

# RAPIDASH: Efficient Detection of Constraint Violations

Zifan Liu  
University of Wisconsin-Madison  
zifan@cs.wisc.edu

Shaleen Deep  
Microsoft  
shaleen.deep@microsoft.com

Anna Fariha  
University of Utah  
afariha@cs.utah.edu

Fotis Psallidas  
Microsoft  
Fotis.Psallidas@microsoft.com

Ashish Tiwari  
Microsoft  
ashish.tiwari@microsoft.com

Avrilia Floratou  
Microsoft  
Avrilia.Floratou@microsoft.com

## ABSTRACT

Denial Constraint (DC) is a well-established formalism that captures a wide range of integrity constraints commonly encountered, including candidate keys, functional dependencies, and ordering constraints, among others. Given their significance, there has been considerable research interest in achieving fast detection of DC violations, especially to support activities related to data exploration and preparation. Despite the significant advancements in the field, prior work exhibits notable limitations when confronted with large-scale datasets: the current state-of-the-art algorithm demonstrates a quadratic (worst-case) time and space complexity relative to the dataset’s number of rows. In this paper, we establish a connection between orthogonal range search and DC violation detection. We then introduce RAPIDASH, a novel algorithm that demonstrates near-linear time and space complexity, representing a theoretical improvement over prior work. To validate the effectiveness of our algorithm, we conduct comprehensive evaluations on both open-source and real-world production datasets, with our production datasets notably being an order of magnitude larger than the datasets employed in prior studies. Our results reveal that RAPIDASH achieves up to 84× faster performance compared to state-of-the-art approaches while also exhibiting superior scalability.

### PVLDB Reference Format:

Zifan Liu, Shaleen Deep, Anna Fariha, Fotis Psallidas, Ashish Tiwari, and Avrilia Floratou. RAPIDASH: Efficient Detection of Constraint Violations. PVLDB, 17(8): 2009-2021, 2024.  
doi:10.14778/3659437.3659454

## 1 INTRODUCTION

Integrity constraints play a pivotal role in a wide range of data analysis tasks such as data exploration [1, 16], data cleaning and repair [19, 37], data synthesis [17], and query optimization [28]. By enforcing integrity constraints, organizations can ensure reliability, consistency, and accuracy of their data, enabling them to make informed decisions, derive meaningful insights, and extract maximum value from their datasets. One class of constraints that is of particular interest is *Denial Constraints* (DCs) [13]. DCs are appealing as they are expressive enough to capture many useful integrity

Table 1: Tax rates for people in different states in the USA.

	SSN	Zip	Salary	FedTaxRate	State	StateCode
$t_1$	100	10108	3000	20%	New York	01
$t_2$	101	53703	5000	15%	Wisconsin	02
$t_3$	102	53703	6000	20%	Wisconsin	02
$t_4$	103	53703	4000	22%	Wisconsin	02

constraints such as functional dependencies, ordering constraints, unique column combinations [2, 39], etc.

**Example 1.** Table 1 shows a sample of a tax dataset that contains information about tax rates for people in different US states. The following two rules are true about this dataset:  $C_1$  : SSN column is a candidate key;  $C_2$  : Zip  $\rightarrow$  State is a functional dependency. However, the rule  $C_3$  : “for all people in the same state, if person A has an equal or larger salary than person B, then A should have an equal or larger tax rate than B”—is not true since the tuples  $t_2$  and  $t_3$  have a higher value for Salary than the tuple  $t_4$  but have tax rates lower than 22% (15% and 20%, respectively). Each of these rules can be expressed as DCs as we will see in Section 2.

We focus on the detection of DC violations on a given dataset. The process involves verifying whether a given DC is satisfied on a specific dataset (we refer to this task as DC verification in Section 2), and is particularly valuable during data exploration, where analysts aim to quickly ascertain the presence or absence of specific patterns within the dataset. For example, a data analyst, while exploring a dataset for the first time, might want to quickly detect which of the constraints (rules) presented in Example 1 hold on the dataset. Additionally, in case a constraint does not hold on the dataset, the analyst might want to enumerate a few tuples that *violate* the constraint (see violation enumeration in Section 2). In Example 1, there are two violations with respect to  $C_3$ : ( $t_2, t_4$ ) and ( $t_3, t_4$ ). These violations can be fed into upstream data cleaning tools [15, 18, 37], which can be used to repair data inconsistencies. Thus, DC violation detection serves as a valuable tool in assessing dataset quality [16].

**Limitations of existing work.** In recent years, substantial advancements have happened in the field of DC violation detection [34, 35]. However, our practical experience in applying some of these approaches to real-world production datasets has unveiled noteworthy limitations of existing methods (refer to Section 5 for comprehensive details). In particular, the best-known algorithm, FACET [34], has a worst-case time and space complexity  $\Omega(|\mathbf{R}|^2)$  on a given relation  $\mathbf{R}$  (with cardinality  $|\mathbf{R}|$  (number of rows), which makes it infeasible for large-scale datasets.

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](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.  
Proceedings of the VLDB Endowment, Vol. 17, No. 8 ISSN 2150-8097.  
doi:10.14778/3659437.3659454

**Our approach.** In this work, we make the connection between the problem of DC violation detection and *orthogonal range search* [8, 9]. Given a set of multidimensional points, orthogonal range search allows efficient searching of all points that lie within an axis-aligned, multidimensional rectangle. The key insight of our work is that range search can be utilized by modeling the tuples in the relation as multidimensional points. We illustrate this with an example.

**Example 2.** Consider the rule  $C_3$  from Example 1 and the three tuples with State=Wisconsin ( $t_2$ ,  $t_3$ , and  $t_4$ ) from Table 1. The three tuples can be visualized as two-dimensional points when considering the columns Salary and FedTaxRate (Figure 1a), which are of interest for rule  $C_3$ . Let us fix our attention on tuple  $t_4$  that has Salary = 4000 and FedTaxRate = 22. According to  $C_3$ , all tuples that have Salary  $\geq 4000$ , must also have a FedTaxRate  $\geq 22$ . In other words, if we are able to find a tuple that has Salary  $\geq 4000$  but a strictly smaller tax rate than 22, we will have found a pair that violates the rule  $C_3$ . The shaded rectangular area in Figure 1b encodes precisely this criteria: the shaded rectangular area represents the region with Salary  $\geq 4000$  but FedTaxRate  $< 22$ . In our example, both  $t_2$  and  $t_3$  lie in the shaded region and thus violate  $C_3$  with respect to  $t_4$ .


To leverage this observation, we would need to employ a data structure that allows fast searching of points for any given multidimensional axis-parallel rectangle. This is precisely the orthogonal range search problem [8, 9] for which there exist practical data structures. Making the connection between DC violation detection and orthogonal range search opens up new possibilities in addressing this challenge, potentially surpassing the capabilities of state-of-the-art methods in violation detection.

In this work, we take the first step in exploring the relationship between these two domains. We demonstrate that there are various challenges in realizing our vision. First, taking an arbitrary DC (formally defined in Section 2) and translating it into an instance of the orthogonal range search problem is not trivial. The complexity arises from having to support constraints with complex structures. Our work is the first one to show how different classes of denial constraints can be mapped into instances of the orthogonal range search problem. Furthermore, we show that directly applying vanilla range search data structures (i.e. without any of the optimizations proposed in this work) leads to poor performance. Figure 2 shows the total running time of verifying four practical DCs from [36] on the TPC-H dataset that has been used by several prior works related to DCs<sup>1</sup>. We find that using vanilla range search is 30 $\times$  worse than FACET (the state-of-the-art) and 100 $\times$  worse than RAPIDASH (our system). In light of these issues, a principled study to design algorithms and optimizations that translate theoretical models for orthogonal range search into tangible performance improvements is imperative. In this work, we propose novel optimizations that provably reduce the time complexity over vanilla range search.


**Our contributions.** Our key contributions are the following:

- (1) Making the connection between orthogonal range search and violation detection and taking the first step in exploring the interplay between these two domains.

<sup>1</sup>See Section 5 for a more detailed analysis of the constraints.



**Figure 1: Geometric representation of the data for Example 2. (a) Salary and FedTaxRate for tuples with State = Wisconsin in table Tax. (b) The shaded area contains the tuples violating  $C_3$  with Salary  $\geq t_4$ .Salary and FedTaxRate  $< t_4$ .FedTaxRate.**



**Figure 2: Total time for verification of four DCs on TPC-H dataset (see Table 3) using different approaches.**

- (2) *A novel DC violation detection algorithm.* We perform a fine-grained classification of DCs and show how they can be mapped to the orthogonal range search problem. We present a near-optimal algorithm RAPIDASH for verifying a given DC on a dataset  $\mathbf{R}$  by leveraging foundational research in the domain of orthogonal range search [8, 31, 41]. Our proposed algorithm has a time complexity of  $O(|\mathbf{R}| \log^{f(\varphi)} |\mathbf{R}|)$ , where  $f(\varphi)$  is a parameter that is dependent only on the characteristics of the DC  $\varphi$  and not on the input dataset  $\mathbf{R}$ . This represents a significant improvement over the best-known verification algorithm [34], which has a worst-case quadratic complexity (both time and space). We also show that in certain scenarios, RAPIDASH can run in linear space while still achieving provably sub-quadratic running time.
- (3) *Experimental evaluation.* We conduct an extensive empirical evaluation over both open-source and production datasets, notably the latter being an order of magnitude larger than the datasets employed in prior studies. We show that RAPIDASH achieves up to 84 $\times$  speedup over the state-of-the-art [34] for DC violation detection. Additionally, our findings show that RAPIDASH exhibits better scalability over previous work.

## 2 BACKGROUND AND PROBLEM STATEMENT

**Relations.** Let  $\mathbf{R}$  be the input relation and  $\text{vars}(\mathbf{R})$  denote the finite set of attributes (i.e. the columns). We use  $|\mathbf{R}|$  to denote the cardinality (number of tuples) of  $\mathbf{R}$ . We will use  $A, B$  to denote attributes,  $s$  and  $t$  to denote tuples, and  $t.A$  to denote the value of attribute  $A$  of tuple  $t$ . Throughout the paper, we assume bag semantics where the relation can have the same tuple present multiple times.

**Denial Constraints (DCs).** DCs express predicate conjunctions to determine conflicting combinations of column values. They generalize other integrity constraints, including unique column combinations, functional dependencies, and order dependencies. We define a predicate  $p$  as the expression  $s.A \text{ op } t.B$  where  $s, t \in \mathbf{R}$ ,  $\text{op} \in \{=, \neq, \geq, >, \leq, <\}$  and  $A, B \in \text{vars}(\mathbf{R})$ . We will refer to  $\neq$  as disequality and  $\geq, >, \leq, <$  as inequalities. All operators except equality will be collectively referred to as non-equality operators. A DC  $\varphi$  is a conjunction of predicates of the following form:

$$\forall s, t \in R, s \neq t : \neg(p_1 \wedge \dots \wedge p_m)$$

A tuple pair  $(s, t)$  is said to be a violation if all predicates in  $\varphi$  evaluate to true. We will say that  $\varphi$  holds on  $\mathbf{R}$  if there is no violation, i.e., the DC is *exact*. An exact DC is said to be minimal if no proper subset of its predicates forms another exact DC. A predicate is said to be *homogeneous* if it is of the form  $s.A \text{ op } t.A$  or  $s.A \text{ op } s.B$ , i.e. it is either defined over a single column  $A$  or it is defined over a single tuple  $s$ , but two different columns; and *heterogeneous* if it is of the form  $s.A \text{ op } t.B$ . We will refer to  $s.A \text{ op } t.A$  as *row-level* homogeneous predicate since it is comparing across two rows.

Since most DCs of interest contain only row-level homogeneous predicates (such as ordering dependencies [39], functional dependencies, candidate keys, etc.), for simplicity, we will use the term homogeneous DC to refer a DC that contains only row-level homogeneous predicates. A heterogeneous DC can contain all types of predicates. Without loss of generality, we will assume that each column of  $\mathbf{R}$  participates in at most one predicate of a homogeneous DC. We will use  $\text{vars}_{\text{op}}(\varphi)$  to denote the set of columns in a homogeneous DC that appear in some predicate with the operator  $\text{op}$ . A DC may also contain *column-level* homogeneous predicates of the form  $s.A \text{ op } s.B$ . However, in the interest of space, we do not discuss them in this paper given that they are not very common in practice<sup>2</sup> and we defer the interested reader to [29].

**Example 3.** Continuing from Example 1, each constraint can be expressed using a DC as follows: (1)  $\varphi_1 : \neg(s.\text{SSN} = t.\text{SSN})$ ; (2)  $\varphi_2 : \neg(s.\text{Zip} = t.\text{Zip} \wedge s.\text{State} \neq t.\text{State})$ ; (3)  $\varphi_3 : \neg(s.\text{State} = t.\text{State} \wedge s.\text{Salary} \leq t.\text{Salary} \wedge s.\text{FedTaxRate} > t.\text{FedTaxRate})$ . The universal quantification is left implicit. Let us fix our attention to  $\varphi_3$ . Note that  $\text{vars}_{=}(\varphi_3) = \{\text{State}\}$ ,  $\text{vars}_{\leq}(\varphi_3) = \{\text{Salary}\}$ , and  $\text{vars}_{>}(\varphi_3) = \{\text{FedTaxRate}\}$ . All the DCs except  $\varphi_3$  hold on the relation Tax defined in Table 1 and are minimal exact DCs.

We are now ready to state the problem considered in the paper.

**Problem Statement.** Given a relation  $\mathbf{R}$  and a DC  $\varphi$ , determine:

- (1) whether  $\varphi$  holds on  $\mathbf{R}$  (DC verification).
- (2) enumerate all the tuple pairs that violate the constraint (violation enumeration).

All complexity results in this paper are based on the standard RAM model [22] of computation.

### 3 LIMITATIONS OF EXISTING SOLUTIONS

We now discuss the limitations of FACET for the problems we aim to address. In Section 5, we experimentally demonstrate some of these limitations using real-world datasets.

<sup>2</sup>None of the experimental evaluations in [33–36] contain a DC with column-level homogeneous predicates.

We begin by giving a brief description of the key ideas underlying FACET [34], the state-of-the-art system for DC violation detection. Let  $\text{tids}$  denote a set of tuple identifiers. All tuples in relation  $\mathbf{R}$  can be represented as  $\text{tids}_{\mathbf{R}} = \{t_1, \dots, t_{|\mathbf{R}|}\}$ . An ordered pair  $(\text{tids}_1, \text{tids}_2)$  represents all tuple pairs  $(s, t)$  such that  $s \in \text{tids}_1, t \in \text{tids}_2, s \neq t$ . FACET processes one predicate of the DC at a time, taking a set of ordered pairs  $(\text{tids}_1, \text{tids}_2)$  as input and generating another set of ordered pairs  $(\text{tids}'_1, \text{tids}'_2)$  that represent tuple pairs that satisfy the predicate as the output. This process is known as *refinement* and FACET refines each predicate using specialized algorithms for each operator. The output of a refinement is consumed as the input for refining the next predicate. At the end of processing all the predicates, we get all tuple pairs that satisfy all the predicates, and, thus, represent the violations.

**Example 4.** Consider the DC  $\varphi_3 : \neg(s.\text{State} = t.\text{State} \wedge s.\text{Salary} \leq t.\text{Salary} \wedge s.\text{FedTaxRate} > t.\text{FedTaxRate})$ . The refinement of predicate  $p_1 : s.\text{State} = t.\text{State}$  produces the set  $\{(\{t_2, t_3, t_4\}, \{t_2, t_3, t_4\})\}$  with a single ordered pair. This ordered pair represents the set of tuple pairs:  $(t_2, t_3), (t_2, t_4), (t_3, t_2), (t_3, t_4), (t_4, t_2), (t_4, t_3)$  since each of them satisfy  $p_1$ . Next, this singleton set is provided as input to predicate  $p_2 : s.\text{Salary} < t.\text{Salary}$  which produces a new set  $\{(\{t_4\}, \{t_2, t_3\}), (\{t_2\}, \{t_3\})\}$  since the Salary for  $t_4$  is smaller than both  $t_2$  and  $t_3$  but Salary for  $t_2$  is smaller than  $t_3$  only. Finally, we process predicate  $p_3 : (s.\text{FedTaxRate} > t.\text{FedTaxRate})$ . The output of the refinement by predicate  $p_3$  would be  $\{(\{t_4\}, \{t_2, t_3\})\}$  which represents the two violations of  $\varphi_3 : (t_4, t_2)$  and  $(t_4, t_3)$ .

FACET contains algorithms that are custom-designed for the different predicate structures. We now highlight the three key sources of inefficiency in FACET:

- (1) **Complexity of IEJoin.** FACET uses IEJoin [27] as the algorithm for processing inequalities. The algorithm is designed to process two inequalities at a time and operates on two sets of tuple pairs simultaneously. The runtime complexity of IEJoin is  $O(|R| \cdot |S|)$  for processing inequality joins between two relations  $R$  and  $S$  (although its space complexity is only  $O(|R| + |S|)$ ). As noted in [34], IEJoin is severely under-performing for predicates of low selectivity.
- (2) **Complexity of Hash-Sort-Merge.** Since IEJoin is designed for at least two predicates with inequality, FACET proposed two novel optimizations to process DCs with a single inequality predicate: **Hash-Sort-Merge (HSM)** and **Binning-Hash-Sort-Merge (BHSM)**. However, it can be shown that both HSM and BHSM still require a quadratic amount of running time and space in the worst-case. Similarly, processing of predicates containing disequality also requires quadratic time and space in the worst-case. In the experimental evaluation, we will demonstrate that this is not just a theoretical argument but manifests itself in reality even for constraints with as few as two predicates.
- (3) Since FACET processes one predicate at a time, it needs to make at least one full pass over the dataset. As we will see later, this is not always necessary for DC verification. Even for enumeration, we will show there exists efficient ways to process the constraints.

---

**Algorithm 1: DC VERIFICATION FOR HOMOGENEOUS CONSTRAINTS WITH EQUALITY PREDICATES**


---

**Input** : Relation  $\mathbf{R}$ , Homogeneous DC  $\varphi$  with only equality predicates.  
**Output**: True/False

- 1  $H \leftarrow$  empty hash table
- 2 **foreach**  $t \in \mathbf{R}$  **do**
  - /\* Project tuple  $t$  on the columns participating in equality predicates \*/
  - 3  $v \leftarrow \pi_{\text{vars}=(\varphi)}(t)$
  - 4 **if**  $v \notin H$  **then**
  - 5 |  $H[v] \leftarrow 0$
  - 6  $H[v] \leftarrow H[v] + 1$
  - 7 **if**  $H[v] > 1$  **then**
  - 8 | /\* Hash collision - violation detected \*/
  - 8 | **return False**
- 9 **return True**

---

## 4 RAPIDASH DESIGN AND ANALYSIS

In this section, we describe the general ideas underlying RAPIDASH followed by specific improvements and optimizations. Our algorithm builds appropriate data structures to store the input data (leveraging existing work on orthogonal range search), and issues appropriate queries to detect violations of a given DC. For ease of exposition and aiding readability, we will focus on the problem of DC verification throughout this section (item (1) in our problem statement in Section 2). At the end of the section, we point out how our main algorithm can be readily modified in a minimal way to support violation enumeration (item (2) in the problem statement). We defer a detailed description to the Appendix in [29].

### 4.1 Foundation

Before we delve into the details of our proposed algorithm and its relationship to orthogonal range search, we will provide some intuition behind the design of the algorithm taking as an example the simplest scenario: homogeneous constraints with equality predicates only. The constraint  $\varphi_1 : \neg(s.SSN = t.SSN)$  presented in the previous section is a qualifying constraint. The verification algorithm should return True over a relation  $\mathbf{R}$ , when SSN is a candidate key. If at least one pair of tuples in  $\mathbf{R}$  shares the same SSN value, then the algorithm should return False.

To evaluate whether such a tuple exists, we can incrementally populate a hash table tuple-by-tuple. The intuition is that if two tuples fall in the same hash partition then they have the same SSN value, and thus, we have identified a violation. The steps are presented in Algorithm 1. For each new tuple, we extract the SSN value and check whether it already exists in the hash table. If not, then we create a new entry with an associated count of 1. If there is already an entry then we increase the count by 1. If the count becomes greater than 1, we have identified two tuples with same SSN value and we return False. For example, in Table 1, the algorithm would create 4 hash table entries (one per SSN value), each with count 1, and thus, the constraint would evaluate to True. The algorithm works similarly with constraints having more than one equality predicates. This algorithm is straightforward and easy to understand, yet it grows more complex with the inclusion of non-equality

---

**Algorithm 2: DC VERIFICATION FOR HOMOGENEOUS CONSTRAINTS WITH INEQUALITY PREDICATES**


---

**Input** : Relation  $\mathbf{R}$ , Homogeneous DC  $\varphi$  with inequalities.  
**Output**: True/False

- 1  $k \leftarrow$  #number of columns appearing in inequality predicates
- 2  $H \leftarrow$  empty hash table
- 3 **foreach**  $t \in \mathbf{R}$  **do**
  - /\* Project tuple  $t$  on the columns participating in equality predicates \*/
  - 4  $v \leftarrow \pi_{\text{vars}=(\varphi)}(t)$
  - 5 **if**  $v \notin H$  **then**
  - 6 |  $H[v] \leftarrow$  new ORTHOGONALRANGESearchTREE( $k$ )
  - /\* Project tuple  $t$  on the columns participating in inequality predicates \*/
  - 7  $z \leftarrow \pi_{\text{vars}(\varphi) \setminus \text{vars}=(\varphi)}(t)$
  - 8 **if**  $\neg H[v].\text{isEmpty}()$  **then**
    - /\* Evaluate violations through two range search queries \*/
    - 9  $L, U, L', U' \leftarrow$  CreateRangeSearchQueries( $z, \phi$ )
    - 10 **if**  $H[v].\text{booleanRangeSearch}(L, U) \vee$
    - 10 |  $H[v].\text{booleanRangeSearch}(L', U')$  **then**
    - 11 | /\* Violation detected \*/
    - 11 | **return False**
    - /\* Insert  $z$  into the range tree \*/
    - 12 |  $H[v].\text{insert}(z)$
  - 13 **return True**
- 14 **PROCEDURE** CreateRangeSearchQueries( $t, \varphi$ )
  - /\*  $L$  and  $U$  are indexed by the non-equality predicates  $p_i$ . Both are of size  $k$  \*/
  - 15  $L \leftarrow (-\infty, \dots, -\infty), U \leftarrow (\infty, \dots, \infty)$
  - /\* Create range search query ( $L, U$ ) \*/
  - 16 **foreach** inequality predicate  $p_i \in \varphi$  **do**
    - 17 **if**  $p_i.\text{op}$  is  $<$  or  $\leq$  **then**
    - 18 |  $U_i \leftarrow \pi_{p_i.\text{col}}(t)$
    - 19 **if**  $p_i.\text{op}$  is  $>$  or  $\geq$  **then**
    - 20 |  $L_i \leftarrow \pi_{p_i.\text{col}}(t)$
  - /\* Create inverted range search query ( $L', U'$ ) \*/
  - 21  $U' \leftarrow L, L' \leftarrow U$
  - 22 flip  $-\infty$  to  $\infty$  and  $\infty$  to  $-\infty$  in  $U'$  and  $L'$  respectively
  - 23 **return**  $L, U, L', U'$

---

predicates. It is at this juncture that orthogonal range search becomes relevant. Before we dive deeper into these scenarios, we will first provide some background on orthogonal range search.


### 4.2 Orthogonal Range Search

In this section, we present some background on orthogonal range search [8, 9]. Given a totally ordered domain  $\mathbb{N}$ , let  $A \subseteq \mathbb{N}^k$ , for some  $k \geq 1$ , be a set of size  $N$ . Let  $L = (\ell_1, \dots, \ell_k)$  and  $U = (u_1, \dots, u_k)$  be such that  $L, U \in \mathbb{N}^k$  and  $\ell_i \leq u_i$  for all  $i \in [k]$ .

**DEFINITION 1.** An orthogonal range search query is denoted by  $(L, U)$ , and its evaluation over  $A$  consists of enumerating the set

$$Q(A) = \{a \in A \mid \bigwedge_{i \in [k]} \ell_i \text{ op}_1 a_i \text{ op}_2 u_i\}$$

where  $\text{op}_1, \text{op}_2 \in \{=, <, \leq\}$ .



**Figure 3: Salary and FedTaxRate for each tuple in Tax. The grey (upper left quadrant centered at  $t_2$ ) and blue shaded areas (lower right quadrant centered at  $t_2$ ) show the regions where the tuples that could form a violation with  $t_2$  lie.**

In other words,  $L$  and  $U$  form an axis-aligned hypercube in  $k$  dimensions, and  $Q(A)$  reports all points in  $A$  that lie on/within the hypercube. The Boolean version of the orthogonal range search problem consists of determining if  $Q(A)$  is empty or not.

**Example 5.** Consider the Table Tax from Example 1. Let  $A$  be the set of two-dimensional points obtained by projecting Tax on Salary and FedTaxRate. Let  $L = (3500, 5)$  and  $U = (4500, 25)$ . Then, the orthogonal range query  $(L, U)$  is asking for all points such that the Salary is between 3500 and 4500, and the FedTaxRate is between 5 and 25. In Table Tax, only  $t_4$  satisfies the criteria (its values of Salary and FedTaxRate are 4000 and 22 respectively). Thus, the result of the orthogonal range search query  $(L, U)$  is  $\{(4000, 22)\}$ .

The two seminal data structures for orthogonal range search are range trees and kd-trees [8]. We will review their complexity and trade-offs when analyzing the complexity of our specific algorithm.

### 4.3 Verification Algorithm

In this section, we present our verification algorithm that leverages prior work on orthogonal range search. The algorithm builds on top of the ideas behind Algorithm 1 but extends them to cover for homogeneous constraints that contain both equalities and inequalities. Without loss of generality, we will assume that none of the predicates contain disequality. This assumption will be removed later. Finally, we assume that the categorical columns in  $R$  have been dictionary-encoded to integers, a standard assumption in line with prior work [33, 35].

The algorithm preserves the core concepts of Algorithm 1, namely the use of a hash table and early termination in case of violations. The primary modification involves incorporating orthogonal range search indexes to identify violations stemming from inequality predicates. The algorithm is presented in Algorithm 2.

We first compute the number  $k$  of columns in inequality predicates (we assume  $k > 0$  as Algorithm 1 covers the case where  $k = 0$ ). We then proceed similarly as before: project each tuple on the columns participating in equality predicates ( $v$  on line 4) and evaluate whether the resulting tuple has been seen before (line 5). If not, we create a new hash table entry whose value now, instead of being an integer counter, is a range search tree of  $k$  dimensions (line 6). This tree will be used to index the  $k$ -dimensional tuples containing the columns participating in inequality predicates and identify violations. Before we further delve into the pseudocode, we explain the main intuition through an example.

**Table 2: Data structure parameter on input of size  $n$  [31].  $k$  is the number of dimensions of the points inserted in the tree.**

DS	Insertion $I(n)$	Answering $T(n)$	Space $S(n)$
Range tree	$O(\log^k n)$	$O(\log^k n)$	$O(n \cdot \log^{k-1} n)$
kd-tree	$O(\log n)$	$O(n^{1-\frac{1}{k}})$	$O(n)$

**Example 6.** Consider the relation Tax from Example 1 and the DC  $\phi_3 : \neg(s.State = t.State \wedge s.Salary \leq t.Salary \wedge s.FedTaxRate > t.FedTaxRate)$ , which contains one equality and two inequality predicates ( $k = 2$ ). For simplicity, we will omit the details of how range search works in this example but instead present it later. Algorithm 2 will first start with the equality predicate, and place  $t_1$  in a hash partition by hashing  $t_1.State = New York$  and initialize a 2-dimensional range search tree for that partition (line 6). Since this tree is empty, we do not perform any violation detection (line 8) and we insert  $(t_1.Salary, t_1.FedTaxRate) = (3000, 20)$  in the tree (line 12).


Next, we process  $t_2$ , which is placed in a different hash partition since  $t_2.State = Wisconsin$  and we initialize a new range tree for this partition. As in the previous step, we then insert  $(t_1.Salary, t_1.FedTaxRate) = (5000, 15)$  in the tree (call this step  $\textcircled{A}$ ). When  $t_3$  is processed, it is placed in the same partition as  $t_2$  since they have the same State value.

At this point, we have two tuples in the same hash partition, and thus we need to consider the inequality predicates in the DC to establish whether there is a DC violation. This is where the orthogonal range search tree is leveraged. Such a violation would occur in two scenarios: (1) if any tuple in the tree ( $t_2$  in this case) has Salary less than  $t_3.Salary = 6000$  but FedTaxRate more than  $t_3.FedTaxRate = 20$ , or (2) if any tuple in the tree has Salary more than  $t_3.Salary$  but FedTaxRate less than  $t_3.FedTaxRate$ .

To identify whether any of the two scenarios above is true, we perform two orthogonal range search queries using the values of  $t_3$  to probe the index. More specifically, we perform a search with  $L = (-\infty, 20)$  and  $U = (6000, \infty)$  (scenario 1). Then, we also search in the inverted range  $L' = (6000, -\infty)$  and  $U' = (\infty, 20)$  (scenario 2). Since  $t_2$  does not lie in the desired range, both range searches return false. Figure 3 visualizes this result. Since there is no violation, we insert  $t_3(6000, 20)$  in the tree (call this step  $\textcircled{B}$ ). Finally,  $t_4$  (highlighted in red in Figure 3) is processed and placed in the same partition as  $t_2$  and  $t_3$  because of same value in State column. We again initiate two range search queries based on  $t_4$  Salary = 4000 and FedTaxRate = 22. The queries would be  $L = (-\infty, 22)$ ,  $U = (4000, \infty)$  and  $L' = (4000, -\infty)$ ,  $U' = (\infty, 22)$ . Then,  $t_2$  and  $t_3$  form a violation wrt.  $t_4$  since both the points represent a higher salary than 4000 but a lower tax rate than 22, and Line 11 returns False.

Algorithm 2 formalizes the process described in the example above. When it comes to evaluating inequality predicates, we generate two appropriate range queries based on the values of the current tuple and the operator type in the inequality predicates (line 9) and then search the range tree (line 10). If any of the range search queries returns True, a DC violation is detected and the algorithm terminates (line 11). Otherwise, the tuple is inserted in the tree (line 12), and the algorithm continues with the next tuple in the relation.

Seminal work by Overmars [31] showed that using range trees and kd-trees, one can design an algorithm with the parameters as



**Figure 4: An illustration of insertion and search in a 2D range tree that stores (Salary, FedTaxRate). The primary tree with rectangle nodes stores Salary, and the secondary trees with circle nodes store FedTaxRate. The Roman numbers denote the visiting order during the search.**

shown in Table 2. We now demonstrate an example of how range trees are used by Algorithm 2 when performing step **A** and step **B** in Example 6.

**Example 7.** In Figure 4, we illustrate insertion and search in a range tree. The tree is two-dimensional and stores (Salary, FedTaxRate) points. We show the process of inserting  $t_2, t_3$  (i.e. performing step **A** and step **B**) from Example 6) and a range search for points whose Salary  $\geq 4000$  and FedTaxRate  $< 22$ . In the range tree, both the primary tree (which stores Salary) and the secondary tree (which stores FedTaxRate) are binary search trees. Leaf nodes store the inserted data, and each internal node stores the smallest value in its right subtree. Each node in the primary tree is linked to a secondary tree, which stores the FedTaxRate value of all the points present in the subtree rooted at this node. When we insert each point, we find the insert position in the primary tree, create a leaf node to store the inserted value, and an internal node to connect the new leaf node and the leaf node at the insert position. We also update all the secondary trees for nodes in the path from the root to the insert position. When we perform the range search, we look for nodes whose values lie in the range by pre-order traversal. For instance, the search  $L = (4000, -\infty), U = (\infty, 22)$  is performed by going from root node *i* to the left child node *ii*. Since the Salary value stored in node *ii* is in the range, we go to the secondary tree linked to it. Finally, we find node *iii* whose FedTaxRate value is less than 22 and return true for the query.

As another example, suppose we also insert  $t_4$  in the tree and search for a point where Salary  $\leq 7000$  and FedTaxRate  $> 20$ . At node *II*, we go to its secondary tree directly instead of traversing its descendants since all the points stored in the subtree rooted at node *II* have Salary  $\leq 5000$ , which are within the search range for Salary. At node *III*, we return true since the minimum FedTaxRate stored in its right subtree is 22, which is greater than 20.

Due to space limitation, we establish the correctness of Algorithm 2, and show the analysis of time and space complexity in our technical report [29]. We state our main result as follows.

**Theorem 1.** Algorithm 2 runs in time  $O(|\mathbf{R}| \cdot (I(|\mathbf{R}|) + T(|\mathbf{R}|)))$  and uses space  $S(|\mathbf{R}|)$  when using range tree or kd-tree with parameters as shown in Table 2.

With range trees, the running time is  $O(|\mathbf{R}| \cdot \log^k |\mathbf{R}|)$  and space usage is  $O(|\mathbf{R}| \log^{k-1} |\mathbf{R}|)$ ; for kd-trees, the running time is  $O(|\mathbf{R}|^{2-\frac{1}{k}})$  and space requirement is  $O(|\mathbf{R}|)$ .

**Comparison with FACET.** Our approach is superior to FACET in three respects. First, we use polynomially less space and time in the worst-case. Second, there exist instances where our algorithm saves a significant amount of time and space by early termination<sup>3</sup>.

**Proposition 1.** For every homogeneous DC  $\phi$  with at least one non-equality predicate, there exists a relation  $\mathbf{R}$  such that Algorithm 2 takes  $O(1)$  time and FACET requires  $\Omega(|\mathbf{R}|)$  time.

Although Proposition 1 may seem like an obscure theoretical argument, the behavior is commonly observed in the real world. We empirically demonstrate this in our experiments. The third aspect concerns the space usage. It turns out that the space requirement of FACET is relation-dependent, with no way of predicting how much memory will actually be needed, an artifact of the lack of nontrivial provable bounds on its time and space requirement. If the machine has only limited amount of memory, FACET will be unable to complete the refinements and fail. On the other hand, our solution allows verification with linear space using kd-trees. This flexibility is important for resource-constrained production scenarios.

#### 4.4 Heterogeneous Predicates

In this section, we present adjustments to Algorithm 2 to handle heterogeneous predicates.

**Example 8.** Continuing our study of Tax from Table 1, consider the following constraint:  $\psi_2 : \neg(s.\text{Salary} < t.\text{FedTaxRate})$  which says that all values of the Salary column must be greater than or equal to any value of FedTaxRate column.

This is an example of a heterogeneous predicate that cannot be handled by Algorithm 2 (notice that both  $s$  and  $t$  are referenced and the columns are different). Along the same lines, heterogeneous comparison constraints over date-related columns have also been found to be useful in our production settings. For example, over the TPC-H schema, [36] identified the heterogeneous constraint  $\neg(s.\text{Receiptdate} \geq t.\text{Shipdate} \wedge s.\text{Shipdate} \leq t.\text{Receiptdate})$  which represents the business logic that a new order is shipped only after all the previously shipped orders are received, i.e., the [Shipdate, Receiptdate] intervals of the orders never overlap. These real-world scenarios underscore the need for going beyond simple homogeneous constraints considered in the previous section.

<sup>3</sup>In this paper, we use the term early termination to mean that the algorithm can avoid scanning the entire relation for some instances.

We now discuss how our algorithm can be extended to handle heterogeneous predicates of the form  $s.A \text{ op } t.B$ , where  $\text{op}$  is  $=, <, \leq, >, \text{ or } \geq$ . We note that an equality predicate  $s.A = t.B$  is equivalent to  $s.A \leq t.B \wedge s.A \geq t.B$  and by supporting inequalities, we can support heterogeneous equality predicates as well.

Now, let's look at how to handle heterogeneous inequality predicates. Suppose that  $\varphi$  has a predicate  $s.C < t.D$ . If  $C = D$ , then the predicate is homogeneous and building a 1-dimensional range search data structure is enough. However, when  $C \neq D$ , we need to index in 2 dimensions ( $C$  and  $D$ ). Additionally, we need to adjust our procedure for computing the two search range queries ( $L, U$ ), ( $L', U'$ ) to consider the different attributes present in the predicate. Let us look at an example.

**Example 9.** Consider the DC  $\psi_2$  from Example 8 that checks that all the Salary values must be greater than any FedTaxRate value. We will create one 2-dimensional range search data structures in which we will store value of (Salary, FedTaxRate). Suppose we are processing tuple  $t_2$  with Salary = 5000 and FedTaxRate = 15. We first do a range search to check if there is a tuple with Salary that is less than  $t_2.$ FedTaxRate denoted as  $L = (-\infty, -\infty), U = (15, \infty)$ . Additionally, we check whether there is a tuple with FedTaxRate that is larger than  $t_2.$ Salary denoted as  $L' = (-\infty, 5000), U' = (\infty, \infty)$ . If any of the two range search queries return True, we have found a violation.

Given this example, it becomes clear that we can easily extend Algorithm 2 to account for heterogeneous predicates by simply adjusting the procedure to create the two range search queries that are used to detect violations. Algorithm 3 shows the updated query generation procedure. The main idea is that if  $\varphi$  has a predicate  $s.C < t.D$ , then when we process a new tuple  $r$ , the upper-bound for attribute  $C$  is set to  $r.D$  in the forward check, and the lower-bound for attribute  $D$  is set to  $r.C$  in the inverted check (because we are comparing attribute  $C$  of  $s$  with attribute  $D$  of  $t$  in the predicate). It is worth noting that this algorithm can be applied for homogeneous constraints as well. Note that, when  $C = D$ , we recover our original procedure presented in Algorithm 2. Thus, we can safely handle both types of constraints by simply replacing the CreateRangeSearchQueries function in Algorithm 2 with the one presented here. Finally, we also note that the new generalization also extends our algorithm to handle the case when attributes occur in more than one predicate.

## 4.5 Supporting disequalities

So far, we have discussed how to support homogeneous and heterogeneous predicates with equalities or inequalities. In this section, we show that it is possible to apply orthogonal range search techniques to support disequalities as well. This type of operator is quite important as it is required for specifying functional dependencies. We demonstrate that with the below example.

**Example 10.** Consider the constraint:  $\psi_2 : \neg(s.\text{Zip} = t.\text{Zip} \wedge s.\text{StateCode} \neq t.\text{StateCode})$ . This DC represents the functional dependency  $\text{Zip} \rightarrow \text{StateCode}$ . The constraint contains a disequality predicate which cannot be handled by any of our previous algorithms.

We now discuss how we can support verification of such predicates. Any predicate  $s.A \neq t.B$  can be written as a union of two predicates:  $(s.A < t.B) \vee (s.A > t.B)$ . Therefore, a DC containing  $\ell$

---

### Algorithm 3: GENERALIZED RANGE SEARCH QUERY GENERATION THAT COVERS BOTH HOMOGENEOUS AND HETEROGENEOUS PREDICATES WITH INEQUALITIES

---

**Input** : A tuple  $r$  from relation  $R$ , DC  $\varphi$

**Output**: Two range search queries (normal and inverted)

```

1 PROCEDURE CreateRangeSearchQueries( $r, \varphi$ )
2    $L, L' \leftarrow (-\infty, \dots, -\infty), U, U' \leftarrow (\infty, \dots, \infty)$  /*  $L, U, L', U'$ 
   are indexed by attributes of  $R$  that appear in inequality
   predicates */
3   foreach inequality predicate  $s.C \text{ op } t.D$  in  $\varphi$  do
4     if  $\text{op}$  is  $<$  or  $\leq$  then
5        $U.C \leftarrow r.D$ 
6        $L'.D \leftarrow r.C$ 
7     if  $\text{op}$  is  $>$  or  $\geq$  then
8        $L.C \leftarrow r.D$ 
9        $U'.D \leftarrow r.C$ 
10  return  $L, U, L', U'$ 

```

---

predicates with  $\text{op}$  as  $\neq$  can be equivalently written as a conjunction of  $2^\ell$  DCs containing no disequality operator.

If the original homogeneous DC contains no inequality predicate, then it is possible to reduce the number of equivalent DCs from  $2^\ell$  to  $2^{\ell-1}$ . The idea is that a violation  $(s, t)$  is symmetric (i.e.  $(t, s)$  is also a violation) if the DC contains only equality and disequality predicates. Therefore, when converting a DC to only have inequalities, it suffices to expand  $(s.A \neq t.A)$  to just  $(s.A < t.A)$  for one last disequality predicate instead of  $(s.A < t.A) \vee (s.A > t.A)$ .

**Proposition 2.** Given a homogeneous DC  $\varphi$  with only equality and  $\ell$  disequality predicates, there exists an equivalent conjunction of  $2^{\ell-1}$  DCs that contain only equality and inequality predicates.

## 4.6 Optimizations

In the previous sections, we have presented how we can support verification of both homogeneous and heterogeneous constraints with equalities, inequalities and disequalities. We now present optimizations for specific type of constraints that can improve the overall performance.

**Optimization for a single inequality.** If a DC has row homogeneous equality predicates and at most one predicate (homogeneous or heterogeneous) containing an inequality, then the verification can be done in linear time. The key idea is that for predicate containing inequality  $s.A \text{ op } t.B$ , it is enough to keep track of the running minimum and maximum values for values seen in columns  $A$  and  $B$  respectively as we process the relation. To illustrate the idea, we use an example.

**Example 11.** Consider the functional dependency  $\psi_1$  from Example 10 and the Tax table. Based on proposition (2), we can convert the disequality predicate into inequality to get the DC  $\neg(s.\text{Zip} = t.\text{Zip} \wedge s.\text{StateCode} < t.\text{StateCode})$ . Since both columns in the inequality predicate are the same (StateCode), we have  $A = B$ . Let us focus on rows of Tax with Zipcode=53703. We will keep track of the min and max value of StateCode for each of these rows. When  $t_2$  is processed, we set  $\text{min} = \text{max} = t_2.\text{Zipcode} = 02$ . For  $t_3$ , we observe that since  $t_3$  also has StateCode=02, the predicate  $t_2.\text{StateCode} < t_3.\text{StateCode}$  is false and thus no violation is found. The values of min

---

**Algorithm 4:** DC VERIFICATION FOR DCs WITH ROW-HOMOGENEOUS EQUALITY AND ONE INEQUALITY PREDICATE
 

---

**Input** : Relation  $R$ , DC  $\varphi$  containing equality predicates of form  $s.C = t.C$  (for some  $C$ ) and **one** inequality predicate  $p$  of form  $s.A \text{ op } t.B$

**Output**: True/False

```

1  $H \leftarrow$  empty hash table
2 foreach  $t \in R$  do
3    $v \leftarrow \pi_{\text{vars}=(\varphi)}(t)$ 
4   if  $v \notin H$  then
5      $H[v] \leftarrow (+\infty, +\infty, -\infty, -\infty)$  /* four-tuple represents
6        $(\min_A, \min_B, \max_A, \max_B)$  for  $v$  */
7     if  $(p.\text{op} \in \{<, \leq\} \wedge H[v].\min_A \text{ op } t.B) \vee (p.\text{op} \in \{>, \geq\} \wedge H[v].\max_A \text{ op } t.B)$  then
8       return false
9     if  $(p.\text{op} \in \{<, \leq\} \wedge t.A \text{ op } H[v].\max_B) \vee (p.\text{op} \in \{>, \geq\} \wedge t.A \text{ op } H[v].\min_B)$  then
10      return false
11     $H[v].\min_A \leftarrow \min\{H[v].\min_A, t[A]\}$  /* modify  $\min_A$  */
12     $H[v].\min_B \leftarrow \min\{H[v].\min_B, t[B]\}$  /* modify  $\min_B$  */
13     $H[v].\max_A \leftarrow \max\{H[v].\max_A, t[A]\}$  /* modify  $\max_A$  */
14     $H[v].\max_B \leftarrow \max\{H[v].\max_B, t[B]\}$  /* modify  $\max_B$  */
15 return true

```

---

and  $\max$  remain unchanged. Finally, for  $t_4$ , we also do not find a violation since  $t_4.\text{StateCode} = \emptyset 2$ . However, if there was a different row  $t'_4$  such that  $t'_4.\text{Zip} = 53703$  and  $t'_4.\text{StateCode} = 03$ ,  $t_2.\text{StateCode} < t'_4.\text{StateCode}$  would become true and thus, we would have found a violation since all predicates are true for a tuple pair.

Algorithm 4 shows the algorithm. Like the previous algorithms, we begin by partitioning the input into a hash table based on the equality predicates. Let the inequality predicate be  $s.A \text{ op } t.B$ . The main idea is to maintain the running minimum and maximum values for the Columns A and B for each partition of the input. Since the comparison is one-dimensional, it is sufficient to compare against the minimum (or maximum) value. The algorithm makes only one pass over the entire dataset and the overall time complexity is  $O(|R|)$ . While this optimization is simple, it has important implications. In particular, popular constraints such as functional dependencies (FD) are DCs that contain exactly one inequality predicate. Algorithm 4 recovers the standard linear time algorithm to verify FDs [23]. However, it is unclear that FACET, in its present form, can achieve the same provable guarantee.

#### 4.7 Enumerating violations

In this section, we discuss violation enumeration ((2) in our problem statement). Recall that the Algorithm 2 with minor modifications already supports the enumeration of the violations. In this section, we describe an optimization that improves the time and space complexity of enumerating the DC violations. The first observation we make is that, unlike verification, enumeration usually requires examining every tuple of the relation. Therefore, we can make use of sort-based optimizations. We demonstrate with an example.

**Example 12.** Consider  $\varphi_3 : \neg(s.\text{State} = t.\text{State} \wedge s.\text{Salary} \leq t.\text{Salary} \wedge s.\text{FedTaxRate} > t.\text{FedTaxRate})$  from Example 3 and

the rows in table Tax with State=Wisconsin. The first step is to sort the three rows in increasing order of Salary, which creates the ordering  $t_4, t_2, t_3$ . The key observation here is that once sorting has been performed, it is guaranteed that any tuple can only form a violation (wrt. the Salary predicate) with a tuple that appears after it in the sort order. At this point, the Salary attribute can be completely removed from consideration as the corresponding predicate will always be satisfied. We can now resume our usual processing by constructing a one dimensional binary search tree on FedTaxRate attribute. We first insert  $t_4.\text{FedTaxRate} = 22$  into the tree. For the next tuple in the order  $t_2$ , which has  $t_2.\text{FedTaxRate} = 15$ , we ask the tree to enumerate all tuples that have value greater than 15, a standard operation on a binary search tree. We find the  $t_4$  is such a tuple and thus report  $(t_4, t_2)$  as a violation to the user. Then, we insert  $t_2.\text{FedTaxRate}$  in the tree. Finally,  $t_3$  is processed and we search for all values in the tree that with FedTaxRate greater than 20 and report  $(t_4, t_3)$  as a violation.

A detailed description of this optimization and the corresponding enumeration algorithm is presented in [29]. We also have the following complexity result.

**Theorem 2.** Let  $k$  be the number of columns that occur in predicates containing an inequality. Then, our enumeration algorithm runs in time  $O(|R| \cdot (I(|R|) + T(|R|) + \log |R|) + K)$  to enumerate  $K$  violations and uses space  $S(|R|)$  when using range tree or kd-tree with parameters as shown in Table 2 with the number of dimensions as  $k-1$ .

#### 4.8 Discussion

So far, we saw how we can leverage orthogonal range search to support DC constraint verification and enumeration. Our approach relies on creating classes of constraints (homogeneous with equalities, homogeneous with both equalities and inequalities, heterogeneous) and devising specialized algorithms and optimizations for each class of constraints. An alternative approach would be to create a  $k$ -dimensional range search index, where  $k$  is the distinct number of column in all predicates (as opposed to inequality predicates only in RAPIDASH). We could then use this index to perform appropriate range search queries and detect violations. We call this approach *vanilla range search* and we experimentally compare it with RAPIDASH in Section 5.

The main advantage of our techniques over vanilla range search is that they reduce the dimensionality of the range search tree (value of  $k$ ) which results in significant savings in performance and space. As shown in Table 2, although range search data structures in their most efficient form only add a  $(\log |R|)^k$  term in the time complexity, the factor is multiplicative. Thus, even for relatively small datasets containing  $1M \sim 2^{20}$  rows and a DC with  $k = 1$  homogeneous predicates (such as an FD), we will have  $(\log_2 2^{20})^1 = 20$  as a multiplicative factor which leads to a blowup in both space usage and time in the worst-case. Thus, reducing the value of  $k$  whenever possible, can lead to improved performance and allow scaling to larger datasets. This intuition based on the theoretical properties of range search trees is further validated in our experiments.

Our approach categorizes constraints and comes up with approaches to reduce (or even eliminate)  $k$ . This is achieved by: 1) leveraging hash tables for homogeneous equality predicates and only pairing them with range trees when inequalities are present, 2) reducing dimensions in the presence of disequalities when possible



**Table 3: List of denial constraints used in experiments for each dataset.**

Dataset	Cardinality	#Columns	DC number	Denial constraint
Tax	1M	12	$c_1$	$\neg(s.\text{AreaCode} = t.\text{AreaCode} \wedge s.\text{Phone} = t.\text{Phone})$
Tax	1M	12	$c_2$	$\neg(s.\text{ZipCode} = t.\text{ZipCode} \wedge s.\text{City} \neq t.\text{City})$
Tax	1M	12	$c_3$	$\neg(s.\text{State} = t.\text{State} \wedge s.\text{HasChild} = t.\text{HasChild} \wedge s.\text{ChildExemp} \neq t.\text{ChildExemp})$
Tax	1M	12	$c_4$	$\neg(s.\text{State} = t.\text{State} \wedge s.\text{Salary} > t.\text{Salary} \wedge s.\text{Rate} < t.\text{Rate})$
TPC-H	1M	12	$c_5$	$\neg(s.\text{Customer} = t.\text{Supplier} \wedge s.\text{Supplier} = t.\text{Customer})$
TPC-H	1M	12	$c_6$	$\neg(s.\text{Receiptdate} \geq t.\text{Shipdate} \wedge s.\text{Shipdate} \leq t.\text{Receiptdate})$
TPC-H	1M	12	$c_7$	$\neg(s.\text{ExtPrice} > t.\text{ExtPrice} \wedge s.\text{Discount} < t.\text{Discount})$
TPC-H	1M	12	$c_8$	$\neg(s.\text{Qty} = t.\text{Qty} \wedge s.\text{Tax} = t.\text{Tax} \wedge s.\text{ExtPrice} > t.\text{ExtPrice} \wedge s.\text{Discount} < t.\text{Discount})$
NCVoter	1M	67	$c_9$	$\neg(s.\text{countyid} = t.\text{countyid} \wedge s.\text{countydesc} \neq t.\text{countydesc})$
NCVoter	1M	67	$c_{10}$	$\neg(s.\text{ageatyearend} > t.\text{birthyear})$
NCVoter	1M	67	$c_{11}$	$\neg(s.\text{statuscd} = t.\text{statuscd} \wedge s.\text{voterdesc} = t.\text{voterdesc} \wedge s.\text{reasoncd} \neq t.\text{reasoncd})$
NCVoter	1M	67	$c_{12}$	$\neg(s.\text{mailzipcode} = t.\text{zipcode} \wedge s.\text{statecd} \neq t.\text{mailstate})$
$D_1$	50M	28	$\varphi_{1,1}$	$\neg(s.A = t.A \wedge s.B = t.B \wedge s.C \neq t.C \wedge s.D \neq t.D)$
$D_1$	50M	28	$\varphi_{1,2}$	$\neg(s.C = t.C \wedge s.E = t.E \wedge s.F = t.F \wedge s.G \neq t.G \wedge s.H \neq t.H)$
$D_1$	50M	28	$\varphi_{1,3}$	$\neg(s.B = t.B \wedge s.I = t.I \wedge s.J = t.J \wedge s.K \neq t.K \wedge s.L \neq t.L)$
$D_1$	50M	28	$\varphi_{1,4}$	$\neg(s.A = t.A \wedge s.I = t.I \wedge s.M > t.M \wedge s.N \neq t.N)$
$D_2$	25M	28	$\varphi_{2,1}$	$\neg(s.A = t.A \wedge s.B = t.B \wedge s.C \geq t.C \wedge s.D \leq t.D \wedge s.E \leq t.E \wedge s.F \geq t.F \wedge s.G > t.G)$
$D_2$	25M	28	$\varphi_{2,2}$	$\neg(s.A \neq t.A \wedge s.B = t.B \wedge s.H \leq t.H \wedge s.F \geq t.F \wedge s.G \geq t.G)$
$D_2$	25M	28	$\varphi_{2,3}$	$\neg(s.A = t.A \wedge s.I \neq t.I \wedge s.D \leq t.D \wedge s.G \geq t.G \wedge s.J = t.J)$
$D_2$	25M	28	$\varphi_{2,4}$	$\neg(s.C \leq t.C \wedge s.D \leq t.D \wedge s.K = t.K)$

(see Proposition (2)), completely eliminating range search trees for constraints with single inequalities (such as FDs) and leveraging running min and max values instead, 3) reducing dimensionality through sorting for violation enumeration.

**Multi-constraint execution.** In this work, we focus on single-core and independent processing of DCs (i.e. no work sharing). Before we end the subsection, we briefly comment on the multi-constraint execution capabilities. FACET supports processing multiple DCs at the same time and parallel processing. Our framework could also be expanded to leverage common prefixes of two or more DCs. If the prefixes are equality predicates, the hash table can be shared. For tree-based data structures, the branching of the internal nodes of the tree can be carefully controlled to reuse the nodes that are common to multiple DCs. Similarly, although range search structures can use parallel processing capabilities [10] in theory, their behavior in practice needs further investigation. Further, since FACET performs processing in a *columnar* fashion but our algorithm is *row-oriented*, the impact of the different design choices needs to be explored rigorously.

## 5 EXPERIMENTAL EVALUATION

In this section, we report the results of our experimental evaluation. In particular, we seek to answer the following questions:

- (Q.1) How does RAPIDASH, FACET and vanilla orthogonal range search compare in terms of performance?
- (Q.2) What is the performance improvement (time and space) of the RAPIDASH verification and enumeration algorithm compared to FACET on open-source datasets used in prior work [34]?

(Q.3) What is the performance and scalability of RAPIDASH and FACET over large scale real-world production datasets with complex constraints?


(Q.4) What is the impact of varying number of violations on FACET and both variants of RAPIDASH?

### 5.1 Experimental Setting

**Datasets, DCs, and Hardware.** We perform experiments on both open-source datasets and production datasets. For open-source datasets, we use the Tax, TPC-H, and NCVoter, with a total of 12 DCs that were identified by prior works [33–35] to be representative of what is usually seen in production settings (as defined by experts or discovered from data). We also use two production datasets (related to banking records and document shipping) from Microsoft customers interested in applying DC verification on their data. Each dataset contains a mix of categorical, numeric, and datetime columns. For both production datasets, we pick 3 DCs by taking a 10% sample of each dataset and discover DCs that are true over the sample. The fourth DC (denoted by  $\varphi_{i,4}$  for dataset  $D_i$ ) holds over the full dataset. Table 3 lists a total of 20 DCs over all datasets that we use in our experiments<sup>4</sup>. Note that constraints  $c_5, c_6, c_{10}, c_{12}$  are examples of heterogeneous constraints. We ran all experiments on an Intel(R) Xeon(R) W-2255 CPU @ 3.70GHz machine with 128GB RAM running Windows 10 Enterprise (version 22H2). All of our experiments are executed over a single core and in the main memory setting.

**Evaluation Metrics.** We perform experiments for both DC verification and enumeration ((1) and (2) in our problem statement respectively) using RAPIDASH and FACET and report the end-to-end

<sup>4</sup>The column names in the DCs have been omitted due to security and privacy concerns.



**Figure 5: Running time for verification for vanilla range search tree, FACET, and RAPIDASH( $\perp$ ).**

running time and space consumption. For enumeration performance, we report the total time to count and return the number of all the violations (same approach as FACET [34]) to avoid output materialization cost and focus on understanding the intrinsic hardness. All reported running times are the trimmed mean of five independent executions after the dataset has been loaded in memory.

**Implementation.** Similar to prior work, RAPIDASH is implemented in Java. We use a standard implementation of orthogonal range trees (referred to as RAPIDASH( $\perp$ )) and kd-trees (referred to as RAPIDASH(kd)). We implemented FACET in Java using the Metanome infrastructure from [11] and [33], with all optimizations enabled as described in [34]. We manually verified that the performance of our implementation is in line with the numbers reported in [34] accounting for hardware differences. Further, to ensure an apples-to-apples comparison, for DC verification, we ensure that FACET execution terminates as soon as the first violation is found.

## 5.2 FACET vs. RAPIDASH vs Vanilla Range Search

In this section, we answer Q.1, i.e., what is the performance of RAPIDASH vs FACET vs vanilla range search (as described in Section 4.8) for constraint verification. Results for constraint enumeration are included in [29] and follow similar trend as verification.

Figure 5 shows the running time of using a range tree directly (i.e. the vanilla approach), FACET, and RAPIDASH( $\perp$ ) on all 12 DCs over the three open-source datasets. For 8 out of 12 DCs, directly applying range search was slower than FACET by up to 100 $\times$ . While this may seem surprising at first glance, it is explained by the blowup due to logarithmic factors as described in Section 4.8, an observation that was also made in [40]. For the remaining 4 DCs, due to a large number of violations, both RAPIDASH( $\perp$ ) and vanilla range search tree terminated early and thus the running time is close. We note here that RAPIDASH( $\perp$ ) still has an advantage because our optimizations ensure that the data structure is created over a smaller number of dimensions. For all DCs, RAPIDASH( $\perp$ ) was the fastest out of all three. This result demonstrates that even after making the connection between DC processing and orthogonal range search, non-trivial optimizations are required to obtain competitive performance through range search algorithms. In the subsequent experiments, we exclude vanilla range search from the comparisons because RAPIDASH already demonstrates superior performance. We defer the interested reader to [29] for more experiments with vanilla range search. Instead, we concentrate on comparing RAPIDASH variants against the state-of-the-art algorithm, FACET.

## 5.3 Evaluation on Open-Source Datasets

In this section, we answer Q.2 by diving deeper into the trade-offs between the two implementation of range trees and kd-trees. Figure 6 shows the running time for FACET and RAPIDASH. Let us fix our attention to Figure 6a. Our first observation is that for  $c_1$  and  $c_5$  DCs (which contain only equality and thus, both FACET and RAPIDASH take a provably linear amount of time in theory), RAPIDASH is faster by 2 $\times$  for verification. This is because FACET requires cardinality estimation for all columns involved in the predicates, followed by creating the refinements which require iterating over the dataset again. RAPIDASH requires no statistics and iterates over the dataset only once. Constraints  $c_6$ ,  $c_7$ , and  $c_8$  all have a large number of violations (on the order of several hundred million). For these constraints, both versions of RAPIDASH are up to 84 $\times$  faster than FACET. The performance gap is directly attributed to Proposition 1. RAPIDASH can find a violation after only looking at a few tuples in the dataset while FACET requires expensive computation. In fact, for  $c_6$  and  $c_7$ , we observed that the size of all ordered pairs<sup>5</sup> after just the first refinement (which are inequality predicates) is 1.2B and 3.6B respectively.

The speedup improvement obtained by RAPIDASH also extends to the violation enumeration problem (Figure 6b). Although there is no early termination possible for counting, RAPIDASH still performs up to an order of magnitude better due to our improved algorithms. FACET performance, on the other hand, degrades further since the last refinement cannot be stopped early as FACET requires *all* refinements to be complete in order to begin counting. Note that both RAPIDASH( $\perp$ ) and RAPIDASH(kd) have the same performance numbers since all constraints contain at most two inequality predicates and thus, both range trees and kd-trees degenerate into a simple binary search tree. For the NCVoter dataset (Figure 7c), we also observe the same behavior. For all DCs, both variants of RAPIDASH are 2 – 200 $\times$  faster.

Figure 7 shows the space usage for both verification and violation enumeration. For FACET, space usage is calculated using the cardinality (in millions) of cluster pairs constructed and we use the number of nodes in the tree constructed for RAPIDASH( $\perp$ ) and RAPIDASH(kd). For every DC, RAPIDASH uses significantly lower space compared to FACET. The high space usage of FACET is directly attributed to the size of ordered pairs generated after refining predicates. The largest gap is observed for  $c_6$ ,  $c_7$ , and  $c_8$ , which is expected since the DCs have a lot of violations, making refinement computation and storage expensive. On the other hand, for each constraint, RAPIDASH(kd) requires only (provably) linear amount of memory and this behavior can be directly observed in practice as well. For NCVoter, FACET memory use ranges from 256MB to 384MB (maximum memory was used by  $c_3$  since it has 25B violations) and RAPIDASH memory use was between 8MB to 71MB. We note that for constraint  $c_{10}$ , an example of a heterogeneous constraint, RAPIDASH uses the inequality heterogeneous predicate optimization from Section 4.6 and constructs a tree on only birthyear instead of both ageatyearend and birthyear, leading to a small memory footprint of only 8MB for both variants of RAPIDASH compared to 256MB for FACET.

<sup>5</sup>Given an ordered pair  $(tids_1, tids_2)$ , its size is defined as  $|tids_1| + |tids_2|$




Figure 6: Running time for evaluation on different datasets.




Figure 7: (Left and middle) Space requirement (cardinality of ordered pairs for FACET and number of nodes in the trees for RAPIDASH in millions) of different algorithms for DC verification on Tax and TPC-H; (Right) Running time (in milliseconds) for DC verification and violation enumeration on NCVoter.

**Constraints with near worst-case behavior.** To demonstrate that for FACET, even simple constraints on a small dataset can have poor performance in practice (as promised in Section 3), we tested the constraint  $\phi = \neg(s.\text{Tax} \neq t.\text{Tax} \wedge s.\text{ExtPrice} \neq t.\text{ExtPrice})$  on TPC-H dataset. For this constraint, FACET runs for a few minutes and the program then crashes due to Java out-of-memory error due to creation of trillions of ordered pairs when refining the second predicate of the constraint. The materialization eventually consumes the entire main memory. However, RAPIDASH can count the number of violations within 5 seconds.

Table 4: Running time (in milliseconds) for violation enumeration on the TPC-H dataset with varying cardinality.

DC	Dataset size	Enumeration		
		FACET	RAPIDASH( $\perp$ )	RAPIDASH(kd)
$c_6$	1M	5693	1075	1090
	2M	11713	1818	1835
	4M	24454	3739	3890
$c_7$	1M	6030	640	652
	2M	13653	1408	1423
	4M	34628	3053	3092
$c_8$	1M	5591	741	793
	2M	17136	1098	1124
	4M	34628	2772	2783


**Scalability.** Table 4 shows the behavior of the algorithms on the TPC-H dataset with varying cardinality. Let us take constraint  $c_8$

(rightmost subfigure) as an example. As the cardinality increases, the speedup obtained increases from 7.5 $\times$  for 1M to 13.86 $\times$  when the dataset size is 4M. This suggests that the running time of FACET grows at a non-constant rate compared to RAPIDASH which grows in line with expectation.

## 5.4 Evaluation on Production Datasets

In this section, we answer Q.3. Figure 6c shows the running time (in log scale) of RAPIDASH for all production datasets and DCs. The speedup obtained by our algorithm is close to an order of magnitude and up to 40 $\times$ . Compared to FACET, both algorithms perform significantly better on all DCs. RAPIDASH( $\perp$ ) performs better than RAPIDASH(kd) on all DCs. This is not surprising since using kd-trees for orthogonal range search is more expensive as shown by their big-O time analysis. However, RAPIDASH(kd) is still faster than FACET by up to 20 $\times$ .

The speedup obtained by RAPIDASH can be attributed to two reasons. First, RAPIDASH can terminate as soon as a violation is discovered, as opposed to FACET, which cannot do early termination in general (see Section 3). The second reason is that RAPIDASH does not require any expensive materialization as opposed to the ordered pair generation that is done by FACET. We also measured the space usage of all systems. For FACET, the space usage is the cardinality of the cluster pairs generated at each stage of the refinement pipelines. For RAPIDASH based algorithms, the space usage refers to the number of points inserted in the tree. For all DCs, FACET used 1.4 – 8 $\times$  more space compared to RAPIDASH( $\perp$ ). RAPIDASH(kd) was a further order of magnitude lower in its space requirement compared to RAPIDASH( $\perp$ ). Interested reader can find the figure in our



**Figure 8: Running time (in seconds) for violation enumeration on production datasets. Red (solid) bars for  $\varphi_{2,1}$  and  $\varphi_{2,2}$  denote an out-of-memory error.**

**Table 5: Running time (in milliseconds) of FACET and RAPIDASH (R( $\perp$ )) and R(kd) on DCs for Tax for varying number of violations.**

DC	# vio	% rows changed	Detection			Enumeration		
			FACET	R( $\perp$ )	R(kd)	FACET	R( $\perp$ )	R(kd)
$c_1$	138K	5%	347	1	1	812	294	294
	264K	10%	352	1	1	835	251	251
	482K	20%	373	1	1	904	262	262
	920K	50%	406	1	1	1042	295	295
$c_2$	8.4M	5%	469	2	2	1196	104	104
	15.7M	10%	566	1	1	1500	136	136
	27M	20%	504	1	1	1500	114	114
	43M	50%	492	2	2	1812	121	121
$c_3$	1.2B	5%	616	12	12	1014	160	160
	2.2B	10%	585	10	10	991	162	162
	3.6B	20%	615	4	4	1030	172	172
	5.4B	50%	589	4	4	1812	133	133
$c_4$	0.8B	5%	1781	16	26	9008	544	544
	1.5B	10%	1625	15	20	16656	630	630
	2.6B	20%	1792	9	29	22094	621	621
	4.1B	50%	1729	15	31	28280	705	705

technical report [29]. Figure 8 shows the running time for violation enumeration. For most constraints, we observe a similar trend in the running time as we saw for verification. The most interesting observation is that for constraints  $\varphi_{2,1}$  and  $\varphi_{2,2}$ , both FACET and RAPIDASH( $\perp$ ) fail to complete due to Java out-of-memory error, as both constraints contain a large number of inequality predicates. However, RAPIDASH(kd) can finish the computation in about 10 minutes, thanks to its linear memory use guarantee. In terms of space usage, we observed very similar behavior as Figure 7. We also observed good scalability behavior on the production datasets. We omit the graphs due to space limitations and refer the interested reader to the technical report [29].

### 5.5 Varying the Number of Violations

We conclude this section by answering Q.4. We analyze the performance for a fixed set of DCs but modify the dataset to vary the number of violations. We chose the Tax dataset for the experiment since all four DCs have zero violations on the unperturbed data. To introduce violations, we take the original table and modify a certain fraction of the rows by replacing the values in the row with another value from the domain of the column. Table 5 shows the number of

violations for each DC as we modify 5%, 10%, 20%, 50% of the dataset and running time. For detection, RAPIDASH is up to two orders of magnitude faster due to benefit of early termination. In the case of enumeration, all systems need to perform more computation. For FACET, we observe that as the number of violations increases, the running time also increases. This behavior is attributed to the cost of materializing cluster pairs which get large as the number of violations increase. The effect is starkly visible for  $c_4$  where the running time increase by over three times for 9s to 28s. The performance of both variants of RAPIDASH is more robust to the change in number of violations since tree based structures allow efficient counting. The space usage of RAPIDASH was also observed to be an order of magnitude smaller than FACET.

## 6 RELATED WORK

**DC Violation Detection and Discovery.** DCs as an integrity constraint language was originally proposed by Chu et al. [13]. We refer the reader to [2, 3] for a general overview. To the best of our knowledge, FACET [34] is the state-of-the-art algorithm for DC violation detection. FACET follows the design of VioFINDER [35] and uses an operator called *refinement* to evaluate DC predicates. Our proposed algorithm has better worst-case time/space complexity than FACET, as well as methods that rely on DBMS to detection violations [15, 18, 37], resulting in significant performance improvements in practice. DC verification and enumeration are closely related to solutions for DC discovery (such as HYDRA [11] and DCFINDER [33]). Both of these systems rely on the two-step process of first building the evidence set, followed by enumerating the DCs. The two-step approach is also been successfully used for other dependency discovery algorithms [32, 38].

**Range Searching.** The connection between geometric algorithms and general join query processing has been made by several prior works [26, 30]. Specifically, range searching has been used for aggregate query processing [25] and CQs involving comparisons [40, 41]. Optimizations introduced in this paper could also be applied to certain subsets of queries considered in [40, 41] since DCs can be expressed as CQs with "short" comparisons but with the added twist that the query is a self-join. A further consideration in our work (but not elaborated in the prior works) is that we have a significant interest in making sure that our algorithms can run in linear space, a requirement for production implementation. Range trees and their variants have also been extensively used in geospatial information systems (see [5–7, 12, 14, 20, 24] for an overview) and indexes for database systems [21]. For an overview of the theoretical aspects of range searching, we refer the reader to [4].

## 7 CONCLUSION

In this paper, we studied the problem of DC violation detection. We presented RAPIDASH, a DC violation detection algorithm with near-linear time complexity with respect to the dataset size that leverages prior work on orthogonal range search. Through empirical evaluation, we demonstrated that our algorithm is faster than the state of the art by up to 84 $\times$  on open-source datasets and large-scale production datasets.

## REFERENCES

- [1] Ziawasch Abedjan, Xu Chu, Dong Deng, Raul Castro Fernandez, Ihab F Ilyas, Mourad Ouzzani, Paolo Papotti, Michael Stonebraker, and Nan Tang. 2016. Detecting data errors: Where are we and what needs to be done? *Proceedings of the VLDB Endowment* 9, 12 (2016), 993–1004.
- [2] Ziawasch Abedjan, Lukasz Golab, and Felix Naumann. 2017. Data Profiling: A Tutorial. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May 14-19, 2017*, Semih Salihoglu, Wenchao Zhou, Rada Chirkova, Jun Yang, and Dan Suciu (Eds.). ACM, 1747–1751. <https://doi.org/10.1145/3035918.3054772>
- [3] Ziawasch Abedjan, Lukasz Golab, Felix Naumann, and Thorsten Papenbrock. 2018. Data Profiling. Morgan & Claypool Publishers. *Synthesis Lectures on Data Management* (2018).
- [4] Pankaj K. Agarwal. 2004. Range Searching. In *Handbook of Discrete and Computational Geometry, Second Edition*, Jacob E. Goodman and Joseph O'Rourke (Eds.). Chapman and Hall/CRC, 809–837. <https://doi.org/10.1201/9781420035315.ch36>
- [5] Lars Arge, Mark de Berg, Herman Haverkort, and Ke Yi. 2008. The priority R-tree: A practically efficient and worst-case optimal R-tree. *ACM Transactions on Algorithms (TALG)* 4, 1 (2008), 1–30.
- [6] Norbert Beckmann, Hans-Peter Kriegel, Ralf Schneider, and Bernhard Seeger. 1990. The R<sup>+</sup>-tree: An efficient and robust access method for points and rectangles. In *Proceedings of the 1990 ACM SIGMOD international conference on Management of data*. 322–331.
- [7] Norbert Beckmann and Bernhard Seeger. 2009. A revised R<sup>+</sup>-tree in comparison with related index structures. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*. 799–812.
- [8] Jon Louis Bentley and Jerome H Friedman. 1979. Data structures for range searching. *ACM Computing Surveys (CSUR)* 11, 4 (1979), 397–409.
- [9] Jon Louis Bentley and James B Saxe. 1980. Decomposable searching problems I. Static-to-dynamic transformation. *Journal of Algorithms* 1, 4 (1980), 301–358.
- [10] Pouya Bisadi and Bradford G Nickerson. 2011. Orthogonal Range Search using a Distributed Computing Model.. In *CCCG*.
- [11] Tobias Bleifuß, Sebastian Kruse, and Felix Naumann. 2017. Efficient denial constraint discovery with hydra. *Proceedings of the VLDB Endowment* 11, 3 (2017), 311–323.
- [12] King Lum Cheung and Ada Wai-Chee Fu. 1998. Enhanced nearest neighbour search on the R-tree. *ACM SIGMOD Record* 27, 3 (1998), 16–21.
- [13] Xu Chu, Ihab F Ilyas, and Paolo Papotti. 2013. Discovering denial constraints. *Proceedings of the VLDB Endowment* 6, 13 (2013), 1498–1509.
- [14] Mark de Berg, Otfried Cheong, Marc J. van Kreveld, and Mark H. Overmars. 2008. *Computational geometry: algorithms and applications, 3rd Edition*. Springer. <https://www.worldcat.org/oclc/227584184>
- [15] Wenfei Fan, Chao Tian, Yanghao Wang, and Qiang Yin. 2021. Parallel discrepancy detection and incremental detection. *Proceedings of the VLDB Endowment* 14, 8 (2021), 1351–1364.
- [16] Anna Fariha, Ashish Tiwari, Arjun Radhakrishna, Sumit Gulwani, and Alexandra Meliou. 2021. Conformance Constraint Discovery: Measuring Trust in Data-Driven Systems. In *SIGMOD '21: International Conference on Management of Data, Virtual Event, China, June 20-25, 2021*. ACM, 499–512. <https://doi.org/10.1145/3448016.3452795>
- [17] Chang Ge, Shubhankar Mohapatra, Xi He, and Ihab F. Ilyas. 2021. Kamino: Constraint-Aware Differentially Private Data Synthesis. *Proc. VLDB Endow.* 14, 10 (2021), 1886–1899. <https://doi.org/10.14778/3467861.3467876>
- [18] Floris Geerts, Giansalvatore Mecca, Paolo Papotti, and Donatello Santoro. 2020. Cleaning data with Llunatic. *The VLDB Journal* 29 (2020), 867–892.
- [19] Stella Giannakopoulou, Manos Karpathiotakis, and Anastasia Ailamaki. 2020. Cleaning denial constraint violations through relaxation. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 805–815.
- [20] Antonin Guttman. 1984. R-Trees: A Dynamic Index Structure for Spatial Searching. In *SIGMOD '84, Proceedings of Annual Meeting, Boston, Massachusetts, USA, June 18-21, 1984*. ACM Press, 47–57. <https://doi.org/10.1145/602259.602266>
- [21] Joseph M. Hellerstein, Jeffrey F. Naughton, and Avi Pfeffer. 1995. *Generalized Search Trees for Database Systems*. Morgan Kaufmann. 562–573 pages. <http://www.vldb.org/conf/1995/P562.PDF>
- [22] John E. Hopcroft and Jeffrey D. Ullman. 1979. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley.
- [23] Toshihide Ibaraki, Alexander Kogan, and Kazuhisa Makino. 1999. Functional dependencies in Horn theories. *Artificial Intelligence* 108, 1-2 (1999), 1–30.
- [24] Kothuri Venkata Ravi Kanth, Siva Ravada, and Daniel Abugov. 2002. Quadtree and R-tree indexes in oracle spatial: a comparison using GIS data. In *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data, Madison, Wisconsin, USA, June 3-6, 2002*. ACM, 546–557. <https://doi.org/10.1145/564691.564755>
- [25] Mahmoud Abo Khamis, Ryan R Curtin, Benjamin Moseley, Hung Q Ngo, Xuan-Long Nguyen, Dan Olteanu, and Maximilian Schleich. 2020. Functional aggregate queries with additive inequalities. *ACM Transactions on Database Systems (TODS)* 45, 4 (2020), 1–41.
- [26] Mahmoud Abo Khamis, Hung Q Ngo, Christopher Ré, and Atri Rudra. 2016. Joins via geometric resolutions: Worst case and beyond. *ACM Transactions on Database Systems (TODS)* 41, 4 (2016), 1–45.
- [27] Zuhair Khayyat, William Lucia, Meghna Singh, Mourad Ouzzani, Paolo Papotti, Jorge-Arnulfo Quiáné-Ruiz, Nan Tang, and Panos Kalnis. 2015. Lightning Fast and Space Efficient Inequality Joins. *Proc. VLDB Endow.* 8, 13 (2015), 2074–2085. <https://doi.org/10.14778/2831360.2831362>
- [28] Jan Kossmann, Thorsten Papenbrock, and Felix Naumann. 2022. Data dependencies for query optimization: a survey. *The VLDB Journal* 31, 1 (2022), 1–22.
- [29] Zifan Liu, Shaleen Deep, Anna Fariha, Fotis Psallidas, Ashish Tiwari, and Avriella Floratou. 2023. Rapidash: Efficient Constraint Discovery via Rapid Verification. *arXiv preprint arXiv:2309.12436* (2023).
- [30] Hung Q Ngo, Dung T Nguyen, Christopher Re, and Atri Rudra. 2014. Beyond worst-case analysis for joins with minesweeper. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*. 234–245.
- [31] Mark H Overmars. 1983. *The design of dynamic data structures*. Vol. 156. Springer Science & Business Media.
- [32] Thorsten Papenbrock and Felix Naumann. 2016. A hybrid approach to functional dependency discovery. In *Proceedings of the 2016 International Conference on Management of Data*. 821–833.
- [33] Eduardo HM Pena, Eduardo C de Almeida, and Felix Naumann. 2019. Discovery of approximate (and exact) denial constraints. *Proceedings of the VLDB Endowment* 13, 3 (2019), 266–278.
- [34] Eduardo HM Pena, Eduardo C de Almeida, and Felix Naumann. 2021. Fast detection of denial constraint violations. *Proceedings of the VLDB Endowment* 15, 4 (2021), 859–871.
- [35] Eduardo HM Pena, Edson R Lucas Filho, Eduardo C de Almeida, and Felix Naumann. 2020. Efficient detection of data dependency violations. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 1235–1244.
- [36] Eduardo HM Pena, Fabio Porto, and Felix Naumann. 2022. Fast Algorithms for Denial Constraint Discovery. *Proceedings of the VLDB Endowment* 16, 4 (2022), 684–696.
- [37] Theodoros Rekatsinas, Xu Chu, Ihab F. Ilyas, and Christopher Ré. 2017. HoloClean: Holistic Data Repairs with Probabilistic Inference. *Proc. VLDB Endow.* 10, 11 (2017), 1190–1201. <https://doi.org/10.14778/3137628.3137631>
- [38] Philipp Schirmer, Thorsten Papenbrock, Ioannis Koumarelas, and Felix Naumann. 2020. Efficient discovery of matching dependencies. *ACM Transactions on Database Systems (TODS)* 45, 3 (2020), 1–33.
- [39] Jaroslaw Szlichta, Parke Godfrey, and Jarek Gryz. 2012. Fundamentals of Order Dependencies. *Proc. VLDB Endow.* 5, 11 (2012), 1220–1231. <https://doi.org/10.14778/2350229.2350241>
- [40] Qichen Wang and Ke Yi. 2022. Conjunctive Queries with Comparisons. In *Proceedings of the 2022 International Conference on Management of Data*. 108–121.
- [41] Dan E Willard. 1996. Applications of range query theory to relational data base join and selection operations. *journal of computer and system sciences* 52, 1 (1996), 157–169.