariate causal discovery.
- We design a novel data structure that leverages ideas from factorized learning to avoid explicitly materializing joins during confounder discovery, and implement a GPU-compatible system to achieve interactive search latency.
- We evaluate the quality of Suna on a data corpus containing 346 datasets collected from NYC open data [7] and Kaggle datasets [25]. Results show that Suna discovers semantically meaningful confounders complying with existing studies, while achieving orders of magnitude runtime improvements over existing baselines. We further use synthetic datasets to quantitatively study the accuracy of Suna.

Note: The paper is self-contained. References to appendices can be disregarded or located in the technical report [6].

2 RELATED WORK

Confounder Discovery. There are two major lines of work for confounder discovery. One direction [34, 39, 54] explores different assumptions on the input data. CoCo [34] aims to identify confounded attributes assuming the data is observed under different contexts. Watson et al. [54] learns causal structures assuming a

known partitions of all attributes with one partition containing non-descendants of the other. However, they still focus on learning global causal structures instead of answering specific causal queries. HypDB [39] solves a similar problem to Suna, while restricting the scope to only consider the input data with limit attributes. Additionally, HypDB assumes that each parent of the treatment attribute is not a neighbor of another parent in the underlying causal diagram.

Another direction [14, 15, 56, 57] utilizes external data sources, but the focus is not on answering causal queries. MESA [56] assumes it is known a priori that there is no causal effect from T to O , and aims to find an admissible adjustment set. Metam [14] proposes a multi-armed bandit algorithm to balance exploring and exploiting datasets in the data corpus to handle a wide range of data tasks, including causal inference. They use existing techniques to answer causal inference queries, which is orthogonal to our approach. Nexus [15] discovers correlations within a data corpus to infer hidden confounders. Finally, a vision is proposed in [57] to first integrate treatment, outcome, and confounders into a unified dataset, followed by building a causal diagram with a large language model. Lastly, all approaches above involve materializing a single dataset that includes the treatment, outcome, and all candidate confounders, which is impractical for large data corpus.

Causal Discovery. The causal discovery literature focuses on relaxing assumptions for parametric and non-parametric approaches to learn the causal diagram. For example, one popular non-parametric causal discovery algorithm is the PC algorithm [47], which makes the *causal sufficiency* assumption that the underlying causal diagram does not contain any unobserved confounders. It begins with a complete undirected graph, iteratively prunes edges if two variables are conditionally independent given some other variables, and heuristically directs the remaining edges. The FCI algorithm [46] relaxes the *causal sufficiency* assumption by allowing unobserved confounders in the underlying causal diagram. As a tradeoff, FCI can only identify a set of causal DAGs, called a partial ancestral graph (PAG), that are compatible with the dataset. However, not all causal diagrams corresponding to a single PAG evaluate a causal query to the same answers [24]. In general, most non-parametric approaches are prohibitively expensive because (1) the number of conditional independence tests is exponential in the number of covariates, and (2) conditional independence tests are both sample inefficient and computationally expensive [43].

On the other hand, parametric causal discovery approaches make assumptions about the data generation process (e.g., additive noise). For instance, LiNGAM [44, 45, 53] assumes a linear causal mechanism with non-Gaussian additive noise, while Hoyer et al. [19] address non-linear causal mechanisms. These methods exploit asymmetries between causes and effects to infer causal relationships and topologically sort the variables in the underlying DAG. A final maximum likelihood estimation step learns a causal structure consistent with the topological ordering. Yet, the sorting and the estimation step are still impractically expensive and do not scale to corpora with many attributes, as is the case in data repositories.

Relational Causal Discovery. Another line of work in causal discovery focus on modeling causal relationships over relational data [29, 32, 33]. Specifically, the presence of many-to-many relationships between entities results in the violation of the *Stable*

Unit Treatment Value Assumption (SUTVA), requiring a different semantics for causal models over relational data. This line of work differs from ours in that it assumes a well-defined entity-relationship schema, which is not available for data corpus in the wild.

Task-based Dataset Search Systems There has been a number of recent work [8, 10, 14, 20, 21, 41, 58] that broadly fall under the category of task-based dataset search systems. On a high level, given a particular task (e.g. machine learning), these systems make use of external data to augment the performance on that task. Suna is one such system that is designed to optimize support for efficient and scalable causal queries.

Deconfounders. Deconfounders [9, 18, 52] infer synthetic latent variables as substitutes for actual confounders in causal analysis. However, these synthetic variables are typically less interpretable than their real counterparts, and thus we consider these approaches orthogonal to our focus on dataset search.

3 BACKGROUND

Data Model. For a given a relation R , we use uppercase letter X to denote an attribute, $R[X]$ be its corresponding column, $\text{dom}(X)$ be its domain, and $S_R = [X_1, \dots, X_m]$ be its schema. A tuple $t \in R$ denotes a tuple in R , with $t[X]$ representing the value of attribute X in tuple t . The domain of R is the Cartesian product of attribute domains $\text{dom}(R) = \text{dom}(X_1) \times \dots \times \text{dom}(X_m)$. Following conventions in statistics literature, we abuse notation by using the same X to denote the random variable representing attribute X . The annotated relational model [16] maps $t \in R$ to a commutative semi-ring $(D, \oplus, \otimes, 0, 1)$ where D is a set, \oplus and \otimes are commutative binary operators closed over D , and $0/1$ are zero/unit elements.

Causal Inference Foundations. The goal of causal inference is to infer the effect of a treatment attribute T on an outcome attribute O in some dataset about a study population. In the simplest case, the treatment received by each unit in the population is bivariate, consisting of two types of treatments: t_0 and t_1 . There are two corresponding potential outcomes, denoted as $O(t_0)$ and $O(t_1)$, that represent the outcome O had the treatment being assigned to t_0 and t_1 . From this point, we will use variables to denote attributes. One popular measurement to assess the causal effect for binary treatment variables is the *average treatment effect (ATE)*; the ATE of T on O can be defined as

$$\text{ATE}(T, O) = \mathbb{E}[O(t_1) - O(t_0)]$$

The expectation of the potential outcome, $\mathbb{E}[O(t_0)]$, is not equivalent to $\mathbb{E}[O \mid T = t_0]$, the expectation of the outcome given treatment t_0 , because the former represents the expected outcome *if all units had been assigned the treatment t_0* , which is a hypothetical scenario that cannot be observed. The *fundamental problem of causal inference* arises because each unit can only receive either treatment t_0 or t_1 . As a result, $\mathbb{E}[O(t_0)]$ cannot be trivially estimated from observational data [36]. To avoid bias and accurately estimate the potential outcome $\mathbb{E}[O(t_1)]$, a key step is to identify a set of variables Z such that, when conditioned upon, the treatment T is independent of the potential outcomes $O(t_0)$ and $O(t_1)$ (i.e., $O(t_0), O(t_1) \perp\!\!\!\perp T \mid Z$). Then, the definition of the average treatment effect can be rewritten using probability axioms:

$$\text{ATE}(T, O) = (\mathbb{E}[O \mid T = t_0, Z = \mathbf{z}] - \mathbb{E}[O \mid T = t_1, Z = \mathbf{z}]) \Pr(Z = \mathbf{z})$$




Figure 1: Example of a causal DAG \mathcal{G} .

Causal DAG. A causal DAG \mathcal{G} is a direct acyclic graph where each node in $\mathbf{V}(\mathcal{G})$ represents a variable, and each directed edge in $\mathbf{E}(\mathcal{G})$ indicates a potential causal relationship between two variables with the head node as cause and the tail node as effect. For example, Figure 1 (adopted from [55]) encodes potential causal relationships between 7 variables – large population size of a country (P) has a negative effect on the HDI of that country (H), which affects the internet penetration rate (I), mean year of schooling (M) in the country; a software developer’s education level (E) is affected by both M and H , and affects his length of work experience (W), which finally is causal factor of the developer’s salary (S).

For each node V_i in the causal diagram, all nodes having a directed arrow towards it are regarded as the parents of V_i , denoted as $\text{Pa}(V_i)$. A path p is a sequence of variables (V_1, \dots, V_k) such that each (V_i, V_{i+1}) is an edge in the undirected \mathcal{G} ; p is a directed path if each $(V_i, V_{i+1}) \in \mathbf{E}(\mathcal{G})$. A collider V_i in p is when both V_{i-1} and V_{i+1} are parents of V_i . For concreteness, consider variable H in Figure 1, $\text{Pa}(H) = \{P\}$ and S is a collider in the path $p = (W, S, H)$. A node V_i is a descendant of V_j if there exists a directed path (V_j, \dots, V_i) . We use $\text{De}(V_i)$ and $\text{ND}(V_i)$ to denote the set of descendants and non-descendants of V_i . Again, consider M in Figure 1, $\text{De}(M) = \{E, S, W\}$ and $\text{ND}(I) = \{H, P\}$.

D-separation and Causal Paths. D-separation offers a graphical criteria based on *open* and *closed* paths to determine (in)dependencies between pairs of variables without examining the data. For a causal diagram \mathcal{G} over variables $\mathbf{V}(\mathcal{G})$ and any pair of variables (V_i, V_j) such that $V_i \neq V_j$, a path $p = (V_1, \dots, V_k)$ in \mathcal{G} is *open* conditioned on a set of variables \mathbf{Z} if, for each triplet (V_{i-1}, V_i, V_{i+1}) in p if (1) V_i is not a collider, then $V_i \notin \mathbf{Z}$, and (2) V_i is a collider, then $\exists V_j \in \text{De}(V_i) \cup \{V_i\}$ such that $V_j \in \mathbf{Z}$.

Otherwise, p is *closed* conditioned on \mathbf{Z} . V_i is dependent on V_j conditioned on \mathbf{Z} if there exists a path $p = (V_i, \dots, V_j)$ that is *open* conditioned on \mathbf{Z} . A causal path between the treatment and outcome variable is a directed path from T to O in the causal DAG.

Causal Confounders and Admissible Adjustment Sets. A *causal confounder* [17] for Q is a variable that, when conditioned upon, reduces spurious correlations in estimating the causal effect. Similarly, we say a variable Z is a causal confounder relative to a set of variables \mathbf{Z} if conditioning on Z in addition to \mathbf{Z} further reduces spurious correlations compared to conditioning on \mathbf{Z} alone. From this point, we use confounders to refer to causal confounders.

An admissible adjustment set is a set of variables \mathbf{Z} such that when conditioned on, all non-causal paths from T to O are closed, and all causal paths from T to O remain open. Semantically, after condition on \mathbf{Z} , the dependency between T and O is only due to causal paths. Once an admissible adjustment set \mathbf{Z} has been identified, conventional machine learning techniques can estimate the

treatment effect. For example, we can train a linear regression model $O \sim \beta\mathbf{Z} + \alpha T$; the coefficient of the treatment T , α , estimates the average treatment effect of T on O .

Back-door Criterion. The back-door criterion [36] in the causal inference literature characterizes admissible adjustment sets using graphical conditions.

DEFINITION 3.1 (BACK-DOOR CRITERION). Given a causal diagram \mathcal{G} , a treatment variable T and an outcome variable O . A set of variables \mathbf{Z} satisfies the back-door criterion if (1) each back-door path, a directed path (T, V_i, \dots, O) where the first directed edge is $T \leftarrow V_i$, is closed conditioned on \mathbf{Z} ; and (2) $V \notin \mathbf{Z}$ if $V \in \text{De}(T)$.

Additive Noise Model. Following the definition of *Structural Causal Model (SCM)* [36], each variable in the causal model can be expressed as $V_i = f_i(\text{Pa}(V_i), \epsilon_i)$ where the causal mechanism f_i is an arbitrary function and ϵ_i is a random noise independent from any variable $V \in \text{Pa}(V_i)$. The *Additive Noise Model (ANM)* [47] is a special case of SCM where the independent noise is additive.

Causal Discovery for LiNGAM. Linear Non-Gaussian Acyclic Model (LiNGAM) is a type of ANM where each causal mechanism f_i is linear, and the random noise ϵ_i is non-Gaussian. Under LiNGAM assumptions, the correct causal direction between any two unconfounded variables with an unknown causal relationship can be determined. This is a special case applying the contrapositive of the Darmois-Skitovitch theorem [13]. In practice, one determines the causal direction between a pair of unconfounded variables A and B by training two linear regression models $M_{A \rightarrow B}$ (A predicts B) and $M_{B \rightarrow A}$, followed by a measure of independence [45]. The direction where there is a statistically significant dependency between the explanatory variable (e.g., A) and the prediction’s residual (e.g., $B - M_{A \rightarrow B}$) is rejected. The unconfoundedness between A and B is crucial, because otherwise the explanatory variable and the prediction residuals could be dependent even in the correct causal direction. We refer to this building block as *bivariate causal discovery (BCD)*. By assuming *causal sufficiency*, a source node S always exists and can be identified from the causal DAG; it is unconfounded with its descendants and independent of its non-descendants. Iteratively, a source node can be discovered and removed to sort $\mathbf{V}(\mathcal{G})$ in topological order.

4 PROBLEM DEFINITION

Problem Formulation. Let $\mathcal{R} = \{R_1, R_2, \dots\}$ be a data corpus with a set of relations. We denote the set of features for each relation R_i as \mathbf{A}_{R_i} , and let $\mathbf{A} = \cup_i \mathbf{A}_{R_i}$ denote the collection of all features across all relations. A data analyst submits a causal query $Q = (T, O)$ in an input relation R , and aims to discover a set of attributes as an admissible adjustment set to study the causal effect of T on O .

PROBLEM 1 (CONFOUNDER DISCOVERY FOR CAUSAL QUERY). Given a causal query $Q = (T, O)$ in input relation R , let \mathcal{G} be an underlying causal diagram with $\mathbf{V}(\mathcal{G}) = \mathbf{A} \cup \{T, O\}$, find a minimal set of attributes \mathbf{Z}^* in \mathcal{R} that is join-compatible with R such that

$$\begin{aligned} \mathbf{Z}^* &= \underset{\mathbf{Z}}{\text{argmin}} |\mathbf{Z}| && \text{minimality} \\ \text{s.t. } & T \perp\!\!\!\perp O \mid_{\mathcal{G}_{\mathbf{Z}}} \mathbf{Z} && \text{unconfoundedness} \\ & \mathbf{Z} \cap \text{De}(T) = \emptyset && \text{no disturbance} \end{aligned}$$




Figure 2: Overview of Suna’s architecture. Suna has an offline and an online component. The offline component preprocesses and computes sketches containing sufficient statistics from each dataset. The online component uses the computed sufficient statistics to iteratively discover causal confounders.

The two constraints correspond to the back-door criterion in Definition 3.1. Since the goal of data search is to find a small set of relevant datasets and attributes that are easy to manage and interpret by the users, we believe that a smaller adjustment set is preferred over a larger one. As a result, we add the minimality constraint in the objective function to favor a small Z .

Semantics. We follow the semantics of causal inference over relational data defined in CaRL [40], considering many-to-one joins between the input relation R as the fact table and each $R_i \in \mathcal{R}$ as a dimension table. Specifically, we preprocess each R_i by applying the mean aggregation (corresponding to the embedding function Ψ in CaRL) to its attributes with respect to each join key. Hence, the materialized join corresponds to the unit table including the treatment T , outcome O , and embedded candidate confounders Ψ_Z^O where standard causal inference approaches apply. To this end, we slightly abuse notation by denoting the embedded attribute Ψ_Z^O for each $Z \in \mathbf{A}$ as Z and each R_i as the preprocessed relation R_i .

Assumptions. We assume that all relations in \mathcal{R} are join-compatible with R where $R_{\bowtie} = R \bowtie \{R_i\}_i$. Due to the lack of well-defined join conditions between R and each R_i , instead of modeling $\{R\} \cup \mathcal{R}$ with a relational causal diagram [29, 32, 33], we treat columns in R_{\bowtie} as nodes in a standard causal diagram satisfying *causal sufficiency* and can be modeled with LiNGAM for confounder discovery. While the non-Gaussianity property from LiNGAM may seem counterintuitive, this property is frequently observed in real-world data, such as financial and sensor data [12, 49]. In fact, the join-compatibility of **all** relations in \mathcal{R} with R can be relaxed to handle multiple causal queries; for each causal query Q_i , the subset of all relations join-compatible with R is modeled with LiNGAM.

5 SYSTEM ARCHITECTURE

The primary contribution of this work is our algorithm for identifying confounders without the need to construct the full causal diagram or materializing $R \bowtie \{R_i\}_{i=1}^n$. Thus, we design Suna to support an efficient and scalable implementation of the algorithm.

Figure 2 gives an overview of the system architecture. Suna has an offline and an online phase. The offline phase involves collecting tabular datasets from open data repositories [7, 25], preprocessing them by aggregating and removing outliers [31], then transforming them into semi-ring sketches that will be used as input to the confounder discovery algorithm in the online phase. A crucial feature of these sketches is that they allow the confounder




Figure 3: Figure 3a is an example causal diagram \mathcal{G} where $V(\mathcal{G}) = \{A, T, V, Z\}$, Figure 3b is a causal diagram representing the projection of \mathcal{G} onto $\{T, A\}$. Figure 3c is a snippet of a causal diagram with $\{A, T, Z, O\} \subseteq V(\mathcal{G})$, dotted directed edges indicate a directed path from the head node to the tail node in \mathcal{G} , i.e., there exists a directed path from Z to T in \mathcal{G} .

discovery algorithm to be run without materializing joins, leading to large performance gains. Suna’s online phase processes analysts’ causal queries. Upon receiving a query, Suna’s *data discovery component* leverages existing techniques to find datasets that are join-compatible with the query input. The query is then processed by the *causal confounder evaluation component*, which employs BCD as an oracle to directly find confounders. Moreover, the *causal confounder evaluation component*’s algorithm is designed to be parallelizable, making it possible for us to implement GPU acceleration. These components work together to let Suna achieve interactive latency in processing, even for large-scale datasets.

6 CAUSAL CONFOUNDER DISCOVERY

We now describe the theoretical and algorithmic foundations to solving Problem 1. We refer to attributes in \mathbf{A} as variables, and describe how to select a subset of \mathbf{A} as an adjustment set. The main challenge is that *no method exists for validating the correctness of an arbitrary Z without referring to a causal diagram*. We present a novel algorithm that uses BCD to identify the topological order for a superset of the adjustment set in the underlying (but not accessible) causal diagram, and iteratively select an adjustment set that satisfies *unconfoundedness* and *no disturbance* in Problem 1.

This section is organized into three parts. Firstly, we present a baseline that extends the LiNGAM causal discovery algorithm to take the causal query $Q = (T, O)$ into account, and analyze its limitations. Secondly, we prove the main theory that bridges BCD with confounder discovery and enables our iterative discovery method. Lastly, we describe our end-to-end algorithm.

6.1 Baseline

The back-door criterion implies that the set of all non-descendants of the treatment variable $ND(T)$ forms an admissible adjustment set for Q because it automatically closes all back-door paths. As a result, one simple extension of LiNGAM’s causal discovery algorithm is to build the adjustment set by starting with an empty adjustment set $Z = \emptyset$, and iteratively add the source node of \mathbf{A} , as discovered by the vanilla LiNGAM algorithm, to Z . Since variables are added to Z in topological order in \mathcal{G} , we can terminate the vanilla LiNGAM algorithm and conclude that the current Z is admissible if the current source node is T . It is guaranteed that $Z \subseteq ND(T)$, so Z forms an admissible adjustment set.

However, this approach falls short in three aspects. First, the iterative process of obtaining a topological order over $\mathbf{A} \cup \{T, O\}$ spends extra time sorting variables that are not confounders. Second, $Z =$

ND(T) can be arbitrarily large, containing attributes spanning over many datasets. This fails to respect the minimality requirement in Problem 1. Finally, the high dimensionality of Z introduces high variance in estimating $\Pr(O \mid T, Z)$, which affects the accuracy in estimating treatment effects.

EXAMPLE 2. Consider $Q = (E, S)$ with Figure 1 as the underlying causal DAG. (P, H, I, M, E, W, S) is a valid topological order over $V(\mathcal{G})$, suggesting $Z = \{P, H, I, M\}$. However, P and I do not close any backdoor path for Q , so $Z = \{H, M\}$ is still admissible.

6.2 The Main Theorem

In this section, we first introduce causal diagram projection, which defines the causal relationship on a subset of variables $V' \in V(\mathcal{G})$. We then present our insights on this projection and state our main theorem, which provides the theoretical support for an iterative algorithm that only identifies confounders.

6.2.1 Connection between Causal Confounder Discovery and Bivariate Causal Discovery. We adopt the definition of causal diagram projection from [50]. At a high level, this definition characterizes the causal relationships among a subset of variables within a larger causal diagram. This approach allows us to focus on the causal structures between variables of interest. Given two variables V_1 and V_2 in $V(\mathcal{G})$ where (1) there are no open non-causal paths between V_1 and V_2 (*unconfounded*), and (2) there exists a causal path from V_1 to V_2 (*causally related*), any projection of $V(\mathcal{G})$ onto a set $\{V_1, V_2\} \subseteq V'$ will include a directed edge $V_1 \rightarrow V_2$ and will not contain a bi-directed edge $V_1 \leftrightarrow V_2$.

EXAMPLE 3. Consider variables T and A in the causal DAG in Figure 3a, the non-causal path, (A, V, Z, T) , between A and T is closed, while the causal path (A, V, T) is open. Thus, the projection of this causal DAG onto $\{A, T\}$ shows a simple cause-and-effect relationship as shown in Figure 3b.

Under the LiNGAM assumption, the projection of a causal DAG onto two unconfounded and causally related variables will always exhibit a simple cause-and-effect relationship satisfying LiNGAM. With this observation, we now introduce our main theorem.

THEOREM 6.1. Fix a causal diagram \mathcal{G} , a causal query $Q = (T, O)$, and a set of variables Z where $Z \cap \text{De}(T) = \emptyset$. If there exists an open non-causal path for Q conditioned on Z in \mathcal{G} , there must exist a variable $Z \in V(\mathcal{G})$ satisfying (1) there is an open causal path from Z to T , and all non-causal paths between Z and T are closed conditioned on Z in the projection of \mathcal{G} onto $\{T, Z\} \cup Z$; and (2) Z is a confounder for Q conditioned on Z .

PROOF. The full proof can be found in the technical report. \square

One direct implication of Theorem 6.1 is that, when $Z = \emptyset$, we can detect the presence of open non-causal paths in the causal diagram for Q , by applying BCD on each (T, Z) pair where $Z \in V(\mathcal{G})$. If such a path exists, there will be at least one Z for which regressing T on Z results in a significant dependency between T and this regression's residuals. Conversely, if such a Z does not exist, we can safely conclude that there is no open non-causal path for Q in \mathcal{G} conditioned on Z ; hence Z forms an admissible adjustment set. However, Theorem 6.1 does not immediately translate to an algorithm

Algorithm 1 Iterative Confounder Discovery

```

1: Input:  $\mathcal{R}, (T, O) \in S_R, \tau$ 
2: Return:  $Z$ 
3:  $Z \leftarrow \{\}$ 
4: while True do
5:    $C \leftarrow \{\}$ 
6:   for all  $R_i \in \mathcal{R}$  do
7:      $\mathcal{B} \leftarrow \text{gen-bootstrap}(\mathcal{R}, R_i)$  ▷ generate bootstrap samples
8:     for all  $Z \in R_i$  do ▷ augment  $R$  w/  $R_i$ 's attrs
9:       for all  $b \in \mathcal{B}$  do
10:         $p_b^{Z+}, p_b^{Z-} \leftarrow \text{bivariate-CD}(b, Z)$ 
11:         $CI \leftarrow CI \cup \{p_b^{Z-} - p_b^{Z+}\}$ 
12:      end for
13:       $C \leftarrow C \cup Z$  if  $q_\tau(CI)$  significant
14:    end for
15:  end for
16:  break if  $C = \{\}$ 
17:   $Z_{opt} \leftarrow \text{heuristic-search}(C)$ 
18:   $Z \leftarrow Z \cup Z_{opt}$ 
19:   $\text{update-corpus}(\mathcal{R}, Z_{opt})$ 
20: end while

```

because (1) empirical estimation error can lead to false positives and false negatives that violates the *no disturbance* condition, (2) it only guarantees to find a superset C containing a confounder, without giving a strategy to select the confounder from C , and (3) it does not address how to account for a non-empty Z in BCD after the first iteration. We now present our solution to these challenges.

6.3 Algorithm

This section presents the full algorithm that iteratively discovers confounders for Q in Algorithm 1, while accounting for the three challenges explained previously. We first give the algorithm overview, then explain individual components in depth. On a high level, we start with an empty adjustment set $Z = \emptyset$ and add variables to Z once at a time. In each iteration, we first find a superset of attributes that contains a confounder, corresponding to the first condition in Theorem 6.1, and then heuristically select the confounder out from the superset, corresponding to the second condition in Theorem 6.1. Semantically, an attribute Z is added to Z in each iteration if $Z \cup Z$ further reduce spurious correlation between T and O . The input of the algorithm is the data corpus \mathcal{R} and an input dataset containing the treatment and outcome pair (T, O) . Each iteration tests whether $Z \in R_i$ is a unconfounded ancestor of T (L6-15), with respect to the current adjustment set Z . `heuristic-search` then explores these confounder candidates and identifies the optimal confounder Z_{opt} (L17). Lastly, we update Z (L18) and \mathcal{R} (L19) to facilitate confounder discovery in future iterations.

6.3.1 Bootstrap Resampling. To reduce estimation error and enhance robustness, `gen-bootstrap` (L7) generates resampled datasets with replacement for conducting hypothesis tests across multiple samples. `bivariate-CD` (L10) implements BCD, returning dependency measures between the explanatory variable and prediction residuals for both forward ($Z \rightarrow T$) and backward ($T \rightarrow Z$) causal directions. We employ mutual information scores to quantify these dependencies [28], defining $p^{Z+} = \text{MI}(Z; \epsilon_T)$ and $p^{Z-} = \text{MI}(T; \epsilon_Z)$, where ϵ_T and ϵ_Z represent prediction residuals in their respective models. Variables f for which the $100 \cdot (1 - \tau)$ -percentile of the bootstrapped distribution ($q_\tau(CI)$) is significant are considered




Figure 4: Figure 4a is a causal diagram consisting treatment T and outcome O , a latent variable ϵ_T representing the noise of T , and the three other variables Z_1, Z_2 and Z_3 . Figure 4b is a causal diagram considering the residuals of T and Z_3 predicted by Z_2 .

as candidates of confounders (L13), where τ is a tunable hyperparameter defaulted to 0.05. Semantically, this difference serves as a score indicating the likelihood of Z being an ancestor of T .

6.3.2 Heuristic Selection. Theorem 6.1 guarantees that there exists a variable $Z \in C$ such that conditioning on Z mitigate spurious correlation between T and O and reduces their dependency. As a result, one approach to selecting a confounder from C is to estimate the mutual information score between T and O with and without conditioning on Z for each $Z \in C$. A non-trivial difference in $MI(T; O | Z) - MI(T; O | Z \cup \{Z\})$ suggests that Z is a confounder. Therefore, once we have identified a superset of variables containing at least one confounder, we can select the confounder based on a drop in mutual information when conditioned upon. Empirically, we can also use the absolute difference in coefficients of T regressing $T \cup Z$ and $\{T, Z\} \cup Z$ on O .

6.3.3 Update Residuals. In iterations where Z is non-empty, our algorithm must identify variables that close remaining open non-causal paths *conditioned on the current Z* . For clarity of illustration, we assume the current adjustment set Z contains a single confounder, though the underlying methodology generalizes seamlessly to adjustment sets of higher cardinality. Let $Z = \{A\}$, due to condition (1) in Theorem 6.1, A is an admissible adjustment set for the causal query $Q = (Z, T)$. Hence, for linear causal mechanisms, we write $T = \alpha Z + \beta A + \epsilon_T$ for arbitrary constants α and β where $A, Z \perp\!\!\!\perp \epsilon_T$. Let $\epsilon_{T,A}$ and $\epsilon_{Z,A}$ denote the prediction residuals of regressing A on T and regressing A on Z , then

$$\epsilon_{T,A} = \alpha \epsilon_{Z,A} + \epsilon_T$$

If we treat $\epsilon_{T,A}$ and $\epsilon_{Z,A}$ as variables, and construct a causal diagram between them, the causal diagram is $\epsilon_{Z,A} \rightarrow \epsilon_{T,A}$ that satisfies the LiNGAM property. Hence, a significant $p^{\epsilon_{Z,A}^-} - p^{\epsilon_{Z,A}^+}$ indicates that Z is a variable satisfying both conditions in Theorem 6.1.

EXAMPLE 4. Consider a causal diagram \mathcal{G} in Figure 4a, where T is the treatment variable and $Z = \{Z_2\}$. In this setup, the causal path from Z_3 to T is open, while the non-causal path (T, Z_1, Z_2, Z_3) is closed conditioned on Z_2 . By computing the residuals ϵ_{T,Z_2} and ϵ_{Z_3,Z_2} after regressing Z_2 on T and Z_3 , the causal diagram regarding ϵ_{T,Z_2} and ϵ_{Z_3,Z_2} is shown in Figure 4b.

For each iteration, we always need to consider the current Z over BCD. That is, we work with $\epsilon_{T,Z}$ and $\epsilon_{Z,Z}$, the residuals of T and Z after regressing on all variables in Z , for each $Z \in \mathbf{Z}$ in each iteration. By applying the Frisch-Waugh-Lovell theorem, those $\epsilon_{T,Z}$ and $\epsilon_{f,Z}$ can be maintained incrementally over the search.

EXAMPLE 5. Let Z_1 and Z_2 be the causal confounders discovered over the first 2 iterations. We want to compute BCD between T and Z_3 conditioned on $Z = \{Z_1, Z_2\}$. Note that:

$$\begin{aligned} T &= \alpha_1 Z_1 + \alpha_2 Z_2 + \epsilon_{T,Z} \\ \epsilon_{T,Z_1} &= \alpha_2 \epsilon_{Z_2,Z_1} + \epsilon_{T,Z} \end{aligned}$$

Hence, we maintain ϵ_{T,Z_1} and ϵ_{Z_2,Z_1} at the end of the first iteration once Z_1 is discovered. Then, we can calculate $\epsilon_{T,Z}$ by regressing ϵ_{Z_2,Z_1} on ϵ_{T,Z_1} and taking the prediction residual.

7 OPTIMIZATIONS

The major performance bottleneck is L6-15 in Algorithm 1, which runs BCD between the treatment variable and every other variable in \mathbf{A} , and update residuals in L19. Suna accelerates these steps by factorizing BCD to avoid join materialization and using GPU-acceleration to answer queries in seconds. In fact, our algorithm extends beyond optimizing BCD and accelerates estimating mutual information between attributes across multiple datasets.

Main Bottlenecks. For every pair of treatment and potential confounder variables (T, Z) , where $T \in S_R$ and $Z \in S_{R_i}$, bivariate-CD consists of four steps: (1) materialize $R \bowtie_{\mathcal{J}} R_i$; (2) train linear regression models $M_{T \rightarrow Z}$ to predict Z from T and $M_{Z \rightarrow T}$ to predict T from Z , and calculate the prediction residuals $\epsilon_Z = Z - M_{T \rightarrow Z}$ and $\epsilon_T = T - M_{Z \rightarrow T}$; (3) assess the dependency level between T and ϵ_Z and between Z and ϵ_T by computing the mutual information scores $MI(T; \epsilon_Z)$ and $MI(Z; \epsilon_T)$; and (4) use Z_{opt} to predict all variables in \mathbf{A} and update them to their respective prediction residuals, once the optimal confounder Z_{opt} is discovered. The steps are computationally expensive because of the many (T, Z) pairs (due to potentially large numbers of candidate confounders), and the high cost of materializing joins, training models, computing residuals and estimating mutual information.

Factorized Learning. To address the challenges of training regression models and updating variables to their respective prediction residuals, Suna leverages *factorized learning* techniques [31] to distribute linear regression training and inference over joins by pre-computing sketches. These sketches contain semi-ring annotations that transform group-by and join operations into addition (\oplus) and multiplication (\otimes) operations within a semi-ring structure $(D, \oplus, \otimes, 0, 1)$. The annotation for group-by $\gamma_{\mathcal{J}} R$ is the sum of the annotations within the group, and the annotation of a tuple $t \in R_1 \bowtie R_2$ is the product of annotations from contributing tuples in R_1 and R_2 . The sketches for regression models are called the variance semi-ring, which consists of the count, sum and sum of squares of attributes for each candidate confounder $Z \in \mathbf{A}$ over \mathcal{J} . The key optimization of *factorized learning* pushes aggregations γ (\oplus) before joins \bowtie (\otimes) to avoid aggregating on the materialized join. During pre-computation, each tuple in $t \in R$ is lifted to a semi-ring annotation via a lift function $g(\cdot) : \text{dom}(R) \rightarrow D$, and aggregated for each group defined by \mathcal{J} . Pre-aggregated sketches have several advantages: (a) they are smaller than the original dataset—linear in the join key’s domain; (b) they are sufficient for training linear regression models over joins; and (c) they can be easily updated to a new variance semi-ring suitable for deriving prediction residuals. When combined with GPU acceleration, Suna can train of over one million linear regression models through joins in $\approx 1\text{sec}$.

Although factorized learning techniques address steps (1,2,4), they do not address (3), because of the inherent challenges in estimating mutual information. Estimating $MI(T; \epsilon_Z)$ for continuous variables T and ϵ_Z requires estimating the differential entropies $h(T)$, $h(\epsilon_Z)$ and joint differential entropy $h(T, \epsilon_Z)$, which in turn depends on estimating $\Pr(T, \epsilon_Z)$. Due to the invariance property of mutual information under invertible transformations, bivariate-CD(T, Z) is equivalent to bivariate-CD(\tilde{T}, \tilde{Z}), where \tilde{T} and \tilde{Z} denote the normalized columns T and Z in R_{\bowtie} . Then, the significance level of a candidate confounder can be rewritten as

$$p^+ - p^- = MI(\tilde{T}; \tilde{Z} - \beta \cdot \tilde{T}) - MI(\tilde{Z}; \tilde{T} - \beta \cdot \tilde{Z}) = MI(\tilde{T}; \tilde{\epsilon}_Z) - MI(\tilde{Z}; \tilde{\epsilon}_T)$$

As pointed out in [23], this is equivalent to estimating $h(\tilde{T}) + h(\tilde{\epsilon}_Z) - h(\tilde{Z}) - h(\tilde{\epsilon}_T)$, which boils down to estimating the marginal pdfs $\Pr(\tilde{T})$, $\Pr(\tilde{Z})$, $\Pr(\tilde{\epsilon}_T)$ and $\Pr(\tilde{\epsilon}_Z)$. However, these distributions can only be computed after materializing R_{\bowtie} . For instance, in Figure 6a, the distribution of $R[T]$ differs from that of $R_{\bowtie}[T]$ in Figure 6b. Further, no semi-ring (sufficient statistics) exists for this estimation, which necessitates join materialization. To avoid this, we develop a two-phase solution: we first adopt equi-width histograms for marginal pdf estimation, then design a novel semi-ring structure that efficiently computes these histogram-based estimates.

7.1 Histogram-based BCD

7.1.1 Equi-width Histogram. For a relation R and an attribute $X \in S_R$, an equi-width histogram over $R[X]$ places $\{t[X]\}_{t \in R}$ into non-overlapping bins of the same width. To minimize the integrated mean squared error between the histogram and the underlying distribution, we follow Scott's rule [42] and set the bin width to be $w_x = (24\sqrt{\pi})^{1/3} \sigma_X n^{-1/3}$, where σ_X is the (empirical) standard deviation of $R[X]$ and $n = |R|$. In fact, any binning strategy based on σ_X and n is supported. We define the *bin assignment function* $f : \mathbb{R} \rightarrow \mathbb{Z}$ that maps each $t[X]$ to its corresponding bin as

$$f : x \mapsto \lfloor (x - \min\{t[X]\}_{t \in R}) / w_x \rfloor$$

Then, the final histogram is an aggregation query that counts the number of tuples within each bin, which can be expressed in relational algebra as $\gamma_{f(X), COUNT}(R)$, or simplified as $\gamma_{f(X)}(R)$. $\Pr(X)$ is then approximated by normalizing the bin counts by n . We aim to answer the queries $\gamma_{f(X)}(R_{\bowtie})$ for $X \in \{\tilde{T}, \tilde{Z}, \tilde{\epsilon}_T, \tilde{\epsilon}_Z\}$. Notably, σ_X and n in R_{\bowtie} come free with the variance semi-ring, and $\min\{t[X]\}_{t \in R}$ can be computed using a tropical semi-ring [16]. For simplicity, we denote $\min\{t[X] \mid t \in R\}$ as α_X and treat these statistics as constants as they are not the focus of this section.

7.1.2 Histogram Semi-ring. We define the histogram semi-ring $(D, \oplus, \otimes, 0, 1)$, where D is the set of ordered pairs in $\mathbb{Z} \times \mathbb{Z}$. To ease expression, we use a key-value pair (k, v) to denote an element of D with $(k, 0) \in D$ if the key k does not exist in D . For two semi-ring annotations $D_1, D_2 \in D$, we define the semi-ring addition \oplus and semi-ring multiplication \otimes as

DEFINITION 7.1 (SEMI-RING ADDITION \oplus). *The semi-ring addition \oplus of $D_{\oplus} = D_1 \oplus D_2$ is defined as $D_{\oplus} = \{(k, v_1 + v_2) \mid \forall (k, v_1) \in D_1, (k, v_2) \in D_2\}$. The additive identity (0) is defined as $\{(0, 0)\}$.*

DEFINITION 7.2 (SEMI-RING MULTIPLICATION \otimes). *The semi-ring multiplication \otimes of $D_{\otimes} = D_1 \otimes D_2$ can be defined as $\{(k, \sum_{k_1+k_2=k} v_1 \cdot$*




Figure 5: Semantics for addition in Figure 5a and multiplication in Figure 5b over histogram semi-ring.

$v_2) \mid (k_1, v_1) \in D_1, (k_2, v_2) \in D_2\}$. The multiplicative identity (1) is defined as $\{(0, 1)\}$.

PROPOSITION 7.1. *The semi-ring add (\oplus) and multiplication (\otimes) satisfy the commutative and distributive law.*

Intuitively, the histogram semi-ring represents local histograms within each join key. We illustrate the semi-ring operators through an example in Figure 5.

EXAMPLE 6. *The \oplus operator in Figure 5a corresponds to the bin-wise addition of histograms. The \otimes operator computes new histogram bins through a multiplication-addition mechanism: combining bin 0 of the left histogram with bin 1 of the right histogram in Figure 5b creates a new bin whose key is the sum of the participating bins' keys $(0 + 1)$, and value is the product of the participating bins' values $(1 \cdot 1)$.*

Next, we use the histogram semi-ring to estimate $\Pr(\tilde{T})$ and $\Pr(\tilde{\epsilon}_Z)$. The same approach applies to $\Pr(\tilde{Z})$ and $\Pr(\tilde{\epsilon}_T)$.

7.1.3 Estimating pdf for \tilde{T} . To compute $\gamma_{f(\tilde{T})}(R_{\bowtie})$ without materializing R_{\bowtie} , we distribute the aggregation through joins using two steps: (1) compute per-join-key histograms of \tilde{T} in R , and (2) obtain per-join-key counts in R_i via $\gamma_{\mathcal{J}}(R_i)$. Leveraging the histogram semi-ring, we define the lift function $g_{\tilde{T}}(\cdot) : \text{dom}(R) \rightarrow D$ as

$$g_{\tilde{T}}(t) = \begin{cases} (f(t[\tilde{T}]), 1) & t \in R \\ (0, 1) & t \in R_i \end{cases}$$

Here, $g_{\tilde{T}}(\cdot)$ maps tuples in R to single-bin histograms $(f(\tilde{T}), 1)$ and tuples in R_i to the semi-ring multiplicative identity. We illustrate the whole procedure using a running example.

EXAMPLE 7. *Consider relations R and R_1 in Figure 6a, with each tuple annotated by the histogram semi-ring for $g_{\tilde{T}}(t)$; for illustration purpose, $T = \tilde{T}$, $Z = \tilde{Z}$. To avoid materializing $R \bowtie R_1$ in Figure 6b, we first compute $\gamma_{\mathcal{J}}R$ and $\gamma_{\mathcal{J}}R_1$ by aggregating semi-ring annotations within each join key. For join key a_1 , the summation $g_{\tilde{T}}(t)$ is $(0, 1) \oplus (1, 1) \oplus (1, 1) = (0, 2), (3, 1)$ as shown in Figure 6c. Next, we combine sketches using the semi-ring \otimes operator – the product of the two semi-ring annotations across a_1 is $\{(0, 1), (1, 2)\} \otimes \{(0, 3)\} = \{(0 + 0, 1 \cdot 3), (0 + 1, 2 \cdot 3)\} = \{(0, 3), (1, 6)\}$. The final histogram $\gamma(\gamma_{\mathcal{J}}(R) \otimes \gamma_{\mathcal{J}}(R_1))$ (Figure 6d) equals the ground truth in Figure 6b.*

7.1.4 Estimating pdf for $\tilde{\epsilon}_Z$. $f(\tilde{T})$ is completely dependent on R , so it easily distribute through R_{\bowtie} . The same trick does not apply to calculating $f(\tilde{\epsilon}_Z)$, because it is unclear how R and R_1 contribute to $\tilde{\epsilon}_Z$. To factorize $\gamma_{f(\tilde{\epsilon}_Z)}(R_{\bowtie})$, we first decompose $f(\tilde{\epsilon}_Z)$ by leveraging its linearity, and the fact that $|\lfloor a + b \rfloor - (\lfloor a \rfloor + \lfloor b \rfloor)| \in \{0, 1\}$.

R			
\mathcal{J}	\mathcal{T}	$g_{\mathcal{T}}(t)$	$g_{\mathcal{T}\mathcal{Z}}(t)$
a_1	-0.65	$\{(0, 1)\}$	$\{(1, 1)\}$
a_1	0	$\{(1, 1)\}$	$\{(0, 1)\}$
a_1	1.3	$\{(1, 1)\}$	$\{(0, 1)\}$
a_2	-1.94	$\{(0, 1)\}$	$\{(1, 1)\}$

R_1			
\mathcal{J}	\mathcal{Z}	$g_{\mathcal{Z}}(t)$	$g_{\mathcal{T}\mathcal{Z}}(t)$
a_1	-0.49	$\{(0, 1)\}$	$\{(0, 1)\}$
a_1	0	$\{(0, 1)\}$	$\{(1, 1)\}$
a_1	0.81	$\{(0, 1)\}$	$\{(1, 1)\}$
a_2	-0.97	$\{(0, 1)\}$	$\{(0, 1)\}$

(a) Relations R and R_1 annotated with lift functions $g_{\mathcal{T}}(t)$ and $g_{\mathcal{T}\mathcal{Z}}(t)$ for $\gamma_{f(\mathcal{T})}(R_{\bowtie})$ and $\gamma_{f(\mathcal{EZ})}(R_{\bowtie})$, respectively.

R_{\bowtie}					
\mathcal{J}	\mathcal{T}	\mathcal{Z}	\mathcal{EZ}	$f(\mathcal{T})$	$f(\mathcal{EZ})$
a_1	-0.65	-0.49	-0.17	0	0
a_1	-0.65	0	0.33	0	1
a_1	-0.65	0.81	1.14	0	2
a_1	0	-0.49	-0.49	1	0
a_1	0	0	0	1	1
a_1	0	0.81	0.81	1	1
a_1	1.3	-0.49	-1.14	1	0
a_1	1.3	0	-0.65	1	0
a_1	1.3	0.81	0.16	1	1
a_2	-1.94	-0.97	0	0	1

$\gamma_{f(\mathcal{T})}(R_{\bowtie})$		$\gamma_{f(\mathcal{EZ})}(R_{\bowtie})$	
$f(\mathcal{T})$	CNT	$f(\mathcal{EZ})$	CNT
0	4	0	4
1	6	1	5
		2	1

(b) Naive join $R_{\bowtie} = R \bowtie_{\mathcal{J}} R_1$ including \mathcal{EZ} , $f(\mathcal{T})$ and $f(\mathcal{EZ})$, along with the query results.

$\gamma_{\mathcal{J}}(R)$			$\gamma_{\mathcal{J}}(R_1)$			$\gamma_{\mathcal{J}}(R) \otimes \gamma_{\mathcal{J}}(R_1)$		
\mathcal{J}	Annotation		\mathcal{J}	Annotation		\mathcal{J}	Annotation	
a_1	$\{(0, 1), (1, 2)\}$		a_1	$\{(0, 3)\}$		a_1	$\{(0, 3), (1, 6)\}$	
a_2	$\{(0, 1)\}$		a_2	$\{(0, 1)\}$		a_2	$\{(0, 1)\}$	

$\gamma_{f(\mathcal{T})}(R_{\bowtie})$			$\gamma_{f(\mathcal{EZ})}(R_{\bowtie})$			$\gamma_{f(\mathcal{T})}(R_{\bowtie}) \otimes \gamma_{f(\mathcal{EZ})}(R_{\bowtie})$		
\mathcal{J}	Annotation		\mathcal{J}	Annotation		\mathcal{J}	Annotation	
a_1	$\{(0, 2), (1, 1)\}$		a_1	$\{(0, 1), (1, 2)\}$		a_1	$\{(0, 2), (1, 5), (2, 2)\}$	
a_2	$\{(1, 1)\}$		a_2	$\{(0, 1)\}$		a_2	$\{(1, 1)\}$	

(c) Aggregation of semi-ring annotation as sketches and combining semi-rings between two sketches.

Est. $\gamma_{f(\mathcal{T})}(R_{\bowtie})$		Est. $\gamma_{f(\mathcal{EZ})}(R_{\bowtie})$	
\mathcal{I}	CNT	\mathcal{I}	CNT
0	4	0	2
1	6	1	6
		2	2

(d) Exact estimation of $\gamma_{f(\mathcal{T})}(R_{\bowtie})$ and an approximation of $\gamma_{f(\mathcal{EZ})}(R_{\bowtie})$.

Figure 6: Example of factorized histogram method.

For each $t \in R_{\bowtie}$, we rewrite $f(t[\mathcal{EZ}])$ as

$$\left\lfloor \frac{t[\mathcal{Z}] - \beta \cdot t[\mathcal{T}] - \alpha_{\mathcal{EZ}}}{w_{\mathcal{EZ}}} \right\rfloor \approx \left\lfloor \frac{t[\mathcal{Z}] - \alpha_{\mathcal{EZ}}}{w_{\mathcal{EZ}}} \right\rfloor - \left\lfloor \frac{\beta \cdot t[\mathcal{T}]}{w_{\mathcal{EZ}}} \right\rfloor$$

This decompose $f((t[\mathcal{Z}] - \beta \cdot t[\mathcal{T}])/\sigma_{\mathcal{EZ}})$ into two parts where the first term (highlighted in red) only involves $t[\mathcal{Z}]$ in R_1 and the second term (highlighted in blue) only involves $t[\mathcal{T}]$ in R , such that they can be computed locally. We define the lift function $g_{\mathcal{EZ}}(\cdot)$ as

$$g_{\mathcal{EZ}}(t) = \begin{cases} \left(\left\lfloor -\frac{\beta \cdot t[\mathcal{T}]}{w_{\mathcal{EZ}}} \right\rfloor, 1 \right) & t \in R \\ \left(\left\lfloor \frac{t[\mathcal{Z}] - \alpha_{\mathcal{EZ}}}{w_{\mathcal{EZ}}} \right\rfloor, 1 \right) & t \in R_1 \end{cases}$$

Here, $g_{\mathcal{EZ}}(\cdot)$ maps tuples in R and R_1 to their contributing factors in $f(\mathcal{EZ})$. We will use the same running example in Figure 6 to illustrate this procedure.

EXAMPLE 8. Consider again relations R and R_1 in Figure 6a, we first obtain $\beta = 0.5$ using the variance semi-ring, and lift each tuple in R and R_1 with the histogram semi-ring through $g_{\mathcal{EZ}}(t)$. Similar to Example 7, we aggregate the semi-ring annotations locally in R and

R_1 and merge them using the semi-ring \otimes operator as shown in the lower half of Figure 6c, the estimated histogram shown in Figure 6d approximates the ground truth in Figure 6b.

7.2 Convergence Analysis

As noted in previous sections, we use histograms to estimate differential entropy and mutual information. In Lemma 7.3 and Theorem 7.2, we demonstrate that the histogram-based approach yields unbiased estimates for both differential entropy and conditional mutual information, which naturally extend to mutual information.

The differential entropy (which can be negative) applies to continuous random variables, while entropy (always non-negative) applies only to discrete random variables. Both definitions can be unified via measure theory; so we only need to focus on proving results for entropy convergence. In particular, probability mass functions (pmf) can be translated to probability density functions (pdf) by using the counting measure (instead of the Lebesgue measure). We show that the histogram semiring-based entropy estimation, over mixed random vectors, converges to the true entropy:

LEMMA 7.3. Given a random vector (X, Y, Z) that contains discrete-continuous mixture random variables, with semiring histogram bins $B = \underline{B} \cup \bar{B}$ and the resulting discretized random vector (X', Y', Z') , where the bins in \bar{B} contain discrete data points (of which every dimension has a discrete value) and $\underline{B} = B \setminus \bar{B}$, we have

$$\lim_{b \rightarrow 0} H(X', Y', Z') = H(X, Y, Z),$$

where $b = \max_{B_j \in \underline{B}} (\mu(B_j))$, and μ is the product measure.

Furthermore, the histogram-based mutual information estimation also converges to the true mutual information:

THEOREM 7.2. Given a mixed random vector (X, Y, Z) ,

$$\lim_{b \rightarrow 0} \lim_{n \rightarrow \infty} \text{MI}^h(X; Y|Z) = \text{MI}(X; Y|Z)$$

almost surely, where n refers to the sample size and b refers to the maximum volume of the semiring histogram volumes for bins in \underline{B} . $\text{MI}^h(X; Y|Z)$ is the mutual information computed via histograms, and $\text{MI}(X; Y|Z)$ is the true mutual information.

Detailed discussion and proof of Lemma 7.3 and Theorem 7.2 can be found in the technical report. Note that Theorem 7.2 that applies to conditional mutual information also applies to the mutual information, since the latter is a special case of the former. Furthermore, it is straightforward to see that the results apply beyond mixed random vectors of size three.

8 EVALUATION

We now present evaluations of the quality and efficacy of Suna. We first evaluate Suna by constructing a large real-world data corpus consisting of datasets from NYC Open Data [7] and Kaggle [25] to determine if Suna can discover semantically meaningful confounders. Then, we generate synthetic datasets to quantitatively assess the adjustment sets discovered by Suna. Finally, we examine the effectiveness of our optimization proposed in Section 7.

Dataset	Query	Suna	MESA	HypDB	Correlation
SO	What is the effect of education level on salary?	Cost of Living & Rent Index	Unemployment Rate	Country	Rent Index
ELA	What is the effect of each school's extra credit performance score on students' ELA score?	Grade 6 Enrollment % Eligible for Free or Reduced Lunches	% of grads Performance Category Score Total Student Response Rate	School Overall Score ID Pct Level 3 and 4	School Overall Score
Ratio	What is the effect of each school's pupil-to-teacher ratio on student's ELA score?	Level 4: % % Students with Disabilities Minimum Class Size	Performance Category Score Math Proficiency ELA % Proficient	Percent HRA Eligible Average Class Size Pct Level 2	% Students with Disabilities
SAT	What is the effect of test takers numbers on SAT score?	# Safety Incidents Grade 9 Enrollment Total Regents #	Safety and Respect Score Graduation Rate # Ontrack at Year 1	Total Regents #	Total Regents #

Table 1: Report of discovered confounders for causal queries.

8.1 Real-world Experiments

We use real-world datasets to evaluate Suna’s ability to reason against spurious correlations. Our evaluation addresses two key questions: **Q1**: Does Suna discover semantically meaningful confounders? and **Q2**: Can Suna deliver causal query results with interactive latency? We cite existing studies in the social science literature to reason about semantic information. This is consistent the methodologies in prior work [39, 56] for evaluating the quality of causal queries over real-world datasets.

8.1.1 Data and Workload. We construct a large corpus of 346 datasets from NYC Open Data [7] and Kaggle [25]. Datasets from NYC open data contain 1462 numerical attributes and share a common District Borough Number column as join key, and datasets from Kaggle contain 91 attributes and share a common Country column. We construct 4 causal queries using the following datasets:

- **SO** [48] contains the results of annual surveys over 2023 and 2024 conducted by Stack Overflow that collect insights from developers worldwide. We convert the values in the *education level* column to numerical values that represent approximate years of formal schooling.
- **ELA** [4] integrates 2006-11 English Language Arts (ELA) test results with annual school progress reports [1] for each District Borough Number (DBN).
- **Ratio** [3] contains 2013-18 ELA test results over DBNs integrated with the annual pupil-to-teacher ratio report [5] for each DBN.
- **SAT** [2] contains 2012 college-bound seniors’ mean SAT scores over each DBN.

For each query, we preprocess all datasets in the data corpus by aggregating attributes with respect to the join key, and impute missing entries with the mean of the column.

8.1.2 Baselines. We compare Suna with the following approaches: (1) MESA [56] discovers attributes that explain confounding bias between two variables with no causal relationship, (2) HypDB [39] is a confounder detection system that discovers confounders for arbitrary treatment and outcome variables, and (3) Correlation identifies the attribute most strongly correlated with the treatment variable. Both MESA and HypDB require a single table as input, so they require a preprocessing step that materializes the join across all join-compatible datasets. Moreover, HypDB only accepts categorical attributes, necessitating a binning procedure for numerical attributes. However, it is optimized for binary features; with more

than two bins per numerical attribute, it fails to terminate within 10 hours on each query. Consequently, we convert all attributes into binary features to run HypDB efficiently. Although this binning transformation discards some information, it preserves the underlying causal diagram.

8.1.3 Quality analysis. In this section, we aim to answer **Q1**. We report the confounders discovered by different methods for each query in Table 1, in cases where Suna discovers more than 3 confounders, we select the 3 most influential confounders toward the treatment effect. Next, we analyze the quality of the discovered confounders by referencing existing studies. We highlight confounders complying with existing studies in [blue](#) in Table 1. Suna discovers semantically meaningful confounders for each causal query.

SO. We examine potential confounders for causal analysis between years of formal schooling (treatment) and yearly salary (outcome). An economic study [35] shows that social factors—such as higher living costs in urban areas—can drive both greater educational attainment and increased salary. Both Suna and Correlation identify relevant confounders consistent with these findings.

ELA. We next investigate confounders for causal analysis between the additional score a school receives in its progress report (treatment), which represents extra points awarded beyond the core performance metrics, and the school’s average ELA score (outcome). Research suggests socioeconomic status is a key confounder. For this query, only Suna discovers *% poverty*.

Ratio. Then, we explore potential confounders for causal analysis between the pupil-to-teacher ratio (treatment) on the school’s average ELA score (outcome). Prior work indicates that students’ socio-economic status impacts both class sizes and performance. In this query, Suna and Correlation identify *% Students with Disabilities*, while HypDB detects *Percent HRA Eligible* (a proxy for the socioeconomic needs of the school community), all in line with existing literature.

SAT. Finally, we examine confounders when studying how the number of SAT test takes (treatment) affect the average math score (outcome) for each high school. Previous research [30, 38] notes that school size and environment can shape test participation and performance. All methods (Suna, MESA, HypDB, and Correlation) identify plausible confounders.

8.1.4 Runtime analysis. We now evaluate the scalability of Suna. Figure 7a shows the cumulative runtime for Suna, MESA, and

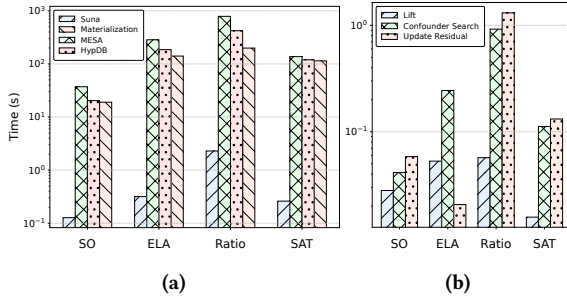


Figure 7: (a) Runtime performance of Suna vs other baselines for real-world causal queries. The x-axis includes all baselines for each query. The y-axis shows runtime. Suna is ≥ 2 orders of magnitude faster than baselines for causal queries. (b) Runtime breakdown of Suna’s algorithmic components for each query.

HypDB, along with the time required to materialize joins for each causal query. Importantly, Suna is ≥ 2 orders of magnitude faster than all baselines for every causal query. In fact, Suna’s completion times of the queries are orders of magnitude faster than MESA and HypDB’s join completion time. This demonstrates the efficiency of Suna’s optimizations. Figure 7b decomposes Suna’s end-to-end runtime into its algorithmic components, including transforming the input dataset into sketches (lift), iteratively searching for confounders with sketches (confounder search), and incrementally maintaining the sketches across iterations (update residual). For all queries, none of these components is a significant runtime bottleneck, as they all finish within ~ 1 s, achieving interactive latency.

8.1.5 Takeaways: *Suna discovers semantically meaningful confounders as validated by existing studies. With optimizations discussed in Section 7, Suna runs orders of magnitude (>100 times) faster than existing confounder discovery systems.*

8.2 Synthetic Experiments

We further assess Suna’s accuracy on synthetic datasets to address three key questions. **Q1:** Does Suna discover an admissible adjustment set that accurately estimates the ATE? **Q2:** How does the total number of variables in the causal diagram affect Suna’s estimate of the average treatment effect, both with and without irrelevant variables as noise? **Q3:** How well do our approximations in Section 7.1 estimate $\Pr(\epsilon_Z)$ and $\Pr(\epsilon_T)$, and how this approximation implies to the estimation of $H(\epsilon_Z)$ and $H(\epsilon_T)$.

8.2.1 Setup. We construct a causal diagram using synthetic datasets following the LiNGAM model. The coefficients of linear mechanisms are drawn randomly within the range of $[1, 2]$; additive noises are sampled from a Uniform distribution $U(0, 1)$. To avoid overflow, we normalize the coefficients that determine each variable by dividing them by the number of parents of that variable and rounding the data to two digits. To construct causal queries, we generate a random pair of treatment and outcome $Q = (T, O)$ where T is a non-descendant of O (i.e., $T \notin \text{ND}(O)$). We generate a relation \mathcal{R} with n tuples and m variables. Then, we partition \mathcal{R} vertically where the input dataset R contains $[\mathcal{J}, T, O]$, and the remaining columns are partitioned into $m - 2$ relations, each contains a single attribute. To support joins, we assign a primary key




Figure 8: MSE and r^2 (normalizing for variations in the ground truth ATE) of the estimated ATE by Suna and different baselines evaluated on over $m \in \{10, 50, 100, 500, 1M\}$. Suna consistently achieves high accuracy across different number of variables, while ND does not scale to large $|A|$.

\mathcal{J} to \mathcal{R} , and distribute it over each of the $m - 2$ relations, \mathcal{R} can be recovered by joining all relations through one-to-one joins over \mathcal{J} .

8.2.2 Accuracy Evaluation. We study **Q1** and **Q2** by randomly generating 100 causal queries with $m \in \{10, 50, 100, 500, 1M\}$. To answer **Q1**, we use the discovered set of confounders \mathbf{Z} to train a linear regression model $O \sim \beta\mathbf{Z} + \alpha T$, and compare α against the ground-truth ATE. To evaluate **Q2**—Suna’s resilience to irrelevant variables—we set $m = 1M$ to construct a causal diagram over 500 variables and fill the remaining ones with random noises. For baselines, we compare Suna with MESA, HypDB, and the simple extension we proposed in Section 6.1, referred to as ND. Since ND does not scale to moderately large m (the topological sort and early stop do not terminate over an hour for $m = 500$ for each causal query), we evaluate ND by using the set of all non-descendants of T obtained from the ground truth causal diagram as the adjustment set. This approach assumes ND will always estimate $\text{ND}(T)$ accurately, an assumption that may not always hold.

Figure 8a reports the squared error (MSE) between the estimated ATE versus the ground truth. To account for the difference in magnitude of the ground truth causal effect, we report the r^2 score in Figure 8b. Suna and ND has similar performance – lowest MSE and highest r^2 score of > 0.99 for each m . On the other hand, MESA and HypDB both fail to discover admissible adjustment sets with r^2 score degrades for larger m . This is because MESA assumes no causal effect between the treatment and outcome, which is easily violated for larger m . When this assumption is violated, MESA may mistakenly select mediators as confounders that leads to wrong adjustment sets. On the other hand, HypDB’s reliance on binning leads to significant information loss, thus hindering its ability to find valid confounders. Interestingly, although ND guarantees to find an admissible adjustment set in theory under infinite sample sizes, it estimates the average treatment effect poorly for $m = 1M$. This is due to estimation error caused by large adjustment sets.

8.2.3 Scalability. We further study **Q1** and **Q2** by examining how the size of the adjustment set and the runtime varies as the number of variables increases. The settings remain the same as Section 8.2.2. Figure 9a shows that Suna discovers a smaller (less than one third of the size than ND), yet still admissible confounder set across all numbers of variables. This makes Suna a more interpretable system for analysts to use, especially when they seek to analyze causal relations in complex scenarios when the number of variables is




Figure 9: Adjustment set size $|Z|$ discovered by Suna and size of $ND(T)$ in \mathcal{G} and Suna’s runtime to iteratively construct Z . The adjustment set discovered by Suna is $>3\times$ smaller.

large. Figure 9b shows that the mean runtime over 100 queries grows slightly as the number of variables increases, due to the preprocessing step to filter out irrelevant variables. By leveraging factorized learning techniques and GPU acceleration (following SET [31]), we manage to keep the overhead minimal ($\sim 3s$ for one million variables), even at a scale where most baselines would time out or take more than days to complete.

8.2.4 Histogram Semi-ring Approximation Evaluation. We investigate **Q3** by examining the divergence between the approximate histogram and the exact histogram (obtained via join materialization), as well as the resulting impact on entropy estimation.


We use the same setup from Section 8.2.1 for data generation. For robustness, we take 10 attributes $\{Z_1, \dots, Z_{10}\}$ and conduct BCD for each pair of (T, Z_i) . Due to symmetry in estimating $\Pr(\tilde{\epsilon}_Z)$ and $\Pr(\tilde{\epsilon}_T)$, we focus our study on the approximation of $\Pr(\tilde{\epsilon}_Z)$ and how such approximation affects $H(\tilde{\epsilon}_Z)$ - a core estimate in BCD-using methods described in Section 7.1. We measure the approximation accuracy of $\Pr(\tilde{\epsilon}_Z)$ using KL divergence from the exact histogram computed by materializing joins, and then quantify how this affects entropy estimates by computing the ℓ_1 distance between the approximated entropy and the ground truth entropy of each $\tilde{\epsilon}_{Z_i}$. The sample size n is varied from $\{100, 1000, 5000, 10000, 50000\}$.

Figure 10a reports that the mean KL-divergence between the approximated and exact histograms converges to 0 as the number of samples increases, all (T, Z_i) pairs. Convergence is also observed in Figure 10b, which shows that as the number of samples increases, the ℓ_1 distance between the approximate and exact histograms approaches 0. In fact, both the KL divergence and ℓ_1 distance fall below a negligible 0.01 at a small sample size ($n = 5000$).

8.2.5 Takeaways. Suna prunes away irrelevant variables and identifies high-quality confounders to accurately answer causal queries in just a few seconds, scaling to 1M variables. Further, empirical evaluations show that the histogram estimation reliably converges as the number of samples increases. Combined with the theoretical analysis in Section 7.2, these results validate our approach in Section 7.

9 DISCUSSION

Discussion. Our main contribution is an algorithm that efficiently discovers confounders in an unobserved causal DAG from a relational data repository—all without materializing joins. In our current implementation, we leverage LiNGAM to analyze linear relationships among numerical attributes under the *causal sufficiency assumption*. Our approach lays a solid foundation for extending



(a) KL Divergence of the approximated versus the exact histograms converges to 0 as the samples n increases. (b) ℓ_1 distance of the entropy estimated using the approximate histograms versus the exact histograms converges to 0 as n increases.

Figure 10: Comparison of KL Divergence and entropy estimates as sample size n increases.

towards more general and complex causal settings. In particular, moving beyond linear models and incorporating latent confounders only requires the capability to perform bivariate causal discovery on the projected causal diagram that includes a variable pair (T, Z) from the current adjustment set.

Limitation and Future Directions. One important yet out-of-scope aspect we did not address is the complexity of handling data quality. In real-world scenarios, missing or contaminated data is common. However, standard data cleaning techniques such as data imputation with mean values distort the underlying causal assumptions. For example, removing entire tuples containing the missing entry introduces artificial selection bias and affects the accuracy of downstream causal analysis. Future work could extend Suna in several promising directions. One avenue is to explore more complex parametric causal mechanisms (e.g. non-linear) capable of handling arbitrary data types, or design a hybrid approach based on assumptions over subsets of datasets. Another direction is to identify priority for data curation over a fixed-sized set of datasets and columns to best improve the accuracy of a causal query.

10 CONCLUSIONS

Suna is a confounder discovery system that searches large corpora to find datasets to improve causal inference task estimation. We prove the connection between bivariate causal discovery, a key building block in parametric causal discovery that exploits the asymmetry between causes and effects, and confounder discovery. This theorem lets Suna directly search for causal confounders without the need for a full causal DAG. To improve scalability, Suna designs a novel data structure inspired by factorized learning to avoid materializing joins and implements a GPU-compatible system to further scale up performance. Suna achieves an r^2 score of >0.99 over a causal diagram of 500 variables where all existing confounder discovery systems fail, and scales to repositories with 1M variables while answering queries within a few seconds.

ACKNOWLEDGMENTS

This research received funding from multiple sources, including National Science Foundation grants (NSF #1845638, #1740305, #2008295, #2106197, #2103794), Simons Foundation (965342, D.A.), as well as from Amazon, Google, Adobe, CAIT, and Intellect Design. The views and conclusions presented here are those of the authors and do not represent the official positions of the funding organizations.

REFERENCES

- [1] [n.d.]. 2007-2008-School-Progress-Reports-All-Schools. <https://data.cityofnewyork.us/Education/2007-2008-School-Progress-Reports-All-Schools/dj4e-3xrn>.
- [2] [n.d.]. 2012 SAT Results. <https://data.cityofnewyork.us/Education/2012-SAT-Results/f9bf-2cp4>.
- [3] [n.d.]. 2013 - 2018 School ELA REsults. <https://data.cityofnewyork.us/Education/2013-2018-School-ELA-REsults/qkpp-pbi8>.
- [4] [n.d.]. 2014-15 To 2016-17 School- Level NYC Regents Report For All Variables. <https://data.cityofnewyork.us/Education/2014-15-To-2016-17-School-Level-NYC-Regents-Report/csp5-2ne9/>.
- [5] [n.d.]. 2018 - 2019 Class Size Pupil to Teacher Ratio. <https://data.cityofnewyork.us/Education/2018-2019-Class-Size-Pupil-to-Teacher-Ratio/axb2-9jkb>.
- [6] [n.d.]. SUNA Technical Report. https://anonymous.4open.science/r/causal_search-3CD5/Suna.pdf.
- [7] 2022. NYC Open Data. <https://opendata.cityofnewyork.us/>.
- [8] Daniel Alabi, Sainyam Galhotra, Shagufa Mehnaz, Zeyu Song, and Eugene Wu. 2025. Privacy and Security in Distributed Data Markets. In *Companion of the 2025 International Conference on Management of Data (SIGMOD/PODS '25)*.
- [9] Ioana Bica, Ahmed Alaa, and Mihaela Van Der Schaar. 2020. Time series deconfounder: Estimating treatment effects over time in the presence of hidden confounders. In *International conference on machine learning*. PMLR, 884–895.
- [10] Nadiia Chepurko, Ryan Marcus, Emanuel Zraggen, Raul Castro Fernandez, Tim Kraska, and David Karger. 2020. ARDA: automatic relational data augmentation for machine learning. *arXiv preprint arXiv:2003.09758* (2020).
- [11] François Chollet. 2019. On the measure of intelligence. *arXiv preprint arXiv:1911.01547* (2019).
- [12] Rama Cont. 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative finance* 1, 2 (2001), 223.
- [13] G. Darmon. 1953. Analyse générale des liaisons stochastiques: étude particulière de l'analyse factorielle linéaire. *Revue de l'Institut International de Statistique / Review of the International Statistical Institute* 21, 1/2 (1953), 2–8. <http://www.jstor.org/stable/1401511>
- [14] Sainyam Galhotra, Yue Gong, and Raul Castro Fernandez. 2023. METAM: Goal-Oriented Data Discovery. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 2780–2793.
- [15] Yue Gong, Sainyam Galhotra, and Raul Castro Fernandez. 2024. Nexus: Correlation Discovery over Collections of Spatio-Temporal Tabular Data. *Proceedings of the ACM on Management of Data* 2, 3 (2024), 1–28.
- [16] Todd J Green, Grigoris Karvounarakis, and Val Tannen. 2007. Provenance semirings. In *Proceedings of the twenty-sixth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 31–40.
- [17] Sander Greenland, Judea Pearl, and James M Robins. 1999. Confounding and collapsibility in causal inference. *Statistical science* 14, 1 (1999), 29–46.
- [18] Tobias Hatt and Stefan Feuerriegel. 2024. Sequential deconfounding for causal inference with unobserved confounders. In *Causal Learning and Reasoning*. PMLR, 934–956.
- [19] Patrik Hoyer, Dominik Janzing, Joris M Mooij, Jonas Peters, and Bernhard Schölkopf. 2008. Nonlinear causal discovery with additive noise models. *Advances in neural information processing systems* 21 (2008).
- [20] Zezhou Huang, Jiaxiang Liu, Daniel Alabi, Raul Castro Fernandez, and Eugene Wu. 2023. Saibot: A Differentially Private Data Search Platform. *arXiv preprint arXiv:2307.00432* (2023).
- [21] Zezhou Huang, Jiaxiang Liu, Haonan Wang, and Eugene Wu. 2023. The Fast and the Private: Task-based Dataset Search. *arXiv preprint arXiv:2308.05637* (2023).
- [22] Paul Hünermund, Jermain Kaminski, and Carla Schmitt. 2022. Causal machine learning and business decision making. *Available at SSRN 3867326* (2022).
- [23] Aapo Hyvärinen and Stephen M Smith. 2013. Pairwise likelihood ratios for estimation of non-Gaussian structural equation models. *The Journal of Machine Learning Research* 14, 1 (2013), 111–152.
- [24] Amin Jaber, Adele Ribeiro, Jiji Zhang, and Elias Bareinboim. 2022. Causal identification under Markov equivalence: calculus, algorithm, and completeness. *Advances in Neural Information Processing Systems* 35 (2022), 3679–3690.
- [25] Kaggle. 2024. Kaggle Datasets Repository. <https://www.kaggle.com/datasets> Accessed: 2024-10-15.
- [26] Aamod Khatiwada, Roei Shraga, Wolfgang Gatterbauer, and Renée J Miller. 2022. Integrating data lake tables. *Proceedings of the VLDB Endowment* 16, 4 (2022), 932–945.
- [27] Aamod Khatiwada, Roei Shraga, and Renée J Miller. 2023. DIALITE: discover, align and integrate open data tables. In *Companion of the 2023 International Conference on Management of Data*. 187–190.
- [28] Alexander Kraskov, Harald Stögbauer, and Peter Grassberger. 2004. Estimating mutual information. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 69, 6 (2004), 066138.
- [29] Sanghack Lee and Vasant G Honavar. 2019. Towards Robust Relational Causal Discovery. In *Conference on Uncertainty in Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:197660683>
- [30] Valerie E Lee and Julia B Smith. 1997. High school size: Which works best and for whom? *Educational evaluation and policy analysis* 19, 3 (1997), 205–227.
- [31] Jiaxiang Liu, Zezhou Huang, and Eugene Wu. 2024. SET: Searching Effective Supervised Learning Augmentations in Large Tabular Data Repositories. In *Proceedings of the Conference on Governance, Understanding and Integration of Data for Effective and Responsible AI*. 26–31.
- [32] Mark Maier. 2014. Causal Discovery for Relational Domains: Representation, Reasoning, and Learning. <https://api.semanticscholar.org/CorpusID:11057134>
- [33] Marc E. Maier, Brian J. Taylor, Hüseyin Oktay, and David D. Jensen. 2010. Learning Causal Models of Relational Domains. *Proceedings of the AAAI Conference on Artificial Intelligence* (2010). <https://api.semanticscholar.org/CorpusID:12408451>
- [34] Sarah Mameche, Jilles Vreeken, and David Kaltenpoth. 2024. Identifying Confounding from Causal Mechanism Shifts. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research)*, Sanjoy Dasgupta, Stephan Mandt, and Yingzhen Li (Eds.), Vol. 238. PMLR, 4897–4905.
- [35] Enrico Moretti. 2013. Real wage inequality. *American Economic Journal: Applied Economics* 5, 1 (2013), 65–103.
- [36] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [37] Judea Pearl. 2010. An introduction to causal inference. *The international journal of biostatistics* 6, 2 (2010).
- [38] Sean F Reardon, Demetra Kalogrides, and Kenneth Shores. 2019. The geography of racial/ethnic test score gaps. *Amer. J. Sociology* 124, 4 (2019), 1164–1221.
- [39] 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.
- [40] Babak Salimi, Harsh Parikh, Moe Kayali, Lise Getoor, Sudeepa Roy, and Dan Suciu. 2020. Causal relational learning. In *Proceedings of the 2020 ACM SIGMOD international conference on management of data*. 241–256.
- [41] Aécio Santos, Aline Bessa, Christopher Musco, and Juliana Freire. 2022. A sketch-based index for correlated dataset search. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2928–2941.
- [42] David W Scott. 1979. On optimal and data-based histograms. *Biometrika* 66, 3 (1979), 605–610.
- [43] Rajen D Shah and Jonas Peters. 2020. The hardness of conditional independence testing and the generalised covariance measure. (2020).
- [44] 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).
- [45] Shohei Shimizu, Takanori Inazumi, Yasuhiro Sogawa, Aapo Hyvärinen, Yoshinobu Kawahara, Takashi Washio, Patrik O Hoyer, Kenneth Bollen, and Patrik Hoyer. 2011. DirectLiNGAM: A direct method for learning a linear non-Gaussian structural equation model. *Journal of Machine Learning Research-JMLR* 12, Apr (2011), 1225–1248.
- [46] Peter Spirtes. 2001. An anytime algorithm for causal inference. In *International Workshop on Artificial Intelligence and Statistics*. PMLR, 278–285.
- [47] Peter Spirtes, Clark Glymour, and Richard Scheines. 2001. *Causation, prediction, and search*. MIT press.
- [48] Stack Overflow. 2024. Stack Overflow Developer Survey. <https://insights.stackoverflow.com/survey/> Accessed: 2024-10-11.
- [49] Hui Yie Teh, Andreas W Kempa-Liehr, and Kevin I-Kai Wang. 2020. Sensor data quality: A systematic review. *Journal of Big Data* 7, 1 (2020), 11.
- [50] Jin Tian and Judea Pearl. 2002. A general identification condition for causal effects. In *Aaai/iaai*. 567–573.
- [51] Clifford H Wagner. 1982. Simpson's paradox in real life. *The American Statistician* 36, 1 (1982), 46–48.
- [52] Yixin Wang and David M Blei. 2019. The blessings of multiple causes. *J. Amer. Statist. Assoc.* 114, 528 (2019), 1574–1596.
- [53] Y Samuel Wang and Mathias Drton. 2023. Causal discovery with unobserved confounding and non-Gaussian data. *Journal of Machine Learning Research* 24, 271 (2023), 1–61.
- [54] David Watson and Ricardo Silva. 2022. Causal discovery under a confounder blanket. In *Conference on Uncertainty in Artificial Intelligence*. <https://api.semanticscholar.org/CorpusID:248721868>
- [55] Brit Youngmann, Michael Cafarella, Amir Gilad, and Sudeepa Roy. 2024. Summarized Causal Explanations For Aggregate Views. *Proceedings of the ACM on Management of Data* 2, 1 (2024), 1–27.
- [56] Brit Youngmann, Michael Cafarella, Yuval Moskovitch, and Babak Salimi. 2023. On Explaining Confounding Bias. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 1846–1859.
- [57] Brit Youngmann, Michael Cafarella, Babak Salimi, and Anna Zeng. 2023. Causal Data Integration. *arXiv preprint arXiv:2305.08741* (2023).
- [58] Erkang Zhu, Dong Deng, Fatemeh Nargesian, and Renée J Miller. 2019. Josie: Overlap set similarity search for finding joinable tables in data lakes. In *Proceedings of the 2019 International Conference on Management of Data*. 847–864.