

Differentially Private Release of Hierarchical Origin/Destination Data with a TopDown Approach

Fabrizio Boninsegna

University of Padova

Padova, Veneto, Italy

fabrizio.boninsegna@phd.unipd.it

Francesco Silvestri

University of Padova

Padova, Veneto, Italy

francesco.silvestri@unipd.it

Abstract

This paper presents a novel method for generating differentially private tabular datasets for hierarchical data, specifically focusing on origin-destination (O/D) trips. The approach builds upon the TopDown algorithm, a constraint-based mechanism developed by the US Census to incorporate invariant queries into tabular data. O/D hierarchical data refers to datasets representing trips between geographical areas organized in a hierarchical structure (e.g., region \rightarrow province \rightarrow city). The proposed method is designed to improve the accuracy of queries covering broader geographical areas, which are derived through aggregation. This feature provides a “zoom-in” effect on the dataset, ensuring that when zoomed back out, the overall picture is preserved. Furthermore, the approach aims to reduce false positive detection. These characteristics can strengthen practitioners’ and decision-makers’ confidence in adopting differential privacy datasets. The main technical contribution of this paper includes a novel TopDown algorithm that employs constrained optimization with Chebyshev distance minimization, with theoretical guarantees on the maximum absolute error. Additionally, we propose a new integer optimization algorithm that significantly reduces the incidence of false positives. The effectiveness of the proposed approach is validated using real-world and synthetic O/D datasets, demonstrating its ability to generate private data with high utility and a reduced number of false positives. Our experiments focus on O/D datasets with a single trip as a unit of privacy; nevertheless, the proposed approach supports other units of privacy and also applies to any tabular data with a hierarchical structure.

Keywords

Differential Privacy, Mobility, Constrained Optimization

1 Introduction

Origin-destination (O/D) data plays a crucial role in policy planning, especially in today’s context, where official statistical agencies release detailed trip records. These records include origin and destination areas, along with various trip attributes such as mode of travel and trip purpose. These detailed data are important for a wide range of planning purposes, from transportation planning [23] to epidemic modeling [16], and are essential for understanding and managing the flow of people and goods in various settings.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

Proceedings on Privacy Enhancing Technologies 2025(3), 29–43

© 2025 Copyright held by the owner/author(s).

<https://doi.org/10.56553/popets-2025-0087>

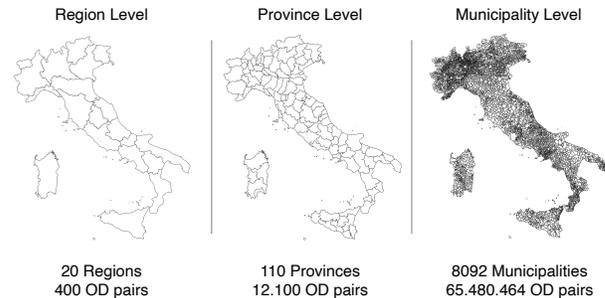


Figure 1: The first three hierarchical levels of Italy, according to ISTAT.

However, the release of mobility data poses significant privacy risks. Individuals’ movement patterns can be sensitive information, potentially revealing personal habits and frequented locations. Differential Privacy (DP) [13] provides a rigorous solution to this challenge. It involves introducing randomness to data in a controlled manner, ensuring that the privacy of individuals in a dataset is preserved while still allowing meaningful analysis. In this paper, we focus on differentially private release of O/D trips with a geographic hierarchy, focusing on obtaining a “zoom-in” effect (which we call *hierarchical accuracy*) and reducing false positives. We will focus on the case where each user contributes one trip and specifically on *bounded* differential privacy [21], where the total number of users is an invariant and does not change under DP.

The Geographic Hierarchy. A geographic hierarchy with g levels can be represented as a tree that captures the relationships between these geographic levels. Each node in the tree corresponds to a geographic area, while its children represent subareas into which that area is divided. For example, the Italian Institute of Statistics (ISTAT)¹ organizes Italy into a geographic hierarchy consisting of regions, provinces, municipalities, section areas, and census sections (with the first three levels illustrated in Figure 1). The regions represent the largest areas (level 1), and each unit at level j is entirely contained within exactly one unit at level $j - 1$. An O/D dataset of trips $D = \{(u_i, v_i)\}_{i=1, \dots, n}$, where u_i, v_i denote the origin and destination at level g of the user i , enables practitioners to compute marginal geographic queries. For example, given an O/D dataset that contains trips between municipalities, it is possible to derive flows (i.e., the number of O/D trips in the dataset) between pairs of provinces or regions.

¹<https://www.istat.it/it/archivio/222527>

Reducing False Positives. In the O/D dataset scenario, the basic mechanism to ensure DP is the addition of Laplace or Gaussian noise to *all* the O/D flows, regardless of the existence of a trip in the data. However, since these datasets are typically sparse (e.g., for trips between cities within a country, only a small subset of possible trips is usually reported), this approach can result in a pathological behavior of false positives, a phenomenon well-documented in the literature [4, 9]. False positives in data analysis refer to instances where there are O/D trips in the DP data, even though they were not present in the real data. For sparse datasets, minimizing false positives is essential to preserve their inherent sparsity, ensuring that the released dataset remains compact in size. Furthermore, reducing false positives can improve confidence in adopting such solutions: a dataset cluttered with numerous insignificant trips (with small flows) is likely to be perceived as low quality by decision-makers, even if the flow of false positive trips might not be statistically significant.

The Hierarchical Accuracy. A geographic hierarchy enables practitioners to aggregate differentially private noisy data flows to obtain insights about broader geographic regions. However, this approach leads to a decrease in precision due to noise accumulation, which significantly reduces the accuracy of the dataset at larger scales. This outcome contrasts with the intuitive expectation that, for general purpose private data, statistics become less accurate as one “zooms-in” to smaller regions. One potential solution is to generate multiple DP datasets, each tailored to a specific geographic level. However, this approach has two major drawbacks: it requires practitioners to manage multiple datasets, and the results derived from these datasets may lack consistency. Therefore, there is a clear need for a single tabular dataset that provides hierarchical accuracy, offering higher precision for broader geographic regions while maintaining consistency and usability.

Our goal is to derive a mechanism to release a unique O/D tabular dataset \tilde{D} at level g with

- (1) (*Privacy*) The dataset \tilde{D} is (ϵ, δ) -differentially private.
- (2) (*Hierarchical accuracy*) The dataset \tilde{D} has a “zoom-in” effect on accuracy. Hence, it has higher accuracy for queries that cover larger geographic areas compared to those covering smaller areas².
- (3) (*Reduce False Positives*) The differentially private algorithm is designed to reduce the occurrence of false positives.

Our algorithm InfTDA rigorously satisfies the first two points and has been experimentally shown to reduce false positives in real and synthetic data sets. InfTDA is built upon the TopDown Algorithm (TDA) developed by the US Census [1] and used to publish the 2020 US census. The main idea of TDA is to use an iterative optimization approach along the hierarchy to ensure hierarchical accuracy in practice. InfTDA distinguishes itself from TDA through a modified optimization approach and a specific focus on addressing false positives, which are effectively minimized using a heuristic integrated into the optimization process.

²In the sense that the estimates related to area A are more accurate than those related to any area nested within A.

1.1 Our Results

Our algorithm InfTDA builds upon TDA with a key modification: it uses the Chebyshev distance (which is the ℓ_∞ distance, so the name of our algorithm) as the objective function in the constrained optimization problem, instead of the Euclidean distance in TDA. Intuitively, minimizing Chebyshev distance is simpler than Euclidean, as the former can be minimized with a linear program. More importantly, to the best of our knowledge, with this choice we are the first to provide an upper bound for the utility of a TDA-like mechanism³. Consequently, we propose IntOpt as an optimizer for the Chebyshev distance that works entirely in the integer domain and is designed with a heuristic to reduce false positives. Finally, we show that any TDA-like algorithm works for a specific tree data structure having some well-defined properties, which we call *non-negative hierarchical tree*. With this generalization, we ensure the broad applicability of InfTDA to any tabular data that can be represented as a non-negative hierarchical tree, making it highly generalizable⁴.

The main results of this paper are

- (1) A demonstration that any O/D dataset can be parsed into a non-negative hierarchical tree, and so it is suitable for a TopDown algorithm.
- (2) A theoretical analysis of the accuracy of InfTDA. The chosen accuracy is the *maximum absolute error* at each level of the hierarchy.
- (3) We propose IntOpt, a fast integer-constrained optimization algorithm that minimizes Chebyshev distance and includes a heuristic to reduce false positives in practical scenarios.

We evaluated InfTDA using both real-world data (a dataset of O/D commuting flows in Italy) and synthetic O/D data. The results show that its utility is not worse than TDA while being faster, simpler, and generates a dataset with fewer false positives. In contrast to TDA, our InfTDA has a theoretical upper bound for the utility, which we now informally show for O/D datasets.

THEOREM 1 (INFORMAL VERSION OF UTILITY OF InfTDA). *Given a O/D dataset with g geographic levels. InfOPT with constant probability returns a differentially private tabular dataset, with maximum absolute error at most $\tilde{O}(\sqrt{\ell^3 g})$, for O/D flows with origin and destination at level $\ell \in \{0, \dots, g\}$.*

A more formal version for any non-negative hierarchical tree is stated in Theorem 11.

1.2 Previous Work

Histogram in DP. Releasing a differentially private O/D dataset is essentially equivalent to the release of differentially private histograms, extensively explored in existing literature [25, 30, 31]. In fact, $D = \{(u_i, v_i)\}_{i=1, \dots, n}$ has a classical histogram representation where each bin counts the occurrence of an O/D pair in D . This can be seen as a tabular representation with origin, destination, and flow (i.e., counts) as columns. The main strategy involves adding Laplace noise to achieve pure differential privacy, as detailed by

³We intend any mechanism that uses an iterative optimization approach along a hierarchy, regardless of the optimization or the differentially private mechanism used.

⁴For example, in a healthcare dataset encompassing diseases and user characteristics, a possible hierarchy might start with counts of diseases, segmented further by gender, then by age groups, and so on.

Dwork et al. [13], or incorporating Gaussian noise for approximate differential privacy [5, 8]. The latter has found practical application in the US Census [1] as Gaussian tail bounds provide a better trade-off between privacy and outliers. These methods return an unbiased estimator of the histogram, but are space-inefficient when the histogram is sparse, as it is necessary to generate independent noise for every count. To address this challenge, more sophisticated algorithms have been developed specifically for sparse histograms. The stability-based method [22, 26, 28] relies on thresholding private counts, getting a biased estimator without false positives. For the sensitivity one scenario, Desfontaines et al. [11] have developed a mechanism that is optimal to reduce the number of false negatives. A method to obtain an unbiased estimator of the histogram has been developed by Cormode et al. [9] using priority sampling at the price of increasing the expected error. Aumüller et al. [4] developed a biased estimator for the counts that matches a lower bound up to a constant factor. All the aforementioned approaches return a differentially private histogram representation of D , mostly by matching lower bounds for counts [28]. However, when range queries are important (i.e., sum of counts), these mechanisms do not offer strong upper bounds due to noise propagation.

Constrained Optimization. The pioneering study by Hay et al. [19] was the first to observe that the accuracy of range queries in differential privacy could be improved through a post-processing stage involving constrained optimization. This was achieved using the Hierarchical mechanism, which is based on the construction of a tree of range queries. In this approach, the outputs of differentially private queries are post-processed such that the result of each query node is derived from an aggregation of the results from its child nodes. The authors also provided an analytical solution to the optimization problem, however it only applies to queries in the real domain, which makes it not appropriate to release counts.

Fiorretto et al. [15] further advanced the field, particularly for mobility data, by developing the Constrained Based Differential Private (CBDP) mechanism. This mechanism specifically addresses the challenge of releasing mobility data with differential privacy, particularly for On-Demand multi-modal transit systems. Building upon the Hierarchical mechanism, CBDP enhances it by integrating the *non-negativity* constraint (i.e., only positive flows are accepted; negative flows may occur due to injection of randomness) and generalizes it by allowing more general constraints than only hierarchical ones. However, this involves solving a unique constrained optimization problem with a potentially prohibitively large number of variables and constraints, especially for O/D data.

The most significant utilization of constrained optimization has been observed in the Disclosure Avoidance System TopDown algorithm (which we called it TDA) implemented by the US Census Bureau to publish the 2020 USA census data [1]. The goal of TDA was to release population histograms, including ethnicity and age distributions, for each census section while maintaining a fixed number of people per state. To achieve this, the US Census Bureau devised a TopDown algorithm that follows a geographic hierarchy that they called *geographic spine* (Nations \rightarrow Regions \rightarrow Divisions \rightarrow States \rightarrow Counties etc.) [27]. Starting with the creation of a differentially private tabulated dataset at the national level, the TopDown approach then proceeds to process data for regions. To

ensure coherence, a constraint optimization problem is solved during post-processing, aiming for the regional tabulated data to be consistent with that of the nation. Essentially, the aggregated attributes of the regions must align with the national attributes. The process is repeated iteratively up to the final level of the geographic spine, ending in the release of micro-data at the census section level. The advantage of a TopDown approach is twofold: it splits the optimization problem into many more feasible problems, and it mimics the sparsity of the data. The latter advantage is an inherent characteristic of the TopDown approach, as identifying false positives is more straightforward in large aggregated datasets. Once an attribute is determined to not exist in a geographic area, it is inferred that the same attribute will also be absent in any subdivision of that area. This inference avoids the introduction of noise in these areas throughout the TopDown process.

Objective Function. All the previously discussed algorithms aim to solve a constrained optimization problem by returning a vector that is as close as possible to the differentially private estimate, measured in terms of the Euclidean distance, referred to as the ℓ_2 distance. Hay et al. [19] raised the possibility of minimizing the ℓ_1 distance (the sum of absolute errors between the differentially private estimate and the released vector). Since this approach does not guarantee a unique solution, the authors chose the ℓ_2 minimization, as it ensures uniqueness. TDA uses a *weighted non-negative* least squares optimization, meaning that a weight could be assigned to each absolute error. In particular, the US Census opted to use the inverse variance of the differentially private random variables as weights. A similar weighted approach was used by Fiorretto et al. [15].

We chose to minimize the Chebyshev distance for two purposes: it produces a simple integer constraint minimization problem, and it gives theoretical guarantees. In contrast, TDA uses a complex two-step optimization algorithm, first solving the problem in the real space with convex optimization and then performing the best integer rounding with linear programming.

1.3 Structure of the Paper

In Section 2, we introduce the core principles of differential privacy, highlighting two key mechanisms: the Gaussian mechanism and the Stability Histogram. Additionally, we define the non-negative hierarchical tree. In Section 3, we define the notation used for the hierarchical structure of O/D data, its reformulation as a non-negative hierarchical tree, and the utility metric used in our analysis. The tree reformulation is then used to formulate our algorithm InfTDA in Section 4, making it of broad applicability. This is followed by a theoretical analysis of InfTDA and the introduction of IntOpt, an optimizer tailored to reduce false positives. The discussion concludes in Section 5, which presents an experimental evaluation of InfTDA compared to baselines and other TDA-like mechanisms on real-world and synthetic datasets, highlighting the advantages of the proposed methodology.

2 Preliminaries

2.1 Differential Privacy

This section introduces the key concepts of differential privacy relevant to this article. In our discussion, we focus on the notion of privacy known as *bounded* [21]. In this framework, two datasets, D and D' , are neighbors (that is, $D \sim D'$) if one can be obtained from the other by substituting a single user. Thus, the name bounded refers to the fact that neighboring datasets have a fixed size. In contrast, within the framework of *unbounded* differential privacy [21], two neighboring datasets may differ by either the addition or removal of a single user, allowing their sizes to vary. In the following, we present the formal definition of differential privacy.

DEFINITION 1 (DIFFERENTIAL PRIVACY (DP) [14]). *Given $\epsilon > 0$ and $\delta \in [0, 1)$. A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ satisfies (ϵ, δ) -differential privacy if for any two neighboring datasets $D, D' \in \mathcal{D}$ and for any subset of outputs $S \subseteq \mathcal{R}$ it holds that*

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta.$$

Another definition, which is more suitable to study the injection of Gaussian noise (used in practical applications, such as that used in the US Census release [1]) is *zero-Concentrated Differential Privacy*. This definition provides a tighter analysis of the aforementioned mechanism and yields simpler composition results.

DEFINITION 2 (ZERO-CONCENTRATED DIFFERENTIAL PRIVACY (zCDP)[7]). *Given $\rho > 0$. A randomized mechanism $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ satisfies ρ -zCDP if for any two neighboring datasets $D, D' \in \mathcal{D}$ and any $\alpha > 1$*

$$D_\alpha(\mathcal{M}(D) \parallel \mathcal{M}(D')) \leq \alpha\rho,$$

where $D_\alpha(\mathcal{M}(D) \parallel \mathcal{M}(D'))$ is the α -Rényi divergence.

Any ρ -zCDP mechanism satisfies also (ϵ, δ) -DP.

LEMMA 2 (FROM ρ -zCDP TO (ϵ, δ) -DP (LEMMA 21 IN [7])). *Let $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ satisfy ρ -zCDP. Then \mathcal{M} satisfies (ϵ, δ) -DP for all $\delta \in (0, 1)$ and*

$$\epsilon = \rho + 2\sqrt{\rho \log(1/\delta)}.$$

This lemma is useful for converting ρ -zCDP guarantees into (ϵ, δ) -DP, which are more commonly used in real-world applications. A significant benefit of employing differential privacy is its resilience to privacy degradation, regardless of the application of any post-processing functions.

LEMMA 3 (POST-PROCESS IMMUNITY (LEMMA 8 [7])). *Let $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ and $f : \mathcal{R} \rightarrow \mathcal{R}'$ be an arbitrary (also randomized) mapping. Suppose \mathcal{M} satisfies ρ -zCDP. Then $f \circ \mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}'$ satisfies ρ -zCDP.*

This characteristic is crucial when data are released to practitioners. If a dataset \tilde{D} is generated using a differentially private algorithm, then running any query on this dataset will not compromise its privacy. Another important property of differential privacy is the composition of privacy budgets, which enables the computation of differential privacy guarantees for the composition of several private algorithms.

LEMMA 4 (COMPOSITION (FROM LEMMA 7 IN [7])). *Let $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$ and $\mathcal{M}' : \mathcal{D} \rightarrow \mathcal{R}'$ be randomized algorithms. Suppose \mathcal{M} satisfies ρ -zCDP and \mathcal{M}' satisfies ρ' -zCDP. Define $\mathcal{M}'' : \mathcal{D} \rightarrow \mathcal{R} \times \mathcal{R}'$ by $\mathcal{M}''(x) = (\mathcal{M}(x), \mathcal{M}'(x))$. Then \mathcal{M}'' satisfies $(\rho + \rho')$ -zCDP.*

Since O/D data can be visualized as a histogram covering all possible O/D trips, we shift our focus to the release of histograms under differential privacy. The histogram for a dataset \mathcal{X}^n , where \mathcal{X} denotes the data universe (e.g., all possible O/D pairs) and n is the number of users, is generated by a counting query $q : \mathcal{X}^n \rightarrow \mathbb{N}^{|\mathcal{X}|}$. This query outputs the absolute frequency of each row in the dataset. A common paradigm for approximating these functions with differentially private mechanisms is through *additive noise mechanisms* calibrated to the p -global sensitivity $\text{GS}_p(q)$ of the function, which is defined as the maximum absolute ℓ_p distance

$$\text{GS}_p(q) = \sup_{D \sim D'} \|q(D) - q(D')\|_p,$$

where the sup is taken over two neighboring datasets. If a user contributes at most to m different trips in an O/D dataset then, under bounded privacy, we have $\text{GS}_1(q) = 2m$ and $\text{GS}_2(q) = \sqrt{2m}$.⁵ The first calibrates the additive noise from a Laplace distribution, while the second calibrates the noise from a Gaussian distribution. We now illustrate two mechanisms for the release of (ϵ, δ) -DP histograms.

2.1.1 Discrete Gaussian Mechanism. If the data universe \mathcal{X} is finite and known, we can achieve zCDP (so approximate DP due to Lemma 2) by adding Gaussian noise to each count.

THEOREM 5 (DISCRETE GAUSSIAN MECHANISM [8]). *Let $q : \mathcal{X}^n \rightarrow \mathbb{N}^{|\mathcal{X}|}$ be a counting query. The discrete Gaussian mechanism applied to a counting query $q(D)$ consisting of the injection of discrete Gaussian noise*

$$\tilde{q}(D) = q(D) + Z \quad Z \sim \mathcal{N}_{\mathbb{Z}}\left(0, \frac{\text{GS}_2(q)^2}{2\rho}\right)^{|\mathcal{X}|}$$

is ρ -zCDP.

The precision of the mechanism is not worse than that of its continuous counterpart [7]

COROLLARY 6 (COROLLARY 9 [8]). *Let $Z \sim \mathcal{N}_{\mathbb{Z}}(0, \sigma^2)$. Then $\text{Var}[Z] \leq \sigma^2$ and $\Pr[Z \geq t] \leq e^{-\frac{t^2}{2\sigma^2}}$ for any $t \geq 0$.*

2.1.2 Stability-Based Histogram. If the data universe \mathcal{X} is unknown, infinite, or very large, we can apply Laplace noise to each positive count of the histogram as long as noisy counts smaller than a certain threshold are set to zero.

THEOREM 7 (SH-STABILITY-BASED HISTOGRAM [6]). *Let $q : \mathcal{X}^n \rightarrow \mathbb{N}^{|\mathcal{X}|}$ be a counting query of 1-global sensitivity equal to 2. The algorithm that first applies Laplace noise to positive queries*

$$\tilde{q}(x_i) = q(x_i) + Z \quad \forall x_i \in \mathcal{X} : q(x_i) > 0, \text{ and } Z \sim e^{-\frac{2}{\epsilon}|x|}$$

and then maps to zero the noisy counts under $t = 1 + \frac{2\log(2/\delta)}{\epsilon}$, is (ϵ, δ) -differentially private.

It is important to stress that this mechanism *does not return false positives by construction* since it only injects noise into positive counts. However, it returns a biased estimator due to thresholding, which might be useless for estimating aggregate queries. For example, if all counts fall below the threshold, any aggregate query would result in zero.

⁵For unbounded privacy, we instead have $\text{GS}_1(q) = m$ and $\text{GS}_2(q) = \sqrt{m}$. The factor of 2 in bounded privacy arises from the definition of neighboring datasets, where D' is considered a neighbor of D if it differs from D by both the addition and removal of two distinct users (i.e., a substitution).

2.2 Non-Negative Hierarchical Tree

The InfTDA algorithm and the analysis provided in this paper are formulated for a tree data structure satisfying a few properties, ensuring broad applicability beyond O/D data. A tree \mathcal{T} of depth T is a tuple $\mathcal{T} = (V, E)$: $V = \cup_{\ell=0}^T V_\ell$ is the set of nodes, with V_ℓ denoting the set of nodes at level $\ell \in [T]$, for $[T] = [0, \dots, T]$; E is the set of edges between consecutive levels. A node at level ℓ is indicated with $u_\ell \in V_\ell$. The set of children to a node is indicated by a function $C : V_\ell \rightarrow 2^{V_{\ell+1}}$ for any level $\ell \in [T]$. Each node of the tree has an attribute $q(u_\ell)$. The specific tree we are interested in this paper is defined as follows.

DEFINITION 3 (NON-NEGATIVE HIERARCHICAL TREE). *A tree $\mathcal{T} = (V, E)$ is said to be non-negative if it contains non-negative attributes $q(u) \geq 0$ for any $u \in V$. The tree is hierarchical if the attribute of u can be computed as the sum of the attributes of its children $C(u)$. Specifically, for every $u \in V$ that is not a leaf, we have:*

$$q(u) = \sum_{v \in C(u)} q(v). \quad (1)$$

In our algorithm, we will use the function $q_C(u_\ell)$, which is the vector containing the attributes of the children of u_ℓ . For the theoretical analysis of the algorithm, we consider a tree with fixed branching factor $b \in \mathbb{N}$, so that $|V_\ell| = b^\ell$. Given any randomized mechanism \mathcal{M} applied to the attributes of the tree, the utility metric is defined for each level $\ell \in [T]$ as the *maximum absolute error*

$$\max_{u_\ell \in V_\ell} |\text{err}(u_\ell)| = \max_{u_\ell \in V_\ell} |\mathcal{M}(q(u_\ell)) - q(u_\ell)|. \quad (2)$$

3 Tree Structure of O/D Data

In this section, we show how any O/D dataset can be parsed into two different non-negative hierarchical trees, which we call the origin and destination trees. This reformulation is useful for describing some queries in the dataset. We start by defining the hierarchy in the geographic space.

Space Partitioning. Let X be a geographic area (e.g., X represents Italy), and assume that X is hierarchically partitioned into g levels (P_1, \dots, P_g) . The dependency among levels is described by the relation $h_\ell : P_\ell \rightarrow P_{\ell-1}$, which are injections mapping areas at level ℓ to the larger areas at level $\ell - 1$, for any $\ell \in [1, \dots, g]$. Note that, according to the previous definition, an area at level ℓ' is included in only one area at level $\ell < \ell'$. In our example in Figure 1, X is all of Italy, while P_1 is the set of regions, P_2 is the set of provinces and P_3 is the set of municipalities. With a slight abuse of notation, we write $v_\ell \in v_{\ell-1}$ to indicate that area v_ℓ is embodied in area $v_{\ell-1}$, so $v_\ell \in h_\ell^{-1}(v_{\ell-1})$. This holds for any geographical inclusion, so $v_{\ell'} \in v_\ell$ if $v_{\ell'}$ is completely contained in the larger region v_ℓ for $\ell < \ell'$.

The O/D Dataset and Range Queries. The dataset we study is a collection of O/D pairs (also called trips) in the finest partition P_g . It is represented as $D = \{(u_{g,i}, v_{g,i})\}_{i=1, \dots, n}$, for $u_g, v_g \in P_g$. For example, considering the hierarchy in Figure 1 that stops at the municipality level (e.g., $g = 3$), the dataset would contain trips between municipalities. These O/D pairs can be counted and aggregated to answer marginal queries, which we refer to as *hierarchical range queries*.

DEFINITION 4 (HIERARCHICAL RANGE QUERY). *Given two levels $\ell_1, \ell_2 \in [g]$ and two areas $u_{\ell_1} \in P_{\ell_1}$ and $v_{\ell_2} \in P_{\ell_2}$, the hierarchical range query is*

$$q(u_{\ell_1}, v_{\ell_2}) = \sum_{u_g \in u_{\ell_1}} \sum_{v_g \in v_{\ell_2}} \sum_{x \in D} \mathbb{1}\{x = (u_g, v_g)\}.$$

Here, $\mathbb{1}\{x = (u_g, v_g)\}$ is an indicator function that is used to count how many O/D pairs (that is, the flow) are in the dataset. Meanwhile, P_0 represents the entire geographic space X . These queries are conceptually equivalent to SQL GROUP BY followed by SUM, so we will refer to them simply as range queries. In particular, we are interested in *intra-level* range queries when $\ell_1 = \ell_2$, and *cross-level* range queries of order one, when $|\ell_1 - \ell_2| = 1$, as they allow us to construct the origin or destination tree, thanks to a *hierarchical consistency*.

OBSERVATION 1 (HIERARCHICAL CONSISTENCY). *Given two levels $\ell_1, \ell_2 \in [g]$ and two areas $u_{\ell_1} \in P_{\ell_1}$ and $v_{\ell_2} \in P_{\ell_2}$, then for any $\ell'_1 \leq \ell_1 \leq g$ and $\ell'_2 \leq \ell_2 \leq g$ we have*

$$q(u_{\ell_1}, v_{\ell_2}) = \sum_{u_{\ell'_1} \in u_{\ell_1}} \sum_{v_{\ell'_2} \in v_{\ell_2}} q(u_{\ell'_1}, v_{\ell'_2}).$$

Following the Italy example, the number of trips from the region Veneto to the region Lombardia has to be the sum of the number of trips among their cities. Another example is that the number of Italians going to Milan must be the sum of the number of trips starting in any Italian region and ending in Milan. Notice that origins and destinations may belong to two completely different geographic spaces X, X' . Before, we considered the case where $X = X'$ and have the same partitions; however, the same arguments can be applied even if $X \neq X'$. For example, X and X' could represent two different countries, such as Italy and Germany. Italy's partitioning may include regions \rightarrow provinces \rightarrow municipalities, while Germany's may consist of states \rightarrow districts \rightarrow municipalities. The dataset might then represent trips between municipalities in Italy and Germany, which can be aggregated at coarser levels within their respective hierarchies. In the next paragraph, we introduce the tree construction in the case X, X' having the same number of partitions $g = g'$.

The Destination and Origin Trees. We explain the construction for the *destination tree*, the origin tree will follow naively. The destination tree is a rooted tree that contains information about *intra* and *cross* level range queries of order one. Any node in the tree contains an origin area u_{ℓ_1} , a destination area v_{ℓ_2} , and a range query $q(u_{\ell_1}, v_{\ell_2})$ with the property that it can be obtained by summing the queries of its child nodes. The root node contains the intra-level query at the zero level, hence the triple $(u_0, v_0, q(u_0, v_0) = n)$. The construction then follows an iterative two-step procedure. Given a node $(u_\ell, v_\ell, q(u_\ell, v_\ell))$:

- (1) create a child node for each finer destination $v_{\ell+1} \in v_\ell$ with attribute $(u_\ell, v_{\ell+1}, q(u_\ell, v_{\ell+1}))$;
- (2) for each child node having destination $v_{\ell+1}$, expand the branch by creating a child node for each finer origin $u_{\ell+1} \in u_\ell$ with attribute $(u_{\ell+1}, v_{\ell+1}, q(u_{\ell+1}, v_{\ell+1}))$.

Each iteration adds two levels in the tree, first by adding a cross-level hierarchical query of order one, then by adding intra-level

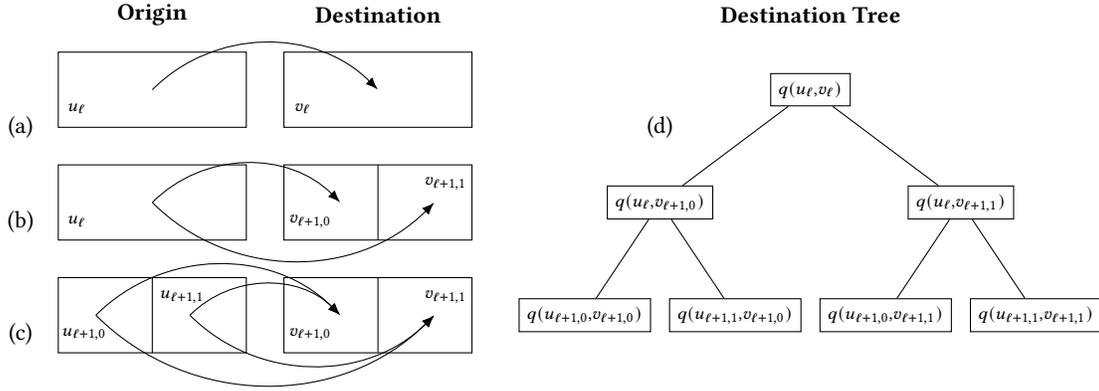


Figure 2: Example of the two-step construction for the destination tree, represented in the left figure from the top to the bottom. In Figure (a), we have two areas at level ℓ , u_ℓ and v_ℓ , and an arrow with attribute $q(u_\ell, v_\ell)$ indicating the flow between them. In Figure (b), the first step is depicted, the destination area v_ℓ is divided into its child areas $v_{\ell+1,0}$ and $v_{\ell+1,1}$ (in this example, we used a bi-partition). The arrows indicate the cross-level range query of order one. In Figure (c), the last step is depicted, the origin area is divided as well, and the arrows indicate the intra-level query of the finer geographic level $\ell + 1$. Figure (d) depicts the destination tree. The links assure hierarchical consistency such that the value of a node can be obtained as the sum of the values of its children.

queries, ending with a tree of $T = 2g + 1$ levels. Figure 2 offers an example of the two-step construction of the destination tree. Using the geographic hierarchy of Italy as an example, the destination tree at level 1 represents a histogram of trips from the country level to the regions, level 2 captures trips between regions, level 3 corresponds to trips from regions to provinces, continuing up to level 6, which contains trips between municipalities. The origin tree can be obtained similarly by selecting finer origins in step (a). The choice of using the destination or origin tree depends on what the practitioners wish to focus on. If cross-level queries starting from origins belonging to larger areas (e.g., regions) and ending to destinations belonging to smaller areas (e.g., provinces) are more important, then the destination tree is the best choice. In the opposite case, we advise to choose the origin tree. Finally, it is important to state that from the non-negative hierarchical tree, we can obtain the original O/D dataset. This is because the leaves of the tree represent the histogram of the dataset D , indicating the absolute frequency with which each O/D pair (u_g, v_g) is observed.

LEMMA 8 (RELATION WITH NON-NEGATIVE HIERARCHICAL TREE).
The destination (origin) tree is a non-negative hierarchical tree.

PROOF. Any node of the destination (origin) tree is a tuple of O/D pairs with non-negative attributes defined in Definition 4. Given a father node (u_ℓ, v_ℓ) , its set of children is $C(u_\ell, v_\ell) = \{(u_\ell, v_{\ell+1}) : \forall v_{\ell+1} \in v_\ell\}$ (for the origin tree the set of children is $C(u_\ell, v_\ell) = \{(u_{\ell+1}, v_\ell) : \forall u_{\ell+1} \in u_\ell\}$). The hierarchical consistency in Observation 1 states that

$$q(u_\ell, v_\ell) = \sum_{v_{\ell+1} \in v_\ell} q(u_\ell, v_{\ell+1}) = \sum_{x \in C(u_\ell, v_\ell)} q(x),$$

which is the hierarchical property in Equation 1. The analysis applies similarly to the next level of the destination tree (as well as to the origin tree). \square

Errors. Let \mathcal{M} be a (ϵ, δ) -DP mechanism that acts on range queries. We are interested in the *maximum absolute error* for any cross-level range queries of order one and intra-level range queries. Therefore, for $(\ell_1, \ell_2) = \{(0, 0), (0, 1), (1, 1), (1, 2), \dots, (g, g)\}$, the error is

$$\max_{(u_{\ell_1}, v_{\ell_2}) \in P_{\ell_1} \times P_{\ell_2}} |\mathcal{M}(q(u_{\ell_1}, v_{\ell_2})) - q(u_{\ell_1}, v_{\ell_2})| \quad (3)$$

Hence, we are interested in the maximum absolute error for any level of the destination tree. The maximum error defined in Equation 3 can be reformulated as in Equation 2.

4 The Top Down Algorithm

In this section, we present InfTDA, but first let us explain why it is appropriate to use a TopDown approach to ensure hierarchical accuracy. The goal is to release a differentially private tabular data \tilde{D} , allowing the data analyzer to perform any marginal query. Under bounded DP, the total number of users, denoted as n , remains fixed within neighboring datasets; thus, it may be disclosed without compromising privacy. However, when applying a differentially private mechanism directly to the histogram representation of D (that is, at the higher level of the hierarchy), the perturbed aggregated \tilde{n} tends to vary around n , with its variance increasing proportionally to the number of point queries. Consider the case $X = X'$, where the origin and destination geographic spaces are the same. When releasing an O/D dataset involving k areas using the Gaussian mechanism with a constant privacy budget, the variance of \tilde{n} is $\text{Var}(\tilde{n}) = O(k^2)$, as the potential number of O/D pairs is k^2 . Due to the cancellation effect of the unbiased estimates produced by the Gaussian mechanism, the maximum error becomes $O(k)$. In contrast, the Stability Histogram does not produce unbiased estimates, resulting in a maximum error of $O(k^2)$. This reasoning applies at any level of the hierarchy, providing a theoretical upper bound on accuracy that does not satisfy the hierarchical accuracy effect we aim to achieve.

PROPOSITION 1 (MAXIMUM ABSOLUTE ERROR PER LEVEL FOR BASELINES). *Given a non-negative hierarchical tree \mathcal{T} with branching factor b , depth T , and a parameter $\beta \in (0, 1)$. The application of the ρ -zCDP Gaussian mechanism at level T achieves for any $\ell \in [T]$ with probability at least $1 - \beta$*

$$\max_{u_\ell \in V_\ell} |\text{err}(u_\ell)| \leq O\left(b^{\frac{T-\ell}{2}} \sqrt{\frac{\ell}{\rho} \log\left(\frac{b}{\beta}\right)}\right). \quad (4)$$

While, for $\beta = n\delta$, the application of the (ϵ, δ) -DP Stability Histogram mechanism at level T achieves with probability at least $1 - \beta$

$$\max_{u_\ell \in V_\ell} |\text{err}(u_\ell)| \leq O\left(\frac{\min(b^{T-\ell}, n) \log(1/\delta)}{\epsilon}\right). \quad (5)$$

The proof of the Proposition can be found in Appendix A. To solve this problem, we could compute the DP estimates of each range query by reallocating the privacy budget among the geographic levels. However, this approach returns inconsistent information about the dataset. For example, the computed flow between two regions might appear smaller than the aggregate flows between their constituent cities. This issue can be addressed by reconciling the estimates with the closest possible values that satisfy certain consistency constraints, as done in the CDBP mechanism [15] and the Hierarchical mechanism [19]. However, the first solves a unique optimization problem for the entire set of queries, yielding a solution that does not scale well, while the latter may return queries with negative values.

We propose a different approach, based on TDA developed by the US Census [1]. For a non-negative hierarchical tree, we iterate a differentially private algorithm starting from the root. At each iteration, an optimization problem is solved using information from the previous level. Unlike TDA, which employs an optimization with an ℓ_2 objective function, our optimization algorithm minimizes an ℓ_∞ objective function, specifically the Chebyshev distance to the noisy vector. We demonstrate both theoretically and experimentally that this is a valid alternative to ensure hierarchical accuracy.

In contrast to TDA, where ℓ_2 minimization leads to a unique solution, minimizing the Chebyshev distance yields multiple optimal solutions. In Section 4.2 we developed an algorithm for integer-constrained optimization that returns an optimal solution that effectively reduces false positives. Another advantage of this approach is that it operates entirely in the integer domain. By comparison, TDA constrained optimization introduced in [1] proceeds in two phases: first, it solves the constrained optimization in the real domain (a relaxation of the integer problem), and then performs a secondary optimization to determine the best rounding. We now present in detail our algorithm.

4.1 InfTDA: TDA with Chebyshev Distance

The TopDown Gaussian Optimized Mechanism with Chebyshev distance optimization InfTDA, operates on the non-negative hierarchical tree, such as the destination tree introduced in Section 3. Similarly to TDA, this method applies discrete Gaussian noise to each level of the tree in a TopDown way, followed by a constrained optimization procedure before descending to the next level.

Since we use bounded privacy, the root attribute, representing the total number of users n in the dataset, can be released without compromising privacy. However, if unbounded privacy were

Algorithm 1 InfTDA

Require: Tree $\mathcal{T} = (V, E)$ of depth T , privacy budget $\rho > 0$.

```

1:  $\tilde{V}_0 \leftarrow \{(u_0, n)\}$ 
2: for  $\ell \in (1, \dots, T)$  do
3:    $\tilde{V}_\ell \leftarrow \{\}$  ▷ DP nodes at level  $\ell$ 
4:   for  $(u, c) \in \tilde{V}_{\ell-1}$  do ▷ Go through the constraints
5:      $\tilde{\mathbf{q}} \leftarrow \mathbf{q}_C(u) + \mathcal{N}_{\mathbb{Z}}(0, T/\rho)^{\dim(\mathbf{q}_C(u))}$  ▷ Apply noise
6:      $\bar{\mathbf{q}} \leftarrow \text{IntOpt}_\infty(\tilde{\mathbf{q}}, c)$  ▷ Solve optimization
7:      $C \leftarrow C(u)$  ▷ Collect set of child nodes of  $u$ 
8:      $X \leftarrow \{(C_j, \bar{q}_j) : \bar{q}_j > 0\}$  ▷ Drop zero attributes
9:      $\tilde{V}_\ell \leftarrow \tilde{V}_\ell \cup X$  ▷ Update level
10: return  $\tilde{\mathcal{T}} \leftarrow (\cup_{\ell=0}^T \tilde{V}_\ell, E)$  ▷ DP Tree

```

required, perturbing the root attribute would become necessary. The algorithm then perturbs the attributes at the first level of the tree using a discrete Gaussian mechanism. The resulting vector is then post-processed to satisfy the hierarchical consistency and non-negativity constraints by solving an integer optimization problem. For each optimized node, the algorithm selects its child nodes, applies a discrete Gaussian mechanism to their attributes, and optimizes them to ensure they are non-negative integers that sum to the attribute of the parent node. The procedure is executed iteratively until the final level T is reached and optimized.

The detailed pseudocode of InfTDA is provided in Algorithm 1. The process begins by constructing the root \tilde{V}_0 of the differentially private tree in line 1. Here, u_0 denotes the root node of the input tree, while n represents its attribute. The algorithm then starts the TopDown loop in line 2. Each iteration aims to construct the set of nodes at level ℓ of the differentially private tree, instantiated in line 3. Each node of the previous level $\ell - 1$ is sampled in line 4 and used as a constraint. In line 5 the discrete Gaussian mechanism with ρ/T privacy budget (for zCDP) is applied to the attributes $\mathbf{q}_C(u)$ of the child nodes of the constraint. Then, in line 6 the private attributes $\tilde{\mathbf{q}}$ are post-processed to satisfy the constraints. The algorithm IntOpt solves the following integer optimization problem by minimizing the Chebyshev distance

$$\begin{aligned} \mathcal{P}(\mathbf{x}, c) &:= \arg \min_{\mathbf{y}} \|\mathbf{x} - \mathbf{y}\|_\infty & (6) \\ \text{s.t. } & y_i \in \mathbb{N}_0 \quad \forall i \in [1, \dots, b] \\ \text{s.t. } & \sum_{i=1}^b y_i = c. \end{aligned}$$

The algorithm IntOpt is described in Section 4.2. In line 7, the set of child nodes of the constraint is constructed, and it is augmented with the corresponding post-processed DP attributes in line 8, dropping nodes with zero attributes. This final step effectively reduces the size of the DP tree and the running time of the algorithm, particularly for sparse datasets. In fact, if there is a node u_ℓ with optimized attribute $\bar{q}(u_\ell) = 0$, then, by consistency, the entire branch of the tree starting at u_ℓ will also have nodes with zero attributes. Finally, in line 9 the set of DP nodes is updated, and this set will serve as constraints in the next iteration. The algorithm outputs a differentially private tree, with optimized attributes. Note that the leaves of

the DP tree constitute the histogram representation of differentially private tabular data.

THEOREM 9 (PRIVACY OF InfTDA). *InfTDA satisfies $\text{GS}_2(q)^2 \frac{\rho}{2}$ -zCDP under bounded privacy.*

PROOF. The attribute of each node at an even level of the tree represents a cross-range query of order one, while nodes at odd levels contain intra-level queries. Hence, an entire level is a histogram of cross or intra-level queries, which have general 2-global sensitivity $\text{GS}_2(q)$. Theorem 5 implies that each iteration of the TopDown loop uses $\text{GS}_2(q)^2 \frac{\rho}{2T}$ privacy budget. As the loop goes through T levels, by the composition and post-processing properties, the algorithm satisfies $\text{GS}_2(q)^2 \frac{\rho}{2}$ -zCDP. \square

Different Privacy Types and Sensitivities. If each user in the dataset contributes m distinct trips, the 2-global sensitivity becomes $\sqrt{2m}$ for bounded privacy, and \sqrt{m} for unbounded privacy. In cases where each user contributes m trips without requiring them to be distinct, the 2-global sensitivity becomes $\sqrt{2m}$ for bounded privacy and m for unbounded privacy. An algorithm that takes into account unbounded privacy is provided in Appendix B. Essentially, it privatizes n in \tilde{V}_0 and rescales the variance of the Gaussian mechanism. Our experiments focus on bounded privacy with $m = 1$.

COROLLARY 10. *When each user in the O/D dataset used to construct the tree contributes a single trip, InfTDA satisfies ρ -zCDP under bounded privacy.*

We now provide an upper bound for the maximum absolute error for each level of the tree.

THEOREM 11 (UTILITY OF InfTDA). *Given a non-negative hierarchical tree \mathcal{T} with branching factor b , depth T , and a parameter $\beta \in (0, 1)$. For each level $\ell \in [1, \dots, T]$, InfTDA with privacy budget $\rho > 0$ achieves with probability at least $1 - \beta$*

$$\max_{u_\ell \in \tilde{V}_\ell} |\text{err}(u_\ell)| \leq O\left(\sqrt{\frac{\ell^3 T}{\rho} \log\left(\frac{b\ell}{\beta}\right)}\right)$$

PROOF. The algorithm applies Gaussian noise in a TopDown way to all attributes of the nodes, except the root. Then, it performs an optimization procedure at each level. Consider a node u_ℓ at level $\ell \in [1, \dots, T]$, with attribute $q(u_\ell)$. Let $\tilde{q}(u_\ell)$ be the attribute returned by the Gaussian mechanism before the optimization is applied, and $\bar{q}(u_\ell)$ after the optimization. By triangle inequality we have

$$|\text{err}(u_\ell)| = |q(u_\ell) - \bar{q}(u_\ell)| \leq |q(u_\ell) - \tilde{q}(u_\ell)| + |\tilde{q}(u_\ell) - \bar{q}(u_\ell)|. \quad (7)$$

The first term is just the absolute value of a Gaussian random variable, so we focus on the second term. Let $u_{\ell-1}$ be the father node of u_ℓ , then $\tilde{q}(u_\ell)$ is an element of the vector solution to the optimization problem $\tilde{\mathbf{q}}_C(u_{\ell-1}) = \mathcal{P}(\tilde{\mathbf{q}}_C(u_{\ell-1}), c)$, where $c = \tilde{q}(u_{\ell-1})$, thus

$$|\tilde{q}(u_\ell) - \bar{q}(u_\ell)| \leq \|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \bar{\mathbf{q}}_C(u_{\ell-1})\|_\infty.$$

Let us consider another vector ξ , called the *offset*, such that $\mathbf{q}_C(u_{\ell-1}) + \xi$ lies within the feasible region of the constrained optimization

problem in Equation 6, then

$$q_{C,j}(u_{\ell-1}) + \xi_j \geq 0 \quad (8)$$

$$\sum_{j=1}^b q_{C,j}(u_{\ell-1}) + \xi_j = \bar{q}(u_{\ell-1}). \quad (9)$$

As the vector $\bar{\mathbf{q}}_C(u_{\ell-1})$ is a solution to the optimization problem, it minimizes the Chebyshev distance with $\tilde{\mathbf{q}}_C(u_{\ell-1})$ under the non-negativity and summation constraints, then by triangle inequality

$$\begin{aligned} \|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \bar{\mathbf{q}}_C(u_{\ell-1})\|_\infty &\leq \|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \mathbf{q}_C(u_{\ell-1}) - \xi\|_\infty \\ &\leq \|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \mathbf{q}_C(u_{\ell-1})\|_\infty + \|\xi\|_\infty. \end{aligned}$$

As $|q(u_\ell) - \tilde{q}(u_\ell)| \leq \|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \mathbf{q}_C(u_{\ell-1})\|_\infty$, the upper bound in Equation 7 becomes

$$|\text{err}(u_\ell)| \leq 2\|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \mathbf{q}_C(u_{\ell-1})\|_\infty + \|\xi\|_\infty. \quad (10)$$

The problem is now to find an upper bound for $\|\xi\|_\infty$.

To upper bound $\|\xi\|_\infty$, we now construct an example of ξ satisfying the constraints and having a bounded ℓ_∞ norm. By construction, from Equation 9 we have that

$$\begin{aligned} \sum_{j=1}^b \xi_j &= \bar{q}(u_{\ell-1}) - \sum_{j=1}^b q_{C,j}(u_{\ell-1}) = \bar{q}(u_{\ell-1}) - q(u_{\ell-1}) \\ &= \text{err}(u_{\ell-1}). \end{aligned}$$

If $\text{err}(u_{\ell-1}) \geq 0$ we can take $\xi_j = \frac{\text{err}(u_{\ell-1})}{b}$ for any $j \in [b]$ as a solution satisfying the summation constraint and the inequality constraint in Equation 8. However, this is not sufficient. If $\text{err}(u_{\ell-1}) < 0$ the inequality constraint might not be satisfied. In this scenario we might consider a solution where $\xi_i = 0$ for any $i \in [b] \setminus \{i^*\}$ where $\xi_{i^*} = -|\text{err}(u_{\ell-1})|$. Any zero element satisfies the constraint in Equation 8 as $q_{C,j}(u_{\ell-1}) \geq 0$. If $\xi_{i^*} \geq -q_{C,i^*}(u_{\ell-1})$ then we finish and obtain an upper bound $\|\xi\|_\infty \leq |\text{err}(u_{\ell-1})|$. In the other case where $\xi_{i^*} < -q_{C,i^*}(u_{\ell-1})$ we need to augment ξ_{i^*} up to meet $-q_{C,i^*}(u_{\ell-1})$. By doing so we increase $\sum_i \xi_i$ making it necessary to decrease some elements of the offset. As we are reducing elements that initially are zero, the new offset still contains only negative elements, and as $\sum_i \xi_i = -|\text{err}(u_{\ell-1})|$ any element cannot be less than $-|\text{err}(u_{\ell-1})|$. Thus, we conclude that there always exists an offset such that $\|\xi\|_\infty \leq |\text{err}(u_{\ell-1})|$.

Now we continue from the upper bound in Equation 10

$$|\text{err}(u_\ell)| \leq 2\|\tilde{\mathbf{q}}_C(u_{\ell-1}) - \mathbf{q}_C(u_{\ell-1})\|_\infty + |\text{err}(u_{\ell-1})|.$$

Completing the recurrence relation by ending at $\text{err}(u_0) = 0$ we get

$$|\text{err}(u_\ell)| \leq 2 \sum_{\kappa=1}^{\ell-1} \|\tilde{\mathbf{q}}_C(u_{\ell-\kappa}) - \mathbf{q}_C(u_{\ell-\kappa})\|_\infty. \quad (11)$$

For $\ell = 0$, we have $\text{err}(u_0) = 0$ under bounded privacy, while for $\ell = 1$ the error is $|\text{err}(u_1)| \leq 2\|\tilde{\mathbf{q}}_C(u_0) - \mathbf{q}_C(u_0)\|_\infty$ which is twice the deviation caused by adding Gaussian noise alone. In Equation 11 we sum the ℓ_∞ norms of Gaussian random vectors with zero mean and variance T/ρ , hence by applying the tail bound in Corollary 6 and a union bound over the dimension b of each vector and ℓ levels, we get for any $\beta \in (0, 1)$

$$\Pr \left[\max_{\kappa \in [\ell]} \|\tilde{\mathbf{q}}_C(u_\kappa) - \mathbf{q}_C(u_\kappa)\|_\infty \geq O\left(\sqrt{\frac{T}{\rho} \log\left(\frac{b\ell}{\beta}\right)}\right) \right] \leq \beta.$$

Algorithm 2 ℓ_∞ Integer Optimization (IntOpt $_\infty$)

Require: $\mathbf{x} \in \mathbb{Z}^d, c \in \mathbb{N}$.

- 1: $\mathbf{z} \leftarrow \max \left(\left\lceil \frac{c - \sum_i x_i}{d} \right\rceil, -\mathbf{x} \right)$
- 2: $t \leftarrow \|\mathbf{z}\|_\infty$
- 3: $I \leftarrow$ sorted indices of \mathbf{x} in ascending order
- 4: $j \leftarrow 0$
- 5: **while** $\sum_i z_i > c - \sum_i x_i$ **do**
- 6: $\Delta \leftarrow \sum_i z_i - c + \sum_i x_i$
- 7: $z_{I[j]} \leftarrow \max(z_{I[j]} - \Delta, -x_{I[j]}, -t)$
- 8: $j \leftarrow (j + 1) \bmod |I|$
- 9: **if** $j = 0$ **then**
- 10: $t \leftarrow t + 1$
- 11: **return** $\mathbf{x} + \mathbf{z}$

Thus, Equation 11 yields the following upper bound with probability at least $1 - \beta$

$$|\text{err}(u_\ell)| \leq O \left(\ell \sqrt{\frac{T}{\rho} \log \left(\frac{b\ell}{\beta} \right)} \right).$$

The claim follows by a union bound over $b^\ell = |V_\ell|$ nodes at level ℓ , leading to an additional $\sqrt{\ell}$ factor. \square

4.2 IntOpt: Integer Optimization with Chebyshev Distance

In this section, we present an algorithm to solve the integer optimization problem aimed at minimizing the ℓ_∞ norm, with special attention to reducing false positives. Given a vector of integers $\mathbf{x} \in \mathbb{Z}^b$, representing the output of a differentially private mechanism, and a natural number $c \in \mathbb{N}$, the integer optimization problem $\mathcal{P}(\mathbf{x}, c)$, as defined in Equation 6, can be reformulated by introducing $\mathbf{z} = \mathbf{y} - \mathbf{x}$. Minimizing $\|\mathbf{z}\|_\infty$ is equivalent to solving the following linear problem

$$\begin{aligned} \min \alpha \quad & \text{s.t.} \quad -\alpha \leq z_i \leq \alpha && \forall i \in [1, \dots, b] \\ & \text{s.t.} \quad z_i \geq -x_i && \forall i \in [1, \dots, b] \\ & \text{s.t.} \quad \sum_{i=1}^b z_i = c - \sum_{i=1}^b x_i. \end{aligned} \quad (12)$$

With this reformulation, we can compute a lower bound for the minimum $\alpha^* = \|\mathbf{z}\|_\infty$ satisfying the constraints of the problem in Equation 12.

LEMMA 12. *Let α^* be the solution of the linear program in Equation 12, then*

$$\alpha^* \geq \max \left(\left\lceil \frac{c - \sum_{i=1}^b x_i}{b} \right\rceil; -\min_i x_i \right). \quad (13)$$

PROOF. From the constraint $-x_i \leq z_i \leq \alpha$, it follows that $\alpha \geq -x_i$ for all $i \in [b]$. Thus, α must satisfy $\alpha \geq \max_i(-x_i) = -\min_i x_i$. Additionally, the equality constraint, combined with $-\alpha \leq z_i \leq \alpha$, implies that $\alpha \geq \left\lceil \frac{c - \sum_i x_i}{b} \right\rceil$. Therefore, we deduce

$$\alpha \geq \max \left(\left\lceil \frac{c - \sum_i x_i}{b} \right\rceil, -\min_i x_i \right),$$

for the relaxed problem in the real domain. Since the feasible region of the relaxed real problem includes the feasible region of the integer problem, the final value of α is obtained by applying the ceiling function. \square

Note also that the solution is not generally unique. For instance, consider $\mathbf{x} = (0, -1, 1)$ and $c = 2$. In this case, there are two possible solutions $\mathbf{y}_1 = (1, 0, 1)$ or $\mathbf{y}_2 = (0, 0, 2)$ both of which have a Chebyshev distance of 1 from \mathbf{x} . In Algorithm 2 we present a simplified version of our optimizer. A faster implementation, guaranteed to run in polynomial time with respect to b , can be found in Appendix C. The core idea is to initialize a solution that satisfies the inequality constraints, achieves a small ℓ_∞ norm, and has a summation exceeding the required value. The algorithm then iteratively reduces the entries of the solution to meet the summation constraint while minimizing any increase in the objective function.

LEMMA 13 (OPTIMALITY). *Algorithm 2 returns a solution that minimizes the Chebyshev distance.*

PROOF. In line 1, we propose our initial solution. First, we demonstrate that its summation exceeds the required value:

$$\sum_{i=1}^b \max \left(\left\lceil \frac{c - \sum_{i=1}^b x_i}{b} \right\rceil, -x_i \right) \geq \sum_{i=1}^b \left\lceil \frac{c - \sum_{i=1}^b x_i}{b} \right\rceil \geq c - \sum_{i=1}^b x_i.$$

Next, we prove that $\|\mathbf{z}\|_\infty$ is at most equal to the lower bound in Equation 13. When $c - \sum_i x_i < 0$, we obtain:

$$\max_i |z_i| \leq \max \left(\left\lceil \frac{c - \sum_{i=1}^d x_i}{d} \right\rceil, -\min_i x_i \right), \quad (14)$$

which follows from $\left\lceil \frac{c - \sum_i x_i}{d} \right\rceil = \left\lceil \frac{c - \sum_i x_i}{d} \right\rceil - 1$ ⁶. Conversely, when $c - \sum_i x_i \geq 0$, $\|\mathbf{z}\|_\infty$ matches the lower bound in Equation 13. Thus, the algorithm starts with a vector that satisfies the inequality constraints, has a small ℓ_∞ norm, and its entries sum to a value larger than what is required. In the loop from lines 5 to 10, each entry in the vector \mathbf{z} is reduced iteratively until its total summation satisfies the constraint. To guarantee optimality, it is crucial to ensure that no entry is excessively reduced, thus avoiding unnecessary increases in $\|\mathbf{z}\|_\infty$. This consideration is addressed in line 2, where the algorithm identifies the smallest possible entry of \mathbf{z} such that $\|\mathbf{z}\|_\infty$ remains unchanged. In line 7, the algorithm updates z_i . When set to $z_i - \Delta$, where Δ is defined in line 6 as the positive remainder, the process terminates, and the result $\mathbf{y} = \mathbf{x} + \mathbf{z}$ is returned. The update respects the inequality constraint $z_i \geq -x_i$ and the optimality condition $z_i \geq -t$. If a solution is achieved in the first round of updates, it is guaranteed to be optimal since t corresponds to the lower bound given in Equation 13. Otherwise, t increases by one, allowing smaller entries and thereby increasing $\|\mathbf{z}\|_\infty$ by 1, which is the minimum increase. A solution and therefore the end of the cycle is always guaranteed as if no updates are possible then $\mathbf{z} = -\mathbf{x}$ and so $\sum_i z_i \leq c - \sum_i x_i$ for $c \geq 0$. \square

Reducing False Positives. The updates in line 7 can be performed iteratively using any permutation I of the indices of \mathbf{z} . In line 3, we propose a specific permutation. The idea behind this choice is that when \mathbf{x} results from a differentially private mechanism, false positives are likely to be associated to attributes with small values.

⁶For $a < 0$, we have $\lceil a \rceil = -\lceil -a \rceil = \lfloor -a \rfloor = \lfloor |a| \rfloor = \lceil |a| \rceil - 1$.

Consequently, the permutation in line 3 prioritizes reducing the smaller elements first, potentially setting them to zero (i.e., $y_i = 0$). Alternatively, an inverse approach can be taken in line 3 by sorting the indices of \mathbf{x} in descending order, which would focus on reducing false negatives instead.

5 Experiments

This section provides an experimental evaluation of InfTDA against various baselines in real-world and synthetic datasets. We start by presenting the baselines used, then we introduce the datasets and the experimental setup.

Baselines. As simple baselines, we use the injection of discrete Gaussian noise, and the application of the Stability-Based histogram, on the leaves’ attributes of the tree. They will be labeled respectively as VanillaGauss and SH. We are aware of more accurate baselines for sparse histograms [4, 11]. However, in terms of maximum absolute error utility, they offer the same asymptotic performance as SH. VanillaGauss offers stronger guarantees for range queries compared to SH, but it is space-inefficient and may produce negative counts. Conversely, SH generates datasets with non-negative counts and no false positives; however, it performs poorly for range queries, especially when the dataset consists predominantly of rare items. These baselines provide experimental evidence underscoring the necessity of using TopDown constraint optimization algorithms like TDA or our proposed InfTDA. We evaluate InfTDA against two variations of TopDown algorithms, each employing a different optimization strategy. The first is a simple implementation of TDA, referred to as TDA_{ℓ_2} , which uses the Euclidean distance as the objective function⁷. The second, denoted as TDA_{ℓ_∞} , incorporates Chebyshev optimization but relies on a black-box solver for the optimization process. This latter baseline allows us to assess the effectiveness of our optimizer, IntOpt, particularly in reducing false positives. We recognize the possibility of employing CBDP and the Hierarchical mechanism as a baseline. However, the first faces significant implementation obstacles due to the vast quantity of O/D pairs in the analyzed dataset, while the second does not work in the integer domain.

Real Dataset. The real-world dataset under examination originates from the Italian National Institute of Statistics (ISTAT) [20]. The dataset contains 28,805,440 commutes for 2011, which are trips departing from the habitual residence and returning to the same place daily. These trips were obtained from survey questionnaires in which each individual was asked to provide the address of their usual place of study or work. As each individual contributes one trip to the dataset, we apply the privacy analysis outlined in Corollary 6. This dataset is structured with a geographic hierarchy consisting of five partitions: regions, provinces, municipalities, section areas, and census sections, with regions representing the highest level (i.e., the largest areas). The first three geographic partitions are shown in Figure 1. The dataset exhibits significant sparsity, containing 362,292 census sections, which theoretically could result in over 100 billion possible O/D pairs. However, it only records 14,287,549

Dataset	Number of users	Number of O/D
Binary Complete	1051271	65536
Binary Dense	734688	32768
Binary Sparse	23302	655
Random Complete	2019580	189225
Random Dense	1003943	95612
Random Sparse	67840	1892

Table 1: Characteristics of synthetic datasets.

actual flows. Due to computational limitations, we considered only O/D pairs up to the municipality level for our experiments. This results in approximately 500,000 O/D pairs out of more than 65 million possible pairs. We generate the destination tree from this dataset, obtaining a tree of depth $T = 6$.

Synthetic Datasets. We generate two types of synthetic partitions, a *binary* partition and a *random* partition, to create six distinct O/D datasets. The binary partition consists of 8 hierarchical levels, where each area is iteratively divided into two smaller areas, resulting in a binary destination tree with a depth of $T = 16$. In contrast, the random partition has 4 levels, where each area is randomly divided into k smaller areas, with k sampled uniformly from 2 to 10. This latter approach simulates real-world scenarios where areas are partitioned unevenly, creating a destination tree with a variable structure. The O/D flows are sampled from a continuous Pareto distribution ($\Pr(x) \sim x^{-\beta}$, where x represents the flow and $\beta > 0$) and rounded, a common pattern observed in mobility and social data [3, 18]. These flows are then assigned as attributes to the leaves of the generated trees. To evaluate the mechanisms under varying levels of sparsity, we simulate three scenarios: *complete*, where all leaves have positive attributes; *dense*, where 50% of the leaves are assigned positive attributes; and *sparse*, where only 1% of the leaves are assigned positive attributes. This allows for a comprehensive testing of the performance of the mechanisms under different sparsity conditions. The number of users generated and the number of O/D in the synthetic dataset can be found in Table 1.

Experimental Setup. The system was developed using open source libraries and Python 3.11. Our approach to differential privacy leverages the OpenDP library⁸ [24], which includes implementations of the discrete Gaussian mechanism and the Stability Histogram. For black-box optimization, we use cvxpy [2, 12]. TDA_{ℓ_2} first minimizes the Euclidean distance for the relaxed program in the real domain using CLARABEL [17], then it rounds and redistributes the exceeding in a similar fashion of InfTDA, hence by prioritizing the elimination of small values. TDA_{ℓ_∞} optimization works completely in the integer domain and uses the GLPK mixed integer optimizer. The tests were conducted using an Intel Xeon Processor W-2245 (8 cores, 3.9GHz), 128GB RAM, and Ubuntu 20.04.3. The experiments and the code are publicly available⁹. Each mechanism is run 10 times, and the error bars in the graphs indicate the maximum and minimum value of the metric considered.

⁷We remark that the real implementation of TDA would require weighted non-negative least squares optimization, and a complex two step optimization to obtain integer values.

⁸<https://github.com/opendp/opendp>

⁹https://github.com/aidaLabDEI/TDA_hierarchical

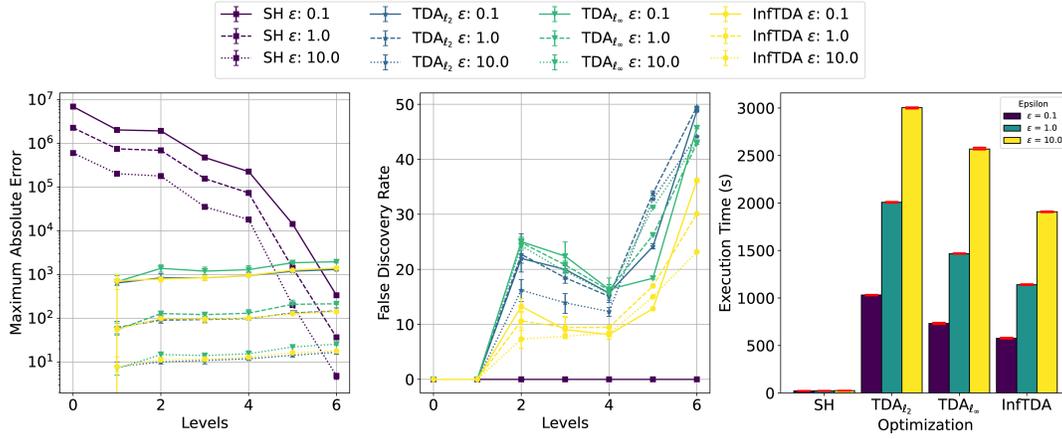


Figure 3: Experiments run for the Italian dataset (from ISTAT). From left to right: maximum absolute error, false discovery rate, and execution time. The error bars indicate maximum and minimum values over 10 experiments.

Privacy budget. The experiments were carried out with $\epsilon \in [0.1, 1, 10]$ and $\delta = 10^{-8}$ (sufficient to ensure that $\delta \ll 1/n$). Privacy budget $\epsilon \leq 1$ provides strong real-world privacy in most cases [29], however, there are examples of deployments with larger values. The US Census of 2020 distributed the population data under $(17.14, 10^{-10})$ -DP [27]. Unfortunately, such a high privacy budget may give meaningless theoretical guarantees in the worst-case scenario¹⁰, but still may offer good privacy under some specific and more realistic attacks. Although we recognize that intelligent allocation of the privacy budget across levels can improve the utility of a differentially private dataset (as done by the US Census [27]), we opted for a uniform distribution of the budget across all levels. We used Lemma 2 to compute the privacy budget for zCDP.

Quality of the DP Dataset. To assess the quality of the DP dataset \tilde{D} we measured two key indicators that can be computed at any level of the tree: the *maximum absolute error* defined in Equation 2 and the *false discovery rate*, which is the percentage of O/D pairs present in \tilde{D} but absent in the original data. In tree notation, for any level $\ell \in [T]$ the false discovery rate is

$$f_d(\ell; \tilde{D}, D) = \frac{|\{u_\ell : \tilde{q}(u_\ell) > 0 \wedge q(u_\ell) = 0\}|}{|\{u_\ell : \tilde{q}(u_\ell) > 0\}|} \cdot 100\%,$$

where D is the original data.

5.1 Discussion

Italian dataset. The analysis of the O/D commuting dataset for Italy is presented in Figure 3, which illustrates three metrics from left to right: the maximum absolute error per level, the false discovery rate per level and the running time. We emphasize that the levels are defined as follows: the zero level represents the total number of users, the second level corresponds to range queries for regions, the fourth level pertains to provinces, and the final level corresponds to

municipalities’ O/D commutes. The VanillaGauss mechanism was not applied to this dataset due to the computational infeasibility of sampling more than 65 million Gaussian noises, one for each potential O/D commute. Instead, this baseline will be evaluated using synthetic data.

We immediately observe that SH performs poorly on range queries, although it provides the most accurate estimates for O/D flows at finer geographic levels. This is not surprising since the mechanism achieves on the leaves a maximum error close to a lower bound [28] (Theorem 5.13). Consistent with our theoretical findings in Theorem 11, InfTDA produces datasets with diminishing accuracy as one moves down the levels of the tree, reflecting improved precision for O/D flows at larger geographic scales. This trend is similarly observed with TDA_{l₂} and TDA_{l_∞}. While TDA_{l₂} shows comparable utility to InfTDA, TDA_{l_∞} underperforms slightly. This is due to the fact that Chebyshev optimization produces multiple optimal solutions, not all of which perform well in practice in terms of utility. This underscores the importance of our optimizer IntOpt, in selecting the most practical solution. This is clearly illustrated in the middle plot of Figure 3, where InfTDA consistently reduces false positive detections compared to TDA_{l_∞} and TDA_{l₂}, regardless of the privacy budget. The reduction in the false discovery rate depends on factors like hierarchical structure and attribute distribution, such as the presence of sparse levels followed by dense ones. Nonetheless, InfTDA consistently outperforms the other mechanisms in reducing false positives.

In terms of running time, SH is clearly the fastest, since it does not need to perform any kind of optimization and runs linearly with the number of users n . InfTDA is the fastest TopDown algorithm. Running time increases with the privacy budget because, in the low-privacy regime, false negatives are more frequent. These occur when attributes that were positive in the sensitive dataset appear as zero in the DP dataset. This reduction in positive attributes typically decreases the size of the returned dataset, leading to fewer optimizations. In conclusion, SH achieves significantly better accuracy, approximately an order of magnitude higher, for flows between municipalities, but performs poorly for range queries. Mechanisms

¹⁰The suspicion of a user’s presence in the dataset for a *strong* attacker—who knows everything about the dataset except their target—grows significantly for large values of ϵ . For instance, with $\epsilon = 10$ and an initial suspicion of 10%, the attacker’s suspicion increases to 99.9% after observing the differentially private dataset. In contrast, for $\epsilon = 1$, the suspicion would increase only to 23%. For further details, refer to [10].

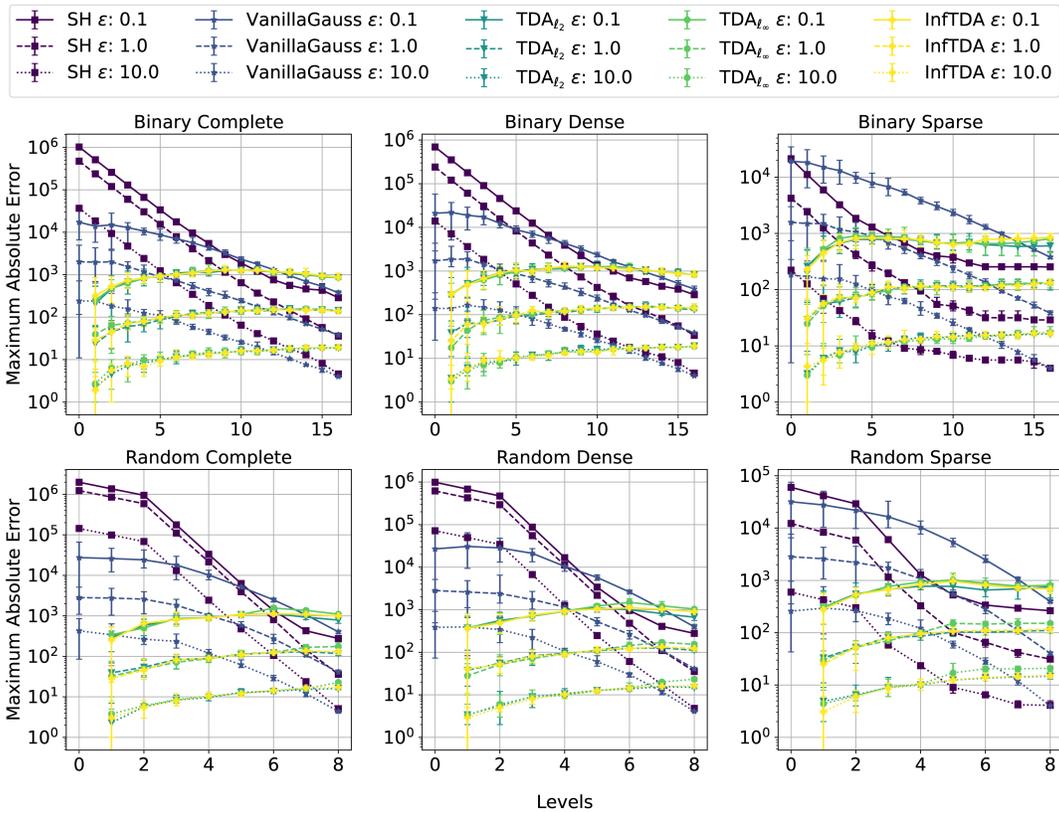


Figure 4: Experiments run for the synthetic datasets with a focus on maximum absolute error. The error bars indicate maximum and minimum values over 10 experiments

based on TDA provide hierarchical accuracy and produce meaningful results even in high privacy regimes. For instance, for $\epsilon = 1$ (resp., $\epsilon = 0.1$), they achieve a maximum absolute error of about 100 (resp., 1000). Additionally, InfTDA generates datasets with fewer false positives and operates more efficiently.

Synthetic Dataset. Figure 4 illustrates the maximum absolute error for the six synthetic datasets we generated. Unlike the previous analysis, this evaluation includes VanillaGauss. As expected, VanillaGauss demonstrates better accuracy than SH for range queries, particularly when the dataset is not highly sparse. In contrast, SH performs better on sparse datasets; however, even in these cases, the error increases significantly for large range queries at the higher levels of the tree. The TopDown algorithms exhibit similar behavior across all scenarios, with one exception: TDA_{l_∞} on the random sparse dataset for $\epsilon \leq 1$. Although this deviation is not substantial compared to the other TDA algorithms, it provides valuable insights consistent with the findings from the Italian dataset. In trees with random branching factors, similar to the structure of the Italy tree, the number of variables in the optimization problem increases. This expansion of the solution space for the Chebyshev distance optimization increases the likelihood of sampling an optimal solution that performs poorly in practice. Conversely, this issue does not arise in binary trees, where each optimization involves

only two variables, making it more likely to sample the best optimal solution for the Chebyshev minimization.

Figure 5 shows the false discovery rate for the sparse datasets. For the binary tree, no significant improvements are observed, indicating that when optimization involves a small number of variables, the choice of the optimization function has minimal impact on the utility of the released datasets. In contrast, for the random tree, InfTDA generates a dataset with fewer false positives compared to the other mechanisms (with the obvious exception of SH). Regarding the execution times of the algorithms, Figure 6 clearly shows that our optimizer is faster than a black-box optimizer. This is not only because InfTDA generates a dataset with fewer false positives but also because it runs faster even for the complete synthetic dataset.

6 Conclusion and Future Directions

We found that a well-designed implementation of the TopDown algorithm can significantly improve the accuracy in the differentially private release of O/D data with a hierarchical structure. This improvement is particularly beneficial when producing tabular datasets where range queries must be more accurate for larger geographical areas. Specifically, this paper explores TopDown algorithms for general tree data structures with non-negative attributes and hierarchical consistency. We illustrate how any O/D dataset

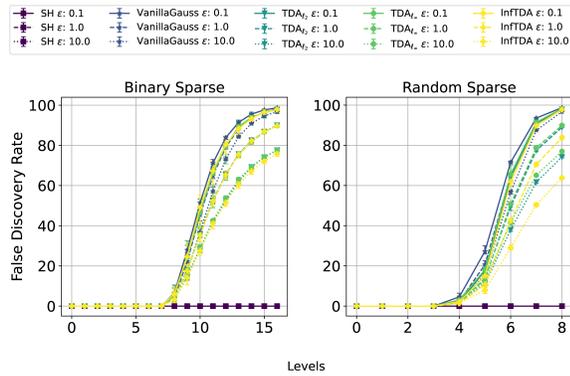


Figure 5: False discovery rate for synthetic datasets

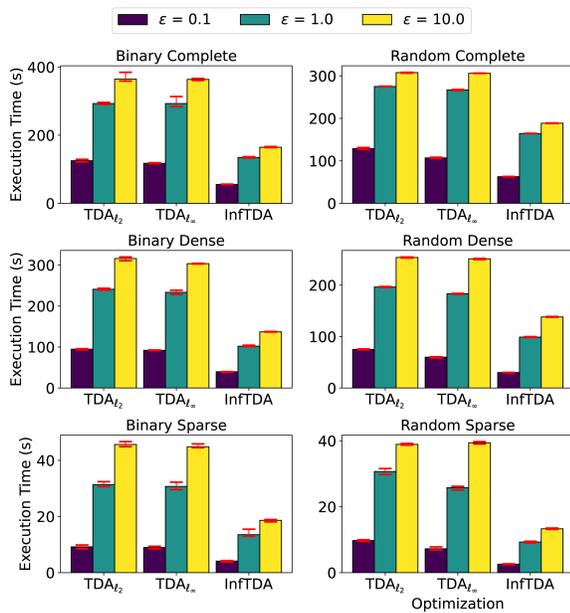


Figure 6: Running time for synthetic datasets and TDA

can be structured in the required tree format. Additionally, we propose a Chebyshev distance constraint optimization problem as an alternative to the commonly used Euclidean distance minimization. This approach yields two key outcomes: a theoretical analysis of the maximum absolute error for a TopDown algorithm and a practical, efficient integer optimization algorithm that has been proven to reduce false positives on both real and synthetic datasets. Our proposed TopDown algorithm, InfTDA, combined with our optimizer, IntOpt, outperforms naive baselines and is not worse than black-box implementations of TDA with different objective functions on both real and synthetic datasets, while being simpler to implement, running faster and effectively reducing the false discovery rate. Given the versatility of our approach, it would be valuable to test it on various real-world datasets beyond O/D data. Any tabular data, such as healthcare data, can be used to construct a non-negative hierarchical tree, provided that a query hierarchy is defined. We

leave this exploration for future research. Furthermore, we believe that studying stability-based algorithms as a source of noise in a TDA would be interesting, particularly in scenarios where the hierarchy consists of unknown partitions. Finally, based on the results of our experiments, we believe that similar theoretical results hold for TDA with Euclidean distance minimization.

Acknowledgments

We would like to thank Martin Aumüller for useful discussions. This work was supported in part by the Big-Mobility project by the University of Padova under the Uni-Impresa call, by the MUR PRIN 20174LF3T8 AHeAD project, and by MUR PNRR CN00000013 National Center for HPC, Big Data and Quantum Computing.

References

- [1] John M. Abowd, Robert Ashmead, Ryan Cumings-Menon, Simson Garfinkel, Micah Heineck, Christine Heiss, Robert Johns, Daniel Kifer, Philip Leclerc, Ashwin Machanavajhala, Brett Moran, William Sexton, Matthew Spence, and Pavel Zhuravlev. 2022. The 2020 Census Disclosure Avoidance System Top-Down Algorithm. *Harvard Data Science Review Special Issue 2* (jun 24 2022). <https://hdsr.mitpress.mit.edu/pub/7evz361>.
- [2] Akshay Agrawal, Robin Verschuere, Steven Diamond, and Stephen Boyd. 2018. A rewriting system for convex optimization problems. *Journal of Control and Decision* 5, 1 (2018), 42–60.
- [3] Laura Alessandretti, Ulf Aslak, and Sune Lehmann. 2020. The scales of human mobility. *Nature* 587, 7834 (2020), 402–407.
- [4] Martin Aumüller, Christian Janos Lebeda, and Rasmus Pagh. 2022. Representing Sparse Vectors with Differential Privacy, Low Error, Optimal Space, and Fast Access. *Journal of Privacy and Confidentiality* 12, 2 (Nov. 2022). <https://doi.org/10.29012/jpc.809>
- [5] Borja Balle and Yu-Xiang Wang. 2018. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. In *International Conference on Machine Learning*. PMLR, 394–403.
- [6] Mark Bun, Kobbi Nissim, and Uri Stemmer. 2016. Simultaneous Private Learning of Multiple Concepts. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science* (Cambridge, Massachusetts, USA) (ITCS '16). Association for Computing Machinery, New York, NY, USA, 369–380. <https://doi.org/10.1145/2840728.2840747>
- [7] Mark Bun and Thomas Steinke. 2016. Concentrated Differential Privacy: Simplifications, Extensions, and Lower Bounds. In *Proceedings, Part I, of the 14th International Conference on Theory of Cryptography - Volume 9985*. Springer-Verlag, Berlin, Heidelberg, 635–658. https://doi.org/10.1007/978-3-662-53641-4_24
- [8] Clement Canonne, Gautam Kamath, and Thomas Steinke. 2022. The Discrete Gaussian for Differential Privacy. *Journal of Privacy and Confidentiality* 12, 1 (Jul. 2022). <https://doi.org/10.29012/jpc.784>
- [9] Graham Cormode, Cecilia Procopiuc, Divesh Srivastava, and Thanh T. L. Tran. 2012. Differentially private summaries for sparse data. In *Proceedings of the 15th International Conference on Database Theory* (Berlin, Germany) (ICDT '12). Association for Computing Machinery, New York, NY, USA, 299–311. <https://doi.org/10.1145/2274576.2274608>
- [10] Damien Desfontaines. 2019. *Differential privacy in (a bit) more detail*. <https://desfontain.es/blog/differential-privacy-in-more-detail.html> Ted is writing things On privacy, research, and privacy research (Personal blog). Accessed: 2025-03-12.
- [11] Damien Desfontaines, James Voss, Bryant Gipson, and Chinmoy Mandayam. 2022. Differentially private partition selection. *Proceedings on Privacy Enhancing Technologies* 1 (2022), 339–352.
- [12] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5.
- [13] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating Noise to Sensitivity in Private Data Analysis. In *Theory of Cryptography*, Shai Halevi and Tal Rabin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 265–284.
- [14] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407. <https://doi.org/10.1561/04000000042>
- [15] Ferdinando Fioretto, Chansoo Lee, and Pascal Van Hentenryck. 2018. Constrained-Based Differential Privacy for Mobility Services. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden) (AAMAS '18). International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1405–1413.

- [16] Sergio Gómez, Alberto Fernández, Sandro Meloni, and Alex Arenas. 2019. Impact of origin-destination information in epidemic spreading. *Scientific reports* 9, 1 (2019), 2315.
- [17] Paul J. Goulart and Yuwen Chen. 2024. Clarabel: An interior-point solver for conic programs with quadratic objectives. arXiv:2405.12762
- [18] Xiao-Pu Han, Qiang Hao, Bing-Hong Wang, and Tao Zhou. 2011. Origin of the scaling law in human mobility: Hierarchy of traffic systems. *Physical Review E* 83, 3 (2011), 036117.
- [19] Michael Hay, Vibhor Rastogi, Gerome Miklau, and Dan Suciu. 2010. Boosting the accuracy of differentially private histograms through consistency. *Proc. VLDB Endow.* 3, 1–2 (sep 2010), 1021–1032. <https://doi.org/10.14778/1920841.1920970>
- [20] Istat. 2011. *Matrici di contiguità, distanza e pendolarismo*. <https://www.istat.it/it/archivio/157423> Accessed: 2025-03-12.
- [21] Daniel Kifer and Ashwin Machanavajjhala. 2011. No free lunch in data privacy. In *Proceedings of the 2011 ACM SIGMOD International Conference on Management of Data (Athens, Greece) (SIGMOD '11)*. Association for Computing Machinery, New York, NY, USA, 193–204. <https://doi.org/10.1145/1989323.1989345>
- [22] Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. 2009. Releasing search queries and clicks privately. In *Proceedings of the 18th International Conference on World Wide Web (Madrid, Spain) (WWW '09)*. Association for Computing Machinery, New York, NY, USA, 171–180. <https://doi.org/10.1145/1526709.1526733>
- [23] Imnook Lee, Shin-Hyung Cho, Kyoungtae Kim, Seung-Young Kho, and Dong-Kyu Kim. 2022. Travel pattern-based bus trip origin-destination estimation using smart card data. *PLOS ONE* 17, 6 (06 2022), 1–20. <https://doi.org/10.1371/journal.pone.0270346>
- [24] Michael Shoemate, Andrew Vyrros, Chuck McCallum, Raman Prasad, Philip Durbin, Silvia Casacuberta Puig, Ethan Cowan, Vicki Xu, Zachary Ratliff, Nicolás Berrios, Alex Whitworth, Michael Eliot, Christian Lebeda, Oren Renard, and Claire McKay Bowen. 2024. *OpenDP Library*. <https://github.com/opendp/opendp>
- [25] Ananda Theertha Suresh. 2019. *Differentially private anonymized histograms*. Curran Associates Inc., Red Hook, NY, USA.
- [26] Marika Swanberg, Damien Desfontaines, and Samuel Haney. 2023. DP-SIPS: A simpler, more scalable mechanism for differentially private partition selection. *Proceedings on Privacy Enhancing Technologies* 4 (2023), 257–268.
- [27] U.S. Census Bureau. 2020. Disclosure Avoidance for the 2020 Census: An Introduction. <https://www.census.gov/library/publications/2021/decennial/2020-census-disclosure-avoidance-handbook.html> Accessed: 2025-03-12.
- [28] Salil Vadhan. 2017. *The Complexity of Differential Privacy*. Springer International Publishing, Cham, 347–450. https://doi.org/10.1007/978-3-319-57048-8_7
- [29] Alexandra Wood, Micah Altman, Aaron Bembenek, Mark Bun, Marco Gaboardi, James Honaker, Kobbi Nissim, David O'Brien, Thomas Steinke, and Salil Vadhan. 2018. Differential Privacy: A Primer for a Non-Technical Audience. *Vanderbilt Journal of Entertainment & Technology Law* 21, 17 (2018). <https://doi.org/10.2139/ssrn.3338027> Berkman Klein Center Research Publication No. 2019-2.
- [30] Jia Xu, Zhenjie Zhang, Xiaokui Xiao, Yin Yang, and Ge Yu. 2012. Differentially Private Histogram Publication. In *2012 IEEE 28th International Conference on Data Engineering*. 32–43. <https://doi.org/10.1109/ICDE.2012.48>
- [31] Xiaojian Zhang, Rui Chen, Jianliang Xu, Xiaofeng Meng, and Yingtao Xie. 2014. Towards accurate histogram publication under differential privacy. In *Proceedings of the 2014 SIAM international conference on data mining*. SIAM, 587–595.

A Additional Proof

PROOF OF PROPOSITION 1. We start by proving the Equation 4. By applying the discrete Gaussian mechanism on each attribute of the final level nodes we can reconstruct the attributes at any level using the hierarchical relation in equation 1. Consider a node u_ℓ at level $\ell \in \{0, \dots, T-1\}$, its private attribute is

$$\tilde{q}(u_\ell) = \sum_{i_{\ell+1} \in \mathcal{C}(u_\ell)} \cdots \sum_{i_T \in \mathcal{C}(u_{T-1})} \tilde{q}(u_T),$$

so the error $\text{err}(u_\ell) = q(u_\ell) - \tilde{q}(u_\ell)$ is

$$\text{err}(u_\ell) = \sum_{u_{\ell+1} \in \mathcal{C}(u_\ell)} \cdots \sum_{u_T \in \mathcal{C}(u_{T-1})} [q(u_T) - \tilde{q}(u_T)]. \quad (15)$$

Each error $q(u_T) - \tilde{q}(u_T)$ at the T level is a Gaussian random variable $\mathcal{N}_{\mathbb{Z}}(0, 1/\rho)$, then the right hand side of equation 15 is a Gaussian random variable $\mathcal{N}_{\mathbb{Z}}(0, b^{T-\ell}/\rho)$, as it is a summation of $b^{T-\ell}$ Gaussian random variables. Then, by applying Corollary 6 and a union

Algorithm 3 InfTDA for unbounded privacy

Require: Tree $\mathcal{T} = (V, E)$ of depth T , privacy budget $\rho > 0$.

```

1:  $\tilde{n} \leftarrow n + \mathcal{N}_{\mathbb{Z}}\left(0, \frac{T}{\rho}\right)$ 
2:  $\tilde{V}_0 \leftarrow \{(u_0, \tilde{n})\}$ 
3: for  $\ell \in (1, \dots, T)$  do
4:    $\tilde{V}_\ell \leftarrow \{\}$  ▷ DP nodes at level  $\ell$ 
5:   for  $(u, c) \in \tilde{V}_{\ell-1}$  do ▷ Go through the constraints
6:      $\tilde{q} \leftarrow \text{qc}(u) + \mathcal{N}_{\mathbb{Z}}\left(0, \frac{T}{\rho}\right)^{\dim(\text{qc}(u))}$  ▷ Apply noise
7:      $\bar{q} \leftarrow \text{IntOpt}_\infty(\tilde{q}, c)$  ▷ Solve optimization
8:      $C \leftarrow C(u)$  ▷ Collect set of child nodes of  $u$ 
9:      $X \leftarrow \{(C_j, \bar{q}_j) : \bar{q}_j > 0\}$  ▷ Drop zero attributes
10:     $\tilde{V}_\ell \leftarrow \tilde{V}_\ell \cup X$  ▷ Update level
11: return  $\tilde{\mathcal{T}} \leftarrow (\cup_{\ell=0}^T \tilde{V}_\ell, E)$  ▷ DP Tree

```

bound over b^ℓ nodes, we obtain for $\beta \in (0, 1)$

$$\Pr \left[\max_{u_\ell \in \tilde{V}_\ell} |\text{err}(u_\ell)| \leq O \left(b^{\frac{T-\ell}{2}} \sqrt{\frac{\ell}{\rho} \log \left(\frac{b}{\beta} \right)} \right) \right] \geq 1 - \beta.$$

We now prove equation 5. Equation 15 is still applicable; however, the noise inserted at level T follows a complicated distribution, which is the result of a symmetric Laplace noise distribution followed by a truncation. So it is not straightforward to compute the distribution of composition of such random variables. However, we have $|\text{err}(u_T)| \leq O \left(\frac{\log(1/\delta)}{\epsilon} \right)$ with probability at least $1 - \delta$. Therefore, by summing the error and applying a union bound on n , since we have at most n random variables to sum, equation 15 leads to $|\text{err}(u_\ell)| \leq O \left(\min(b^{T-\ell}, n) \frac{\log(1/\delta)}{\epsilon} \right)$ with probability at least $1 - n\delta$. As $\delta \ll 1/n$ follows a constant probability upper bound. \square

B InfTDA for unbounded privacy

THEOREM 14. InfTDA in Algorithm 3 satisfies $\left(\frac{m^2}{T} + \text{GS}_2(q)^2 \right) \frac{\rho}{2}$ -zCDP under unbounded privacy.

PROOF. Each iteration of the TopDown loop consumes $\text{GS}_2(q)^2 \frac{\rho}{2T}$ privacy budget. Since the loop goes through T levels, by composition and post-processing properties, the total privacy budget consumed is $\frac{\text{GS}_2(q)^2}{2} \rho$. Additionally, the privatization of the root attribute in line 1 increases the overall privacy budget. Given that n represents a point query and the addition or removal of a single user alters n by m , the ℓ_2 sensitivity is m . Consequently, line 1 contributes an additional privacy budget of $\frac{m^2}{T} \rho$. \square

C Fast IntOpt $_\infty$

This is the algorithm that we implemented in our experiments. One key modification is that the while loop operates only on the indices of elements in \mathbf{z} that can still be reduced, as specified in lines 4 and 11. The main modification is in line 12, where it is computed the smallest entry of \mathbf{z} that, if $z_i = -t - r$ for any $i \in I$, then we would have

$$c - \sum_{i=1}^d x_i \leq \sum_{i=1}^b z_i \leq c - \sum_{i=1}^b x_i + |I|,$$

Algorithm 4 Polynomial time ℓ_∞ Integer Optimization (IntOpt $_\infty$)

Require: $\mathbf{x} \in \mathbb{Z}^b, c \in \mathbb{N}$.

- 1: $\mathbf{z} \leftarrow \max\left(\left\lceil \frac{c - \sum_i x_i}{b} \right\rceil, -\mathbf{x}\right)$
- 2: $t \leftarrow \|\mathbf{z}\|_\infty$
- 3: $I \leftarrow$ sorted indices of \mathbf{x}
- 4: $I \leftarrow (i \in I : z_i > -x_i)$
- 5: $j \leftarrow 0$
- 6: **while** $\sum_i z_i > c - \sum_i x_i$ **do**
- 7: $\Delta \leftarrow \sum_i z_i - c + \sum_i x_i$
- 8: $z_{I[j]} \leftarrow \max(z_{I[j]} - \Delta, -x_{I[j]}, -t)$
- 9: $j \leftarrow (j + 1) \bmod |I|$
- 10: **if** $j = 0$ **then**
- 11: $I \leftarrow (i \in I : z_i > -x_i)$
- 12: $r \leftarrow \lfloor \frac{1}{|I|} (\sum_i z_i - c + \sum_i x_i) \rfloor$
- 13: $t \leftarrow t + \max(1, r)$
- 14: **return** $\mathbf{x} + \mathbf{z}$

ensuring that only one additional reduction round is required to satisfy the equality constraint. If all elements are clipped to $-t - r$, the total summation decreases by $\sum_i z_i - r \cdot |I|$. The first inequality is derived as follows

$$\begin{aligned} \sum_{i=1}^b z_i - r \cdot |I| &= \sum_{i=1}^b z_i - \left\lfloor \frac{1}{|I|} \left(\sum_i z_i - c + \sum_i x_i \right) \right\rfloor \cdot |I| \\ &\geq \sum_{i=1}^b z_i - \frac{1}{|I|} \left(\sum_{i=1}^b z_i - c + \sum_{i=1}^b x_i \right) \cdot |I| = c - \sum_{i=1}^b x_i, \end{aligned}$$

while the second comes from

$$\begin{aligned} \sum_{i=1}^b z_i - r \cdot |I| &\leq \sum_{i=1}^b z_i - \left\lfloor \frac{1}{|I|} \left(\sum_{i=1}^b z_i - c + \sum_{i=1}^b x_i \right) - 1 \right\rfloor \cdot |I| \\ &= c - \sum_{i=1}^b x_i + |I|, \end{aligned}$$

If not all elements are clipped to $-t - r$, at least one element is clipped to $-x_i$, reducing the cardinality of I by at least 1. Consequently, after the first reduction loop (line 10), the algorithm either transitions to the second-to-last loop (where all elements are clipped to $-t - r$) or reduces $|I|$ to $|I| - 1$. In the worst-case scenario, where $|I| = b$, it takes $O(b)$ iterations for $|I|$ to decrease by 1. Thus, the overall runtime of the algorithm is bounded by $O(b^2)$.