Fast and Reliable Missing Data Contingency Analysis with Predicate-Constraints

Xi Liang
University of Chicago
xiliang@uchicago.edu

Zechao Shang
University of Chicago
zcshang@cs.uchicago.edu

Sanjay Krishnan
University of Chicago
skr@cs.uchicago.edu

Aaron J. Elmore
University of Chicago
aelmore@cs.uchicago.edu

Michael J. Franklin
University of Chicago
mjfranklin@uchicago.edu

ABSTRACT

Today, data analysts largely rely on intuition to determine whether missing or withheld rows of a dataset significantly affect their analyses. We propose a framework that can produce automatic contingency analysis, i.e., the range of values an aggregate SQL query could take, under formal constraints describing the variation and frequency of missing data tuples. We describe how to process SUM, COUNT, AVG, MIN, and MAX queries in these conditions resulting in hard error bounds with testable constraints. We propose an optimization algorithm based on an integer program that reconciles a set of such constraints, even if they are overlapping, conflicting, or unsatisfiable, into such bounds. Our experiments on real-world datasets against several statistical imputation and inference baselines show that statistical techniques can have a deceptively high error rate that is often unpredictable. In contrast, our framework offers hard bounds that are guaranteed to hold if the constraints are not violated. In spite of these hard bounds, we show competitive accuracy to statistical baselines.

CCS CONCEPTS

- **Information systems** → **Database query processing**;
- **Theory of computation** → **Sketching and sampling**.

KEYWORDS

approximate query processing; data sketching; missing data; aggregate estimation; contingency analysis

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 ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMOD'20, June 14–19, 2020, Portland, OR, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6735-6/20/06...\$15.00

<https://doi.org/10.1145/3318464.3389785>

ACM Reference Format:

Xi Liang, Zechao Shang, Sanjay Krishnan, Aaron J. Elmore, and Michael J. Franklin. 2020. Fast and Reliable Missing Data Contingency Analysis with Predicate-Constraints. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD'20)*, June 14–19, 2020, Portland, OR, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3318464.3389785>

1 INTRODUCTION

The data stored in a database may differ from real-world truth in terms of both completeness and content. Such issues can arise due to data entry errors, inexact data integration, or software bugs [5]. As real-world data are rarely perfectly clean or complete, data scientists have to reason how potential sources of error may affect their analyses. Communicating these error modes and quantifying the uncertainty they introduce into a particular analysis is arguably as important as timely execution [14].

For example, suppose a data analyst has collected data from a temperature sensor over the span of several days. She is interested in computing the number of times that the sensor exceeded a temperature threshold. The data are stored in 10 partitions; one of which failed to load into the database due to parsing errors. The analyst can still run her query on the 9 available partitions, however, she needs to determine whether the loss of that partition may affect her conclusions.

Today, analysts largely rely on intuition to reason about such scenarios. The analyst in our example needs to make a judgment about whether the lost partition correlates with the attributes of interest, such as temperature, in any way. Such intuitive judgments, while commonplace, are very problematic because they are based on assumptions that are often not formally encoded in any code or documentation. Simply reporting an extrapolated result does not convey any measure of confidence in how (in)accurate the result might be, and could hide the fact that some of the data were not used.

This paper defines a framework for specifying beliefs about the missing rows in a dataset in a logical constraint language and an algorithm for computing a range of values

an aggregate query can take under those constraints (hereafter called a result range).¹ This framework, which we call the Predicate-Constraint (PC) framework, facilitates several desirable outcomes: (1) the constraints are efficiently testable on historical data to determine whether or not they held true in the past, (2) the result range is calculated deterministically and guaranteed to bound the results if the constraints hold true in the future, (3) the framework can reconcile interacting, overlapping, or conflicting constraints by enforcing the most restrictive ones, and (4) the framework makes no distributional assumptions about past data resembling future data other than what is specified in the constraints. With this framework, an analyst can automatically produce a contingency analysis, i.e., the range of values the aggregate could take, under formally described assumptions about the nature of the unseen data. Since the assumptions are formally described and completely determine the result ranges, they can be checked, versioned, and tested just like any other analysis code—ultimately facilitating a more reproducible analysis methodology. The constraints themselves, called Predicate-Constraints, are logical statements that constrain the range of values that a set of rows can take and the number of such rows within a predicate. We show that deriving the result ranges for a single “closed” predicate-constraint set can be posed as a mixed-integer linear program (MILP). The solver itself contains a number of novel optimizations, which we contribute such as pruning of unsatisfiable search paths.

To the best of our knowledge, a direct competitor framework does not exist. While there is a rich history of what-if analysis [8] and how-to analysis [17], which characterize a database’s behavior under hypothetical updates, analyzing the effects of constrained missing rows on aggregate queries has been not been extensively studied. The closest framework is the m-table framework [24], which has a similar expressiveness but no algorithm for computing aggregate result ranges. Likewise, some of the work in data privacy solves a simplified version of the problem where there are no overlapping constraints [26]. In summary, we contribute:

- (1) A formal framework for contingency analysis over missing or withheld rows of data, where users specify constraints about the frequency and variation of the missing rows.
- (2) An optimization algorithm that reconciles a set of such constraints, even if they are overlapping, conflicting, or unsatisfiable, into a range of possible values that SUM, COUNT, AVG, MIN, and MAX SQL queries can take.

- (3) Optimization that improve accuracy and/or optimization performance such as pruning unsatisfiable constraint paths.

This paper is an abridged version of a full technical report² that includes several important extensions to queries that involve joins and additional experiments. In particular, we show links to the Fractional Edge Cover bounds employed in the analysis of worst-case optimal joins [19].

2 BACKGROUND

In this paper, we consider the following user interface. The system is asked to answer SQL aggregate queries over a table with a number of missing rows. The user provides a set of constraints (called predicate-constraints) that describes how many such rows could be missing and a range of possible attribute values those missing rows could take. The system should integrate these constraints into its query processing and compute the maximal range of results (aggregate values) consistent with those constraints.

All of our queries are of the form:

```
SELECT agg(attr)
FROM R
WHERE ....
GROUP BY ....
```

We consider SUM, COUNT, AVG, MIN, and MAX aggregates with predicates. Because GROUP-BY clause can be considered as a union of such queries without GROUP-BY. In the rest of the paper, we focus on queries without GROUP-BY clause.

2.1 Example Application

Consider a simplified sales transaction table of just three attributes:

```
Sales(utc, branch, price)
Nov-01 10:20, New York, 3.02
Nov-01 10:21, Chicago, 6.71
...
Nov-16 6:42, Trenton, 18.99
```

Over this dataset, a data analyst is interested in calculating the total number of sales:

```
SELECT SUM(price)
FROM Order
```

Suppose that between November 10 and November 13 there was a network outage that caused data from the New York and Chicago branches to be lost. How can we assess the effect of the missing data on the query result?

¹We use this term to differentiate a deterministic range with probabilistic confidence intervals.

²The full report can be found at <http://arxiv.org/abs/2004.04139>

Figure 1: Simple Extrapolation could introduce significant error when missing rows are correlated (e.g., tend to have the highest values).

Simple Extrapolation: One option is to simply extrapolate the SUM for the missing days based on the data that is available. While this solution is simple to implement, it leads to subtle assumptions that can make the analysis very misleading. Extrapolation assumes that the missing data comes from roughly the same distribution as the data that is present. Figure 1 shows an experiment on one of our experimental datasets. We vary the amount of missing data in a way that is correlated with a SUM query. Even if the exact amount of missing data is known, the estimate becomes increasingly error-prone. More subtly, extrapolation returns a single result without a good measure of uncertainty—there is no way to know how wrong an answer might be.

Better Extrapolation: A smarter approach might be to build a probabilistic model that identifies trends and correlations in the data (e.g., a Gaussian Mixture Model that identifies weekly patterns) and use that model to extrapolate. If a user mis-specifies her belief in the data distribution or sampling process, any inference would be equally fallible as simple extrapolation. The probabilistic nature of the inference also makes potential failure models hard to interpret—errors could arise due to modeling, sampling, or even approximation error in the model fitting process.

Our Approach: In light of these issues, we propose a fully deterministic model for quantifying the uncertainty in a query result due to missing data. Like the probabilistic approach, we require that the user specify her belief about the distribution of missing data. Rather than specifying these beliefs in terms of probability distributions, she specifies the beliefs in terms of hard constraints.

For example, there are no more than 300 sales each day in Chicago. Or, the most expensive product costs 149.99 and no more than 5 are sold each day. We collect a set of such constraints and solve an optimization problem to find the maximal sum possible for all missing data instances that satisfy the constraints. This formalism acts as a programming framework that can be used to test the effects of different scenarios. It crucially enforces that there are testable constraints

that are recorded during the decision making process. We will use this as an example throughout the paper.

3 PREDICATE-CONSTRAINTS

The formal problem setting is defined as follows. Let R be a relation with a known “certain” partition denoted by R^* and unknown “missing” partition $R^?$, where $R = R^* \cup R^?$. The user defines a set of constraints $\pi_1, \dots, \pi_n$ over the possible tuples that could be in $R^?$ and their multiplicity. The computational problem is to derive a result range, namely, the min and max value that one of the supported aggregate queries given all possible instances of $R^?$ that are valid under the constraints. This section describes our language for expressing constraints. Suppose, this relation is over the attributes $A = \{a_1, \dots, a_p\}$. The domain of each attribute a_i is a set denoted by $\text{dom}(a_i)$.

3.1 Predicate-Constraint

If $R^?$ could be arbitrary, there is clearly no way to bound its effect on an aggregate query. Predicate-constraints restrict the values of tuples contained in $R^?$ and the total cardinality of such tuples. A single predicate-constraint defines a condition over the $R^?$, for tuples that satisfy the condition a range of attribute values, and the minimum and maximum occurrence of this predicate. As an example predicate-constraint in the sales dataset described in the previous section, “the most expensive product in Chicago costs 149.99 and no more than 5 are sold”. This statement has a predicate (“in Chicago”), a constraint over the values (“cost at most 149.99”), and occurrence constraint (“no more than 5”). We will show how to express systems of such constraints in our framework.

Predicate: A predicate ψ is a Boolean function that maps each possible row to a True and False value $\psi : \mathcal{D} \mapsto \mathbb{Z}_2$. For efficient implementation, we focus on predicates that are conjunctions of ranges and inequalities. This restriction simplifies satisfiability testing, which is an important step in our algorithms introduced in Section 4.1.

Value Constraint: A value constraint specifies a set of ranges that each attribute can take on. A range of the attribute a_i is defined as a row of two elements $(l, h) \in \text{dom}(a_i)$ where $l \leq h$. A value constraint v is a set of ranges for each of the p attributes:

$$v = \{(l_1, h_1), \dots, (l_p, h_p)\}$$

v defines a Boolean function as above $v : \mathcal{D} \mapsto \mathbb{Z}_2$ that checks whether a row satisfies all the specified ranges. Since we focus on bounding aggregates it is sufficient to assume that the attribute ranges are over numerical attributes.

Frequency Constraint: Associated with each predicate is additionally a frequency constraint. This bounds the number

of times that rows with the predicate appear. The frequency constraint is denoted as $\kappa = (k_l, k_u)$. k_l and k_u specify a range of the frequency constraint, i.e., there are at least k_l rows and at most k_u rows that satisfy the predicates in $R^?$. Of course, k_l, k_u must be non-negative numbers and $k_l \leq k_u$.

Predicate Constraint: A predicate-constraint is a three-tuple of these constituent pieces $\pi = (\psi, \nu, \kappa)$, a predicate, a set of value constraints, and a frequency constraint. The goal of π is to define constraints on relations that satisfy the above schema. Essentially a predicate constraint says if R is a relational instance that satisfies the predicate-constraint π “**For all rows that satisfy the predicate ψ , the values are bounded by ν and the number of such rows is bounded by κ** ”. Formally, we get the definition below.

Definition 3.1 (Predicate Constraint). A predicate constraint is a three-tuple consisting of a predicate, a value constraint, and a frequency constraint $\pi = (\psi, \nu, \kappa)$. Let R be a relational instance over the attributes A . R satisfies a predicate constraint denoted by $R \models \pi$ if:

$$(\forall r \in R : \psi(r) \implies \nu(r)) \wedge k_l \leq |\{r \in R : \psi(r)\}| \leq k_u$$

Let us now consider several examples of predicate constraints using the sales data example in the previous section on two days worth of missing rows. For example, if we want to express the constraint “the most expensive product in Chicago costs 149.99 and no more than 5 are sold”:

$$c_1 : (\text{branch} = \text{'Chicago'}) \implies (0.00 \leq \text{price} \leq 149.99), (0, 5)$$

In this example, the predicate is $(\text{branch} = \text{'Chicago'})$, the value constraint is $(0.00 \leq \text{price} \leq 149.99)$ and the frequency constraint is $(0, 5)$. If the aforementioned predicate constraint is describing a sales dataset, then it specifies that there are at most 5 tuples in the dataset with value of the *branch* attribute equals to ‘Chicago’ and the range of the *price* attribute of these tuples are between 0.0 and 149.99 (inclusive).

We could tweak this constraint to be “the most expensive product in ALL branches costs 149.99 and no more than 100 are sold”:

$$c_2 : \text{TRUE} \implies (0.00 \leq \text{price} \leq 149.99), (0, 100)$$

Suppose one wanted to define a simple histogram over a single attribute based the number of sales in each branch, we could express that with a tautology:

$$\begin{aligned} (\text{branch} = \text{'Chicago'}) &\implies (\text{branch} = \text{'Chicago'}), (100, 100) \\ (\text{branch} = \text{'New York'}) &\implies (\text{branch} = \text{'New York'}), (20, 20) \\ (\text{branch} = \text{'Trenton'}) &\implies (\text{branch} = \text{'Trenton'}), (10, 10) \end{aligned}$$

Observe how c_1 and c_2 interact with each other. Some of the missing data instances allowed by only c_2 are disallowed by c_1 (Chicago cannot have more than 5 sales at 149.99). This interaction will be the main source of difficulty in computing result ranges based on a set of PCs.

3.2 Predicate-Constraint Sets

Users specify their assumptions about missing data using a set of predicate constraints. A predicate-constraint set is defined as follows:

$$S = \{\pi_1, \dots, \pi_n\}$$

S gives us enough information to bound the results of common aggregate queries when there is *closure*: every possible missing row satisfies at least one of the predicates.

Definition 3.2 (Closure). Let S be a predicate constraint set with the elements $\pi_i = (\psi_i, \nu_i, \kappa_i)$. S is closed over an attribute domain $\mathcal{D}$ if for every $t \in \mathcal{D}$:

$$\exists \pi_i \in S : \psi_i(t)$$

Closure is akin to the traditional closed world assumption in logical systems, namely, the predicate constraints completely characterize the behavior of the missing rows over the domain.

To understand closure, let us add a new constraint that says “the most expensive product in “New York” is 100.00 and no more than 10 of them are sold:

$$c_3 : (\text{branch} = \text{'New York'}) \implies (0.00 \leq \text{price} \leq 100.00), (0, 10)$$

Closure over c_1, c_3 means that all of the missing rows come from either New York or Chicago.

4 CALCULATING RESULT RANGES

This section focuses on a simplified version of the bounding problem. We consider a single table and a single attribute aggregates. Let q denote such an aggregate query. The problem to calculate the upper bound is:

$$u = \max_R q(R) \tag{1}$$

$$\text{subject to: } R \models S$$

We will show that our bounds are tight—meaning that the bound found by the optimization problem is a valid relation that satisfies the constraints.

Throughout the rest of the paper, we only consider the maximal problem. Unless otherwise noted, our approach also solves the lower bound problem:

$$l = \min_{R \in \mathcal{R}} q(R)$$

$$\text{subject to: } R \models S$$

Specifically, we solve the lower bound problem in two settings. In a general setting where there is no additional constraint, the minimal problem can be solved by maximizing the negated problem: we first negate the value constraints, solve the maximizing problem with negated weights, and negate the final result.

In a special but also common setting, all the frequency constraints’ lower bounds are 0 (i.e., each relation has no

Figure 2: Predicates in a predicate-constraint set are possibly overlapping. The first step of the algorithm is to decompose a set of predicate-constraints into disjoint cells.

minimum number of missing rows) and the value constraints' lower bounds are 0 (i.e., all attributes are non-negative), the lower bound is easily attained by the absent of missing row.

4.1 Cell Decomposition

Intuitively, we can think about the optimization problem as an allocation. We strategically assign rows in each of the predicates to maximize the aggregate query. However, the first challenge is that a row may fall in multiple Predicate-Constraints' predicates, so this row may "count towards" multiple Predicate-Constraints. As the first step of the solution, we decompose the potential overlapping Predicate-Constraints' predicates into disjoint cells.

An example that illustrates the decomposition is depicted in Figure 2. Each predicate represents a sub-domain. For each subset of predicates, a cell is a domain that only belongs to these predicates and not others. Thus, for n predicates in a predicate-constraint set there are at most $O(2^n)$ cells. The cells take the form of conjunctions and negations of the predicates of each of the n predicate constraints:

$$\begin{aligned}
 c_0 &= \psi_1 \wedge \dots \wedge \psi_n \\
 c_1 &= \psi_1 \wedge \dots \wedge \neg\psi_n \\
 c_2 &= \psi_1 \wedge \dots \wedge \neg\psi_{n-1} \wedge \psi_n \\
 &\dots \\
 c_{2^n-1} &= \neg\psi_1 \wedge \dots \wedge \neg\psi_{n-1} \wedge \neg\psi_n
 \end{aligned}$$

For each c_i , we have to reconcile the active predicate constraints (not negated above). Each cell is thus assigned the most restrictive upper and lower value bounds, and upper and lower cardinality bounds in the set of active constraints. Not all possible cells will be satisfiable—where there exists a row $t \in \mathcal{D}$ that satisfies the new predicate-constraint. As in Figure 2, there are 7 possible subsets, but there are only 5 satisfiable cells. We use the Z3 [7] solver to prune all the cells that are not satisfiable.

Optimization 1. Predicate Pushdown: Cell decomposition is a costly process for two reasons. First, there is a potentially exponential number of cells. Second, determining whether a cell is satisfiable is not always easy (each check is on the order of 10's of ms). One obvious optimization is to push down the predicates of the target query into the decomposition process. When the target query has predicates, we exclude all cells that do not overlap with the query's predicate.

Optimization 2. DFS Pruning: The naive solution is to simply sequentially iterate through all of the cells and test each for satisfiability. Note that for a problem with n PCs, each logical expression describing a cell is a conjunction of n predicates. Conjunctions can be short-circuited if any of their constituent elements evaluate to false. Therefore, the process of evaluating the satisfiability of the cells can be improved by using a Depth First Search that evaluates prefixes of growing length rather than a sequential evaluation. With DFS, we can start from the root node of all expressions of length 1, add new PCs to the expression as we traverse deeper until we reach the leaf nodes which represent expressions of length n . As we traverse the tree, if a sub-expression is verified by Z3 to be unsatisfiable, then we can perform pruning because any expression contains the sub-expression is unsatisfiable.

Optimization 3. Expression Re-writing: To further improve the DFS process, we can re-write the logical expressions to have an even higher pruning rate. There is one simple re-writing heuristic we apply:

$$(X \wedge \neg(X \wedge Y)) = \text{True} \implies X \wedge \neg Y = \text{True}$$

It means if we have verified a sub-expression X to be satisfiable, and we also verified that after adding a new PC Y , the expression $X \wedge Y$ becomes unsatisfiable. We can conclude that $X \wedge \neg Y$ is satisfiable without calling Z3 to verify. As shown in the experiment section, the DFS pruning technique combined with the rewriting can prune over 99.9% cells in real-world problems.

Optimization 4. Approximate Early Stopping: With the DFS pruning technique, we always get the correct cell decomposition result because we only prune cells that are verified as unsatisfiable. We also propose an approximation that can trade range tightness for a decreased run time. The idea is to introduce 'False-Positives', i.e., after we have used DFS to handle the first K layers (sub-expressions of size K), we stop the verification and consider all cells that have not been pruned as satisfiable. These cells are then admitted to the next phase of processing where we use them to formulate a MILP problem that can be solved to get our bound. Admitting unsatisfiable cells introduces 'False-Positives' that would make our bound loose, but it will not violate the correctness of the result (i.e. the result is still a bound) because:

(1) the ‘true-problem’ with correctly verified cells is now a sub-problem of the approximation and (2) the false-positive cells does not add new constraints to the ‘true-problem’.

4.2 Integer-Linear Program

We assume that the cells in C are ordered in some way and indexed by i . Based on the cell decomposition, we denote C_i as the (sub-)set of Predicate-Constraints that cover the cell i . Then for each cell i , we can define its maximal feasible value $U_i(a) = \min_{p \in C_i} p.v.h_a$, i.e., the minimum of all C_i ’s value constraints’ upper bounds on attribute a .

A single general optimization program can be used to bound all the aggregates of interest. Suppose we are interested in the SUM over attribute a , we slightly abuse the terms and define a vector U where $U_i = U_i(a)$. Then, we can define another vector X which represents the decision variable. Each component is an integer that represents how many rows are allocated to the cell.

The optimization problem is as follows. To calculate the upper bound:

$$\begin{aligned} & \max_X U^T X \\ \text{subject to: } & \forall j, k_l^{(j)} \leq \sum_{i: j \in C_i} X[i] \leq k_u^{(j)} \\ & \forall i, X[i] \text{ is integer} \end{aligned} \quad (2)$$

As an example, the first constraint in Figure 2 is $k_l^1 \leq x_1 + x_2 \leq k_u^1$ for π_1 , so on and so forth. This MILP can be solved with a commercial optimization solver such as Gurobi or CPLEX.

Given the output of the above optimization problem, we can get bounds on the following aggregates:

COUNT: The count of cardinality can be calculated by setting U as the unit vector ($U_i = 1$ for all i).

AVG: We binary search the average result: to testify whether r is a feasible average, we disallow all rows to take values smaller than r and invoke the above solution for the maximum SUM and the corresponding COUNT. If the overall average is above r , then we test $r' > r$, otherwise we continue testing $r' < r$.

MAX/MIN: Assuming all cells are feasible, the max is the largest of all cells’ upper bound. Min can be handled in a similar way.

Faster Algorithm in Special Cases In the case that the Predicate-Constraints have disjoint predicates, the cell decomposition problem becomes trivial as each predicate is a cell. The MILP problem also degenerates since all the constraints, besides each variable being integer, do not exist at all. Thus, if we take the SUM problem as an example, the solution is simply the sum of each Predicate-Constraints’ maximum sum, which is the product of the maximum value

and the maximum cardinality. This disjoint case is related to the work in data privacy by Zhang et al. [26].

4.3 Complexity

Next, we analyze the complexity of the predicate-constraint optimization problem. Suppose each predicate-constraint is expressed in Disjunctive Normal Form with at most p clauses. Suppose, we are given a set of N such predicate-constraints. The data complexity of the predicate-constraint optimization problem is the number of computational steps required to solve the optimization problem for a fixed p and a variable N . For $p \geq 2$, we show that the problem is computationally hard in N based on a reduction:

PROPOSITION 4.1. *Determining the maximal sum of a relation constrained by a predicate-constraint is NP-Hard.*

SKETCH. We prove this proposition by reduction to the maximal independent set problem, which is NP-Hard. An independent set is a set of vertices in a graph, no two of which are adjacent. We can show that every independent set problem can be described as a predicate-constraint maximal sum problem. Let $G = (V, E)$ be a graph. We can think of this graph as a relational table of “vertex” rows. For example: $(V1, V2)$. For each vertex $v \in V$, we define a single Predicate-Constraint with a max value of 1 and a max frequency of 1 ($x = v, [0, 1], [0, 1]$). For each edge between v and v' , we define another predicate constraint that is equality to either of the vertices also with a max frequency of 1:

$$(x = v \vee x = v', [0, 1], [0, 1]).$$

This predicate constraint exactly contains the two vertex constraints. Therefore, the cells after the cell-decomposition step perfectly align with those vertices. The optimization problem allocates rows to each of those cells but since the edge constraint has a max frequency is 1 only one of the vertex cells can get an allocation. Since all of the vertex constraints have the same max value, the optimization problem finds the most number of such allocations that are consistent along the edges (no two neighboring vertices both are allocated). This is exactly the definition of a maximal independent set. Since the maximal sum problem is more expressive than the maximal independent set its complexity is greater than NP-Hard in the number of cells. $\square$

The hardness of this problem comes from the number of predicate constraints (i.e., data complexity). While, we could analyze the problem for a fixed N and variable p (i.e., query complexity), we find that this result would not be informative and is highly problem specific with very large constant-factors dependent on N .

4.4 Numerical Example

To understand how this optimization problem works, let's consider a few simple numerical examples using our example dataset. Suppose, we have a query that wants to calculate the total number of sales over a period:

```
SELECT SUM(price)
FROM Order
WHERE utc >= Nov-11 0:00 AND
      utc <= Nov-13 0:00
```

For the sake of simplicity, let's assume that all of this data was missing and the only information we have is what is described in the PCs. Suppose, we define two PCs that describe the price of the least/most expensive items sold on the day, and that between 50 and 100 items were sold:

$$t_1: Nov-11 \leq utc < Nov-12 \implies 0.99 \leq price \leq 129.99, (50, 100)$$

$$t_2: Nov-12 \leq utc < Nov-13 \implies 0.99 \leq price \leq 149.99, (50, 100)$$

Since the PCs are disjoint, we can trivially compute the result range for the total sales:

$$[50 \times 0.99 + 50 \times 0.99, 100 \times 129.99 + 100 \times 149.99] \\ = [99.00, 27998.00] \blacksquare$$

Now, suppose, we had overlapping PCs. This makes the solution much harder to manually reason about since we have to account for the interaction between the constraints.

$$t_1: Nov-11 \leq utc < Nov-12 \implies 0.99 \leq price \leq 129.99, (50, 100)$$

$$t_2: Nov-11 \leq utc < Nov-13 \implies 0.99 \leq price \leq 149.99, (75, 125)$$

This requires a cell decomposition to solve. There are 3 possible cells:

$$c_1 := (Nov-11 \leq utc < Nov-12) \wedge (Nov-11 \leq utc < Nov-13)$$

$$c_2 := \neg(Nov-11 \leq utc < Nov-12) \wedge (Nov-11 \leq utc < Nov-13)$$

$$c_3 := (Nov-11 \leq utc < Nov-12) \wedge \neg(Nov-11 \leq utc < Nov-13)$$

c_3 is clearly not satisfiable, so we can discount this cell. Then, we need to figure out how much to allocate to c_1 and c_2 . The lower bound can be achieved by an allocation of 50 tuples to c_1 and 25 to c_2 . The upper bound can be achieved by an allocation of 50 tuples to c_1 and 75 tuples to c_2 . Note that the optimal allocation does not maximize tuples in c_1 .

$$[50 \times 0.99 + 25 \times 0.99, 50 \times 129.99 + 75 \times 149.99] = \\ [74.25, 17748.75] \blacksquare$$

5 RESULTS

Now, we evaluate the PC framework in terms of accuracy and efficiency at modeling missing data rows.

5.1 Experimental Setup

We set up each experiment as follows: (1) summarize a dataset with n Predicate-Constraints, (2) each competitor framework gets a similar amount of information about the dataset (e.g., a statistical model with $O(n)$ statistical parameters), (3) we estimate a query result using each framework with calibrated error bounds. Our comparative evaluation focuses on SUM, COUNT queries as those are what the baselines support. In the full technical report, we show similar results on MIN, MAX, AVG queries using our framework and additional experiments on more real-world datasets.

It is important to note that each of these frameworks will return a confidence interval and not a single result. Thus, we have to measure two quantities: the **failure rate** (how often the true result is outside the interval) and the tightness of the estimated ranges. We measure tightness by the ratio between the upper bound and the result (we denote this as the **over estimation rate**). A ratio closer to 1 is better on this metric, but is only meaningful if the failure rate is low.

5.1.1 Sampling. To construct a sampling baseline, we assume that the user provides actual unbiased example missing data records. While this might be arguably much more difficult for a user than simply describing the attribute ranges as in a predicate constraint, we still use this baseline as a measure of accuracy. In our estimates, we use *only* these examples to extrapolate a range of values that the missing rows could take.

Uniform Sampling We randomly draw n samples (US-1) and $10n$ samples (US-10) from the set of missing rows.

Stratified Sampling We also use a stratified sampling method which performs a weighted sampling from partitions defined by the PCs that we use for a given problem. Similarly, we denote n samples as (ST-1) and $10n$ samples (ST-10).

The goal is to estimate the result range using a statistical confidence interval. Commonly, approximate query processing uses the Central Limit Theorem to derive bounds. These confidence intervals are parametric as they assume that the error in estimation is Normally distributed. Alternatively, one could use a non-parametric method, which does not make the Normal assumption, to estimate confidence intervals like those described in [9] (we use this formula unless otherwise noted). We denote confidence interval schemes in the following way: US-1p (1x sample using a parametric confidence interval), US-10n (10x sample using a non-parametric confidence interval), ST-10p (10x stratified sample using a parametric confidence interval), etc.

5.1.2 Equiwidth Histogram. We build a histogram on all of the missing data on the aggregate attribute with N buckets and use it to answer queries. We use standard independent assumptions to handle queries across multiple attributes.

5.1.3 PCs. The accuracy of the PC framework is dependent on the particular PCs that are used for a given task. In our “macro-benchmarks”, we try to rule out the effects of overly tuned PCs. We consider two general schemes: **Corr-PC**, even partitions of attributes correlated with the aggregate of interest, and **Rand-PC**, randomly generated PCs.

For **Corr-PC**, we identify the (other than the aggregation attribute) most correlated attributes with the aggregate to partition on. We divide the combined space into equi-cardinality buckets where each partition contains roughly the same number of tuples. We omit the details of this scheme for the sake of brevity. For **Rand-PC**, we generate random overlapping predicate constraints over the same attributes (not necessarily equi-cardinality). We take extra care to ensure they adequately cover the space to be able to answer the desired queries. We see these as two extremes: the reasonably best performance one could expect out of the PC framework and the worst performance. We envision that natural use cases where a user is manually defining PCs will be somewhere in the middle in terms of performance.

5.2 Intel Wireless Dataset

The Intel wireless dataset [21] contains data collected from 54 sensors deployed in the Intel Berkeley Research lab in 2004. This dataset contains 3 million rows, 8 columns including different measurements like humidity, temperature, light, voltage as well as date and time of each record. We consider aggregation queries over the light attribute. For this dataset, Corr-PC is defined as $n = 2000$ predicate constraints over the attributes device_id and time. Rand-PC defines random PCs over those same attributes. Missing rows are generated from the dataset in a correlated way—removing those rows maximum values of the light attribute.

We compare Corr-PC and Rand-PC with 3 baselines we mentioned in Section 5.1: US-1n, ST-1n, and Histogram. We vary the fraction of missing rows r and evaluate the accuracy and failure rate of each technique. Figure 3 and Figure 4 illustrate the experimental results: (1) as per our formal guarantees, both Corr-PC, Rand-PC (and Histograms) do not fail if they have accurate constraints, (2) despite the hard guarantee, the confidence intervals are not too “loose” and are competitive or better than those produced by a 99.99% interval, (3) informed PCs are an order of magnitude more accurate than randomly generated ones. There are a couple of trends that are worth noting. First, the failure rate for sampling techniques on the SUM queries is higher than the expected 1 in 10000 failures stipulated by the confidence intervals. In this missing data setting, a small number of example rows fail to accurately capture the “spread” of a distribution (the extremal values), which govern failures in

estimation. Queries that require values like SUM and AVG are very sensitive to these extrema.

It is important to note that these experiments are idealized in the sense that all the baselines get true information about the missing data and have to summarize this information into $O(n)$ space and measure how useful that stored information is for computing the minimal and maximal value a workload of aggregate queries could take. There is a subtle experimental point to note. If a query is fully covered by the missing data, we solve the query with each baseline and record the results; if a query is partially covered by the missing data, we solve the part that is missing with each baseline then combine the result with a ‘partial ground truth’ derived from the existing data; finally, if a query is not overlapping with the missing data then we can get an accurate answer for such a query. Such issues are common to all the baselines, in the upcoming experiments we only consider the accuracy of representing the missing data (and not partially missing data) to simplify questions of correlation and accuracy.

Figure 3: Performance of baselines given different missing ratio, evaluated with 1000 COUNT(*) query on the Intel Wireless dataset.

Figure 4: Performance of baselines given different missing ratio, evaluated with 1000 SUM query on the Intel Wireless dataset.

5.3 Detailed Sampling Comparison

One might argue that 99.99% is an overly conservative confidence interval. In this experiment, we evaluate the performance of the uniform sampling baseline in terms of failure rate and accuracy as function of that confidence interval setting. Our results show there is not a clear way to calibrate the confidence intervals to either minimize failures or avoid inaccuracy. Results in Table 1 show a clear trade-off between failure rate and accuracy: when we increase the confidence interval, the over-estimation ratio increases, and the failure rate decreases. However, even with a 99.99% confidence interval derived from Chernoff Bound, the failure rate is non-zero. We believe that the PC framework provides the user with a competitive accuracy but the guarantee of no failures.

		Failure Rate %						
Conf (%)		80	85	90	95	99	99.9	99.99
US-1n		20.1	15.6	11.4	6.9	3.4	2.4	0.8
Corr-PC		–0–						
		Over Estimation Rate						
US-1n		1.07	1.08	1.11	1.13	1.2	1.27	3.13
Corr-PC		–2.23–						

Table 1: Trade-off between failure rate and accuracy of an uniform sampling baseline given different confidence interval vs. Corr-PC.

5.3.1 Sampling More Data. In all of our experiments above, we consider a 1x random sample. Where the baseline is given the same amount of data compared to the number of PCs. PCs are clearly a more accurate estimate in the “small data” regime, what if the sample size was larger. In Figure 5, we use the Intel Wireless dataset to demonstrate the performance of the non-parametric bounds using different sample sizes. And the results demonstrate a clear trend of convergence as we increase the sample size. If we consider data parity (1x), the confidence interval is significantly less accurate than a well-designed PC. One requires 10x the amount of data to cross over in terms of accuracy. We envision that PCs will be designed by a data analyst by hand, and thus the key challenge is to evaluate the accuracy of the estimation with limited information about the distribution of missing values.

5.4 Scalability

In this subsection, we evaluate the scalability of PCs and optimization techniques. As mentioned in earlier sections, the complete process of solving a PC problem including two parts. First, we need to perform cell decomposition to find out which cells are valid and second, we need to formalize a MILP problem with the valid cells that can be solved to find

Figure 5: Performance of the uniform sampling baseline with different sample size.

Figure 6: Our optimizations reduce the number of cells evaluated during the cell decomposition phase by over a 1000x.

the optimal bound. The vanilla version described above is intractable because the number of sub-problems we need to solve in the cell decomposition phase is exponential to the number of PCs.

We presented a number of optimizations in the paper to improve this time. We will show that naive processing of the PCs leads to impractical running time. We generate 20 random PCs that are very significantly overlapping. Figure 6 plots the number of cells evaluated as well as the run time of the process. The naive algorithm evaluates the SAT solver on more than 1000x more cells than our optimized approach.

Since cell decomposition is really the most expensive step we can prune and save about 99.9% of the solving time by using DFS (early termination for cell decomposition) and the rewriting heuristic. Without these optimizations, PCs are simply impractical at large scales. In the technical report, we show that non-overlapping PCs scale effectively.

6 RELATED WORK

The overarching challenge addressed in the PC framework is related to the concept of “reverse data management” proposed by Meliou et al. [17, 18]. Meliou et al. argue that as data grow in complexity, analysts will increasingly want to know not what their data currently says but what changes have to happen to the dataset to force a certain outcome. Such *how-to*

analyses are useful in debugging, understanding sensitivity, as well as planning for future data. Meliou et al. build on a long line of *what-if* analysis and data provenance research, which study simulating hypothetical updates to a database and understanding how query results might change [3, 8]. While we address similar concerns to this line of work in spirit, our focus on aggregate queries and confidence intervals leads to a very different set of technical contributions. The PC framework should be evaluated much more like a synopsis data structure than a data provenance reasoning systems.

Therefore, our experiments largely focus on evaluations against other synopsis structures and how to extract confidence intervals from them [6]. While Approximate Query Processing (AQP) has been studied for decades [20], it is well-known that the confidence intervals produced can be hard to interpret [13]. This is because estimating the spread of high dimensional data from a small sample is fundamentally hard, and the most commonly used Central-Limit Theorem-based confidence intervals rely on estimated sample variance. Especially for selective queries, these estimates can be highly fallible—a 95% confidence interval may “fail” significantly more than 5% of the time [1]. Unfortunately, as confidence intervals become more conservative, e.g., using more general statistical bounding techniques, their utility drops [9]. In a sense, our optimization algorithm automatically navigates this trade-off. The algorithm optimizes the tightest bound given available information in the form of PCs. We interpret PCs as generalized histograms with overlapping buckets and uncertain bucket counts. Despite these differences with AQP, we do believe that the connections between uncertainty estimation and dirty data (like missing rows) are under-studied [15, 16]. We also believe that in future work mixed systems with both PCs and samples can have the best of both worlds, e.g., augmenting Quickr with PC estimation [12].

Deterministic alternatives to AQP have been studied in some prior work. Potti et al. propose a paradigm called DAQ [23] that does reason about hard ranges instead of confidence intervals. DAQ models uncertainty at relation-level instead of predicate-level like in PCs and DAQ does not handle cardinality variation. In the limited scenario of windowed queries over time-series data, deterministic bounds have been studied [2]. The technical challenge arises with overlapping constraints and complex query structures (like join conditions and arbitrary predicates). Similarly, we believe that classical variance reduction techniques for histograms could be useful for PC generation in future work [22], since histograms are a dense 1-D special case of our work.

c-tables are one of the classical approaches for representing missing data in a relation [10]. Due to the frequency constraints in Predicate-Constraint sets, we can represent

cases that go beyond the typical closed-world assumption (CWA) is required in c-tables, where all records are known in advance and null cells are specifically annotated. There is also recent work that studies missing rows from databases. m-tables study variable cardinality representations to go beyond the CWA. In m-tables, cardinality constraints are specified per-relation. We specify frequency constraints per predicate. However, Sundarmurthy et al. [24] do not consider the problem of aggregate query processing on uncertain relations. There is similarly related work that studies intentionally withholding partitioned data for improved approximate query performance [25]. We believe that the novelty of our framework is the efficient estimation of aggregate query confidence intervals. Similarly, the work by Burdick et al. is highly related where they study databases with certain “imprecise” regions instead of realized tuples [4]. It is important to note, that our objective is not to build the most expressive language to represent uncertain data but rather one that we can pragmatically use to bound aggregate queries.

The privacy literature has studied a version of this problem: bounding aggregate queries on uncertain data [11, 26]. In fact, Zhang et al. can be seen as solving the partitioned version of our problem [26]. However, they do not need to consider the overlapping case in the way that our work does.

7 CONCLUSION

We proposed a framework that can produce automatic contingency analysis, i.e., the range of values an aggregate SQL query could take, under formal constraints describing the variation and frequency of missing data tuples. There are several interesting avenues for future work. First, we are interested in studying these constraints in detail to model dirty or corrupted data. Rather than considering completely missing or dirty rows, we want to consider rows with some good and some faulty information. From a statistical inference perspective, this new problem statement likely constitutes a middle ground between sampling and Predicate-Constraints. Second, we would like to further understand the robustness properties of result ranges computed by Predicate-Constraints as well as other techniques. Understanding when result ranges are meaningful for real-world analytics will be an interesting study. Finally, we would like to extend the Predicate-Constraint framework to be more expressive and handle a broader set of queries.

ACKNOWLEDGMENTS

This work was supported by an NSF Award CCF-1139158, a Google DAPA Research Award, the CERES Center for Unstoppable Computing, a US Army TATRC Research Award, the University of Chicago’s Research Computing Center, and Intel.

REFERENCES

- [1] S. Agarwal, H. Milner, A. Kleiner, A. Talwalkar, M. Jordan, S. Madden, B. Mozafari, and I. Stoica. Knowing when you're wrong: building fast and reliable approximate query processing systems. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, pages 481–492. ACM, 2014.
- [2] J. Brito, K. Demirkaya, B. Etienne, Y. Katsis, C. Lin, and Y. Papakonstantinou. Efficient approximate query answering over sensor data with deterministic error guarantees. *arXiv preprint arXiv:1707.01414*, 2017.
- [3] P. Buneman, S. Khanna, and T. Wang-Chiew. Why and where: A characterization of data provenance. In *International Conference on Database Theory*, pages 316–330. Springer, 2001.
- [4] D. Burdick, P. M. Deshpande, T. Jayram, R. Ramakrishnan, and S. Vaithyanathan. Efficient allocation algorithms for olap over imprecise data. In *Proceedings of the 32nd international conference on Very large data bases*, pages 391–402. VLDB Endowment, 2006.
- [5] X. Chu, I. F. Ilyas, S. Krishnan, and J. Wang. Data cleaning: Overview and emerging challenges. In *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data*, pages 2201–2206. ACM, 2016.
- [6] G. Cormode, M. Garofalakis, P. J. Haas, C. Jermaine, et al. Synopses for massive data: Samples, histograms, wavelets, sketches. *Foundations and Trends® in Databases*, 4(1–3):1–294, 2011.
- [7] L. de Moura and N. Björner. Z3: An efficient smt solver. In C. R. Ramakrishnan and J. Rehof, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
- [8] D. Deutch, Z. G. Ives, T. Milo, and V. Tannen. Caravan: Provisioning for what-if analysis. In *CIDR*, 2013.
- [9] J. M. Hellerstein, P. J. Haas, and H. J. Wang. Online aggregation. In *Proceedings of the 1997 ACM SIGMOD International Conference on Management of Data*, pages 171–182, 1997.
- [10] T. Imieliński and W. Lipski Jr. Incomplete information in relational databases. In *Readings in Artificial Intelligence and Databases*, pages 342–360. Elsevier, 1989.
- [11] N. Johnson, J. P. Near, and D. Song. Towards practical differential privacy for sql queries. *Proceedings of the VLDB Endowment*, 11(5):526–539, 2018.
- [12] S. Kandula, A. Shanbhag, A. Vitorovic, M. Olma, R. Grandl, S. Chaudhuri, and B. Ding. Quickr: Lazily approximating complex adhoc queries in bigdata clusters. In *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data*, pages 631–646. ACM, 2016.
- [13] O. Kennedy and B. Glavic. Analyzing uncertain tabular data. In *Information Quality in Information Fusion and Decision Making*, pages 243–277. Springer, 2019.
- [14] T. Kraska. Northstar: An interactive data science system. *Proceedings of the VLDB Endowment*, 11(12):2150–2164, 2018.
- [15] S. Krishnan, J. Wang, M. J. Franklin, K. Goldberg, and T. Kraska. Stale view cleaning: Getting fresh answers from stale materialized views. *Proceedings of the VLDB Endowment*, 8(12):1370–1381, 2015.
- [16] S. Krishnan, J. Wang, M. J. Franklin, K. Goldberg, T. Kraska, T. Milo, and E. Wu. Sampleclean: Fast and reliable analytics on dirty data. *IEEE Data Eng. Bull.*, 38(3):59–75, 2015.
- [17] A. Meliou, W. Gatterbauer, and D. Suciu. Reverse data management. *Proceedings of the VLDB Endowment*, 4(12), 2011.
- [18] A. Meliou and D. Suciu. Tiresias: the database oracle for how-to queries. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 337–348. ACM, 2012.
- [19] H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case optimal join algorithms. *Journal of the ACM (JACM)*, 65(3):16, 2018.
- [20] C. Olston, B. T. Loo, and J. Widom. Adaptive precision setting for cached approximate values. In *Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data*, pages 355–366. New York, NY, USA, 2001. ACM.
- [21] W. H. Peter Bodik et al. Intel wireless dataset. <http://db.csail.mit.edu/labdata/labdata.html>, 2004.
- [22] V. Poosala, P. J. Haas, Y. E. Ioannidis, and E. J. Shekita. Improved histograms for selectivity estimation of range predicates. *ACM SIGMOD Record*, 25(2):294–305, 1996.
- [23] N. Potti and J. M. Patel. Daq: a new paradigm for approximate query processing. *Proceedings of the VLDB Endowment*, 8(9):898–909, 2015.
- [24] B. Sundarmurthy, P. Koutris, W. Lang, J. Naughton, and V. Tannen. m-tables: Representing missing data. In *20th International Conference on Database Theory (ICDT 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [25] B. Sundarmurthy, P. Koutris, and J. Naughton. Exploiting data partitioning to provide approximate results. In *Proceedings of the 5th ACM SIGMOD Workshop on Algorithms and Systems for MapReduce and Beyond, BeyondMR'18*, pages 5:1–5:5. New York, NY, USA, 2018. ACM.
- [26] Q. Zhang, N. Koudas, D. Srivastava, and T. Yu. Aggregate query answering on anonymized tables. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 116–125. IEEE, 2007.