

# HDPView: Differentially Private Materialized View for Exploring High Dimensional Relational Data

FumiYuki Kato  
Kyoto University  
fumiYuki@db.soc.i.kyoto-u.ac.jp

Tsubasa Takahashi  
LINE Corporation  
tsubasa.takahashi@linecorp.com

Shun Takagi  
Kyoto University  
takagi.shun.45a@st.kyoto-u.ac.jp

Yang Cao  
Kyoto University  
yang@i.kyoto-u.ac.jp

Seng Pei Liew  
LINE Corporation  
sengpei.liew@linecorp.com

Masatoshi Yoshikawa  
Kyoto University  
yoshikawa@i.kyoto-u.ac.jp

## ABSTRACT

How can we explore the unknown properties of high-dimensional sensitive relational data while preserving privacy? We study how to construct an explorable privacy-preserving materialized view under differential privacy. No existing state-of-the-art methods simultaneously satisfy the following essential properties in data exploration: workload independence, analytical reliability (i.e., providing error bound for each search query), applicability to high-dimensional data, and space efficiency. To solve the above issues, we propose HDPVIEW, which creates a differentially private materialized view by well-designed recursive bisected partitioning on an original data cube, i.e., count tensor. Our method searches for block partitioning to minimize the error for the counting query, in addition to randomizing the convergence, by choosing the effective cutting points in a differentially private way, resulting in a less noisy and compact view. Furthermore, we ensure formal privacy guarantee and analytical reliability by providing the error bound for arbitrary counting queries on the materialized views. HDPVIEW has the following desirable properties: (a) *Workload independence*, (b) *Analytical reliability*, (c) *Noise resistance on high-dimensional data*, (d) *Space efficiency*. To demonstrate the above properties and the suitability for data exploration, we conduct extensive experiments with eight types of range counting queries on eight real datasets. HDPVIEW outperforms the state-of-the-art methods in these evaluations.

### PVLDB Reference Format:

FumiYuki Kato, Tsubasa Takahashi, Shun Takagi, Yang Cao, Seng Pei Liew, and Masatoshi Yoshikawa. HDPView: Differentially Private Materialized View for Exploring High Dimensional Relational Data. PVLDB, 15(9): 1766 - 1778, 2022.

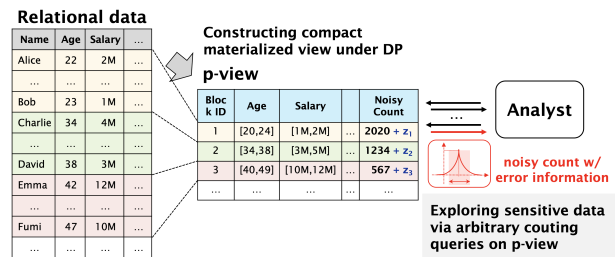
doi:10.14778/3538598.3538601

## 1 INTRODUCTION

In the early stage of data science workflows, exploring a database to understand its properties in terms of multiple attributes is essential to designing the subsequent tasks. To understand the properties, data analysts need to issue a wide variety of range counting queries. If the database is highly sensitive (e.g., personal healthcare records),

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. 15, No. 9 ISSN 2150-8097.  
doi:10.14778/3538598.3538601



**Figure 1: Data exploration through a *privacy-preserving materialized view (p-view for short)* of a multidimensional relational data. The p-view works as an independent query system. Analysts can explore sensitive and multidimensional data by issuing any range counting queries over the p-view before downstream data science workflows.**

data analysts may have little freedom to explore the data due to privacy issues [38, 40].

How can we explore the properties of high-dimensional sensitive data while preserving privacy? This paper focuses on guaranteeing differential privacy (DP) [15, 16] via random noise injections. As Figure 1 shows, we especially study how to construct a *privacy-preserving materialized view (p-view for short)* of relational data, which enables data analysts to explore arbitrary range counting queries in a differentially private way. Note that once a p-view is created, the privacy budget is not consumed any more for publishing counting queries, different from interactive differentially private query systems [18, 22, 23, 37, 41], which consume the budget every time queries are issued. In this work, we describe the desirable properties of the p-view, especially in data exploration for high-dimensional data, and fill the gaps of the existing methods.

Several methods for constructing a p-view have been studied in the existing literature. The most primitive method is to add Laplace noise [15] to each cell of the count tensor (or vector) representing the original histogram and publish the perturbed data as a p-view. While this noisy view can answer arbitrary range counting queries with a DP guarantee, it accumulates a large amount of noise. Data-aware partitioning methods [27, 29, 34, 42, 44, 46, 47] are potential solutions, but they focus only on low-dimensional data due to the high complexity of discovering the optimal partitioning when the data have multiple attributes. Additionally, these methods require exponentially large spaces as the dimensionality of the data increases due to the count tensor representation, which

**Table 1: Only the proposed method achieves all requirements in private data exploration for high-dimensional data. Each competitor represents a baseline [15], data partitioning [29, 42, 46], workload optimization [27, 30, 31], and generative model [17, 35, 45], respectively.**

	Identity <sup>1</sup> [15]	Privtree [46]	HDMM [31]	PrivBayes [45]	HDPVIEW (ours)
Workload independence	✓	✓		✓	✓
Analytical reliability	✓	✓	✓		✓
Noise resistance on high-dimensional		only low-dimensional	✓	✓	✓
Space efficiency			✓	✓	✓

can easily make them impractical. Workload-aware optimization methods [27, 29–31] are promising techniques for releasing query answers for high-dimensional data; however, they cannot provide query-independent p-views needed in data exploration.

In addition, one of the most popular approaches these days is differentially private learning of generative models [13, 17, 19, 20, 25, 35, 39, 45, 48]. Through the training of deep generative models [17, 19, 20, 25, 39] or graphical models [45, 48], counting queries and/or marginal queries can be answered directly from the model or indirectly with synthesized data via sampling. These methods are very space efficient because the synthetic dataset or graphical model can be used to answer arbitrary counting queries. However, these families rely on complex optimization methods such as DP-SGD [7], and it is very difficult to quantitatively estimate the error of counting queries using synthetic data, which eventually leads to a lack of reliability in practical use. Unlike datasets often used in the literature, the data collected in the practical field may be completely unmodelable. Table 1 summarizes a comparison between the most related works and our method, HDPVIEW, i.e., **High-Dimensional Private materialized View**. Each method in the table is described in more detail in Section 2.

Our target use case is privacy-preserving data exploration on high-dimensional data, for which the p-view should have the following four properties:

- **Workload independence:** Data analysts desire to issue arbitrary queries for exploring data. These queries should not be predefined.
- **Analytical reliability:** For reliability in practical data exploration, it is necessary to be able to estimate the scale of the error for arbitrary counting queries.
- **Noise resistance on high-dimensional data:** For high-dimensional data, range counting queries accumulate significant noise injected for DP. This noise accumulation makes the query answers useless. To avoid this issue, we need a robust mechanism for dimensionality increases.
- **Space efficiency:** It is necessary to generate spatially efficient views even for count tensors with a large number of total domains on various datasets.

**Our proposal.** To satisfy all the above requirements, we propose a simple yet effective recursive bisection method on a high-dimensional count tensor, HDPVIEW. Our proposed method has the same principle as [29, 42, 46] of first partitioning a database into

<sup>1</sup>Identity adds noise to the entries of the count vector by the Laplace mechanism [15] and cannot directly perturb high-dimensional datasets due to the domains being too large; we measure estimated error using the method described in [31].

small blocks and then averaging over each block with noise. Unlike the existing methods, HDPVIEW can efficiently perform error-minimizing partitioning even on multidimensional count tensors instead of conventional 1D count vectors. HDPVIEW recursively partitions multidimensional blocks at a cutting point chosen in a differentially private manner while aiming to minimize the sum of aggregation error (AE) and perturbation error (PE). Compared to Privtree [46], our proposed method provides a more data-aware flexible cutting strategy and proper convergence of block partitioning, which results in smaller errors in private counting queries and much better spatial efficiency of the generated p-views. Our method provides a powerful and practical solution for constructing a p-view under the DP constraint. More importantly, the p-view generated by HDPVIEW can work as a query processing system and expose the estimated error bound at runtime for any counting query without further privacy consumption. This error information ensures reliable analysis for data explorers.

**Contributions.** Our contributions are threefold. First, we design a *p-view* and formalize the segmentation for a multidimensional count tensor to find an effective p-view as error minimizing optimization problem. P-view can be widely used for data exploration process on multidimensional data and is a differentially private approximation of a multidimensional histogram that can release counting queries with analytical reliability. Second, we propose HD-PVIEW described above to find a desirable solution to the optimization problem. Our algorithm is more effective than conventional algorithms due to finding flexible partitions and more efficient due to making appropriate convergence decisions. Third, we conduct extensive experiments, whose source code and dataset are open, and show that HDPVIEW has the following merits. (1) Effectiveness: HDPVIEW demonstrates smaller errors for various range counting queries and outperforms the existing methods [15, 29, 31, 45, 46] on multi-dimensional real-world datasets. (2) Space efficiency: HD-PVIEW generates a much more compact representation of the p-view than the state-of-the-art (i.e., Privtree [46]) in our experiment.

**Preview of result.** We present a summary previewing of the experimental results. Table 2 shows the average relative root mean squared error against (RMSE) of HDPVIEW in eight types of range counting queries on eight real-world datasets and the average relative size of the p-view generated by the algorithms. With Identity, we obtain a p-view by making each cell of the original count tensor a converged block. HDPVIEW yields the smallest error score on average. This is a desirable property for data explorations. Furthermore, compared to that of Privtree [46], the p-view generated by HDPVIEW is more space efficient.

**Table 2: HDPVIEW provides low-error counting queries in average on various workloads and datasets, and high space-efficiency of privacy-preserving materialized view (p-view) when  $\epsilon = 1.0$ . (N/A is due to HDMM and PrivBayes do not create p-view.)**

	Identity <sup>1</sup> [15]	Privtree [46]	HDMM [31]	PrivBayes [45]	HDPVIEW (ours)
Average relative RMSE	$1.94 \times 10^7$	7.05	35.34	3.79	<b>1.00</b>
Average relative size of p-view	$4.59 \times 10^{17}$	5578.27	N/A	N/A	<b>1.00</b>

## 2 RELATED WORKS

In the last decade, several works have proposed differentially private methods for exploring sensitive data. Here, we describe the state-of-the-arts related to our work.

**Data-aware partitioning.** Data-aware partitioning is a conventional method that aims to directly randomize and expose the entire histogram for all domains (e.g., count vector, count tensor); thereby, it can immediately compose a p-view that answers all counting queries. A naïve approach to constructing a differentially private view is adding Laplace noise [15] to all values of a count vector; this is called the Identity mechanism. This naïve approach results in prohibitive noise on query answers through the accumulation of noise over the grouped bins used by queries. DAWA [29] and AHP [47] take data-aware partitioning approaches to reduce the amount of noise. The partitioning-based approaches first split a raw count vector into bins and then craft differentially private aggregates by averaging each bin and injecting a single unit of noise in each bin. However, these approaches work only for very low (e.g., one or two)-dimensional data due to the high complexity of discovering the optimal solution when the data have multiple attributes. DPCube [42] is a two-step multidimensional data-aware partitioning method, but the first step, obtaining an accurate approximate histogram, is difficult on high-dimensional data with small counts in each cell.

Privtree [46] and [14] perform multidimensional data-aware partitioning on count tensors, mainly targeting the spatial decomposition task for spatial data. Unlike our method, this method uses a fixed quadtree as the block partitioning strategy, which leads to an increase in unnecessary block partitioning as the dimensionality increases. As a result, it downgrades the spatial efficiency and incurs larger perturbation noise. In addition, this method aims to partition the blocks such that the count value is below a certain threshold, while our proposed method aims to minimize the AE of the blocks and reduce count query noise.

**Optimization of given workloads.** Another well-established approach is the optimization for a given workload. Li et al. [30] introduced a matrix mechanism (MM) that crafts queries and outputs optimized for a given workload. The high-dimensional MM (HDMM) [31] is a workload-aware data processing method extending the MM to be robust against noise for high-dimensional data. PrivateSQL [27] selects the view to optimize from pre-given workloads. In the data exploration process, it is not practical to assume a predefined workload, and these methods are characterized by a loss of accuracy when optimized for a workload of wide variety of queries.

**Private data synthesis.** Private data synthesis, which builds a privacy-preserving generative model of sensitive data and generates synthetic records from the model, is also useful for data exploration. Note that synthesized dataset can work as a p-view

by itself. PrivBayes [45] can heuristically learn a Bayesian network of data distribution in a differentially private manner. DPPro [43], Priview [35] and PrivSyn [48] represent distribution by approximation with several smaller marginal tables. While these methods provide a partial utility guarantee based on randomized mechanisms such as Laplace mechanisms or random projections, they face difficulties in providing an error bound for arbitrary counting queries on the synthesized data. Differentially private deep generative models have also attracted attention [9, 20, 25, 39], but most of the works focus on the reconstruction of image datasets. Fan et al. [17] studied how to build a good generative model based on generative adversarial nets (GAN) for tabular datasets. Their experimental results showed that the utility of differentially private GAN was lower than that of PrivBayes for tabular data. [36] provides a solution for high-dimensional data in a local DP setting.

As mentioned in Section 1, the accuracy of these methods has improved greatly in recent years, but it is difficult to guarantee their utility for analysis using counting queries, and there are large gaps in practice. DPPro [43] utilizes a random projection [24] that preserve L2-distance to the original data in an analyzable form to give a utility guarantee, but this is different from the guarantee for each counting query. CSM [49] gives a utility analysis for queries, however, their analysis ignores the effect of information loss due to compression, which may not be accurate. Also, as shown in their experiments, they apply intense preprocessing to the domain size and do not show the effectiveness for high-dimensional data. Our proposed method provides an end-to-end error analysis for arbitrary counting queries by directly constructing p-views from histograms without any intermediate generative model.

**Querying and programming framework.** PrivateSQL [27] is a differentially private relational database system for multirelational tables, where for each table, it applies an existing noise-reducing method such as DAWA. Unlike our method, PrivateSQL needs a given workload to design private views to release. Flex [23], Google’s work [41] and APEx [18] are SQL processing frameworks under DP. They issue queries to the raw data, which can consume an infinite amount of the privacy budgets. Hence, we believe that these DP-query processing engines are not suitable for data exploration tasks where many instances of trial and error may be possible. Our method generates p-view, which can be used as a differentially private query system that allows any number of range counting queries to be issued.

## 3 PRELIMINARIES

This section introduces essential knowledge for understanding our proposal. We first describe notations this paper uses. Then, we briefly explain DP.

### 3.1 Notation

Let  $X$  be the input database with  $n$  records consisting of an attribute set  $A$  that has  $d$  attributes  $A = \{a_1, \dots, a_d\}$ . The domain  $\text{dom}(a)$  of an attribute  $a$  has a finite ordered set of discrete values, and the size of the domain is denoted as  $|\text{dom}(a)|$ . The overall domain size of  $A$  is  $|\text{dom}(A)| = \prod_{i \in [d]} |\text{dom}(a_i)|$ , where  $[d] = \{1, \dots, d\}$ . In the case where attribute  $a$  is continuous, we transform the domain into a discrete domain by binning, and in the case where attribute  $a$  is categorical, we transform it into an ordered domain. Then,  $\text{dom}(a)$  can be represented as a range  $r[s_a, e_a]$  where for all  $p_a \in \text{dom}(a)$ ,  $s_a \leq p_a \leq e_a$ . For ranges  $r_1, r_2$ ,  $|r_1 \cap r_2|$  means the number of value  $p_a$  satisfies  $s_a \leq p_a \leq e_a$ .

We consider transforming the database  $X$  into the  $d$ -mode count tensor  $\mathcal{X}$ , where given  $d$  ranges  $r_1, \dots, r_d$ ,  $\mathcal{X}[r_1, \dots, r_d]$  represents the number of records where  $(a_1 \in r_1, \dots, a_d \in r_d) \in X$ . We utilize  $x \in \mathcal{X}$  as a count value in  $\mathcal{X}$ ; this corresponds to a cell of the count tensor. We denote a subtensor of  $\mathcal{X}$  as block  $\mathcal{B} \subseteq \mathcal{X}$ .  $\mathcal{B}$  is also a  $d$ -mode count tensor, but its domain in each dimension is smaller than or equal to that of the original count tensor  $\mathcal{X}$ ; i.e., each attribute  $a_i$  ( $i \in [d]$ ),  $r[s_{a_i}, e_{a_i}]$  of  $\mathcal{B}$  and  $r[s'_{a_i}, e'_{a_i}]$  of  $\mathcal{X}$  satisfy  $s'_{a_i} \leq s_{a_i}$  and  $e_{a_i} \leq e'_{a_i}$ . We denote the domain size of  $\mathcal{B}$  as  $|\mathcal{B}|$ .

Last, we denote  $q$  as a counting query and  $\mathbf{W}$  as a workload.  $\mathbf{W}$  is a set of  $|\mathbf{W}|$  counting queries, where  $\mathbf{W} = \{q_1, \dots, q_{|\mathbf{W}|}\}$ , and  $q(\mathcal{X})$  returns the counting query results for count tensor  $\mathcal{X}$ .

### 3.2 Differential Privacy

DP [15] is a rigorous mathematical privacy definition that quantitatively evaluates the degree of privacy protection when we publish outputs. DP is used in broad domains and applications [11, 12, 33]. The importance of DP is supported by the fact that the US census announced '2020 Census results will be protected using "differential privacy", the new gold standard in data privacy protection' [8].

**DEFINITION 1 ( $\epsilon$ -DIFFERENTIAL PRIVACY).** A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{Z}$  satisfies  $\epsilon$ -DP if, for any two inputs  $D, D' \in \mathcal{D}$  such that  $D'$  differs from  $D$  in at most one record and any subset of outputs  $Z \subseteq \mathcal{Z}$ , it holds that

$$\Pr[\mathcal{M}(D) \in Z] \leq \exp(\epsilon) \Pr[\mathcal{M}(D') \in Z].$$

We define databases  $D$  and  $D'$  as *neighboring* databases.

Practically, we employ a randomized mechanism  $\mathcal{M}$  that ensures DP for a function  $f$ . The mechanism  $\mathcal{M}$  perturbs the output of  $f$  to cover  $f$ 's sensitivity, which is the maximum degree of change over any pair of datasets  $D$  and  $D'$ .

**DEFINITION 2 (SENSITIVITY).** The sensitivity of a function  $f$  for any two neighboring inputs  $D, D' \in \mathcal{D}$  is:

$$\Delta_f = \sup_{D, D' \in \mathcal{D}} \|f(D) - f(D')\|.$$

where  $\|\cdot\|$  is a norm function defined in  $f$ 's output domain.

When  $f$  is a histogram,  $\Delta_f$  equals 1 [21]. Based on the sensitivity of  $f$ , we design the degree of noise to ensure DP. The Laplace mechanism and exponential mechanism are well-known as standard approaches. The Laplace mechanism can be used for randomizing numerical data. Releasing a differentially private histogram is a typical use case of this mechanism.

**DEFINITION 3 (LAPLACE MECHANISM).** For function  $f : \mathcal{D} \rightarrow \mathbb{R}^n$ , the Laplace mechanism adds noise  $f(D)$  as:

$$f(D) + \text{Lap}(\Delta_f/\epsilon)^n. \quad (1)$$

where  $\text{Lap}(\lambda)^n$  denotes a vector of  $n$  independent samples from a Laplace distribution  $\text{Lap}(\lambda)$  with mean 0 and scale  $\lambda$ .

The exponential mechanism is the random selection algorithm. The selection probability is weighted based on a score in a quality metric for each item.

**DEFINITION 4 (EXPONENTIAL MECHANISM).** Let  $q$  be the quality metric for choosing an item  $y \in Y$  in the database  $D$ . The exponential mechanism randomly samples  $y$  from  $Y$  with weighted sampling probability defined as follows:

$$\Pr[y] \sim \exp\left(\frac{\epsilon q(D, y)}{2\Delta_q}\right). \quad (2)$$

Quantifying the privacy of differentially private mechanisms is essential for releasing multiple outputs. Sequential composition and parallel composition are standard privacy accounting methods.

**THEOREM 1 (SEQUENTIAL COMPOSITION [15]).** Let  $\mathcal{M}_1, \dots, \mathcal{M}_k$  be mechanisms satisfying  $\epsilon_1, \dots, \epsilon_k$ -DP. Then, a mechanism sequentially applying  $\mathcal{M}_1, \dots, \mathcal{M}_k$  satisfies  $(\sum_{i \in [k]} \epsilon_i)$ -DP.

**THEOREM 2 (PARALLEL COMPOSITION [32]).** Let  $\mathcal{M}_1, \dots, \mathcal{M}_k$  be mechanisms satisfying  $\epsilon_1, \dots, \epsilon_k$ -DP. Then, a mechanism applying  $\mathcal{M}_1, \dots, \mathcal{M}_k$  to the disjoint datasets  $D_1, \dots, D_k$  in parallel satisfies  $(\max_{i \in [k]} \epsilon_i)$ -DP.

## 4 PROBLEM FORMULATION

### 4.1 Segmentation as Optimization

This section describes the foundation of multidimensional data-aware segmentation that seeks a solution for the differentially private view  $\tilde{\mathcal{X}}$  from the input count tensor  $\mathcal{X}$ . Every count  $\tilde{x} \in \tilde{\mathcal{X}}$  is sanitized to satisfy DP. We formulate multidimensional block segmentation as an optimization problem.


**Foundation.** Given a count tensor  $\mathcal{X}$ , we consider partitioning  $\mathcal{X}$  into  $m$  blocks  $\pi = \{\mathcal{B}_1, \dots, \mathcal{B}_m\}$ . The blocks satisfy  $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$  where  $i, j \in [m]$ ,  $j \neq i$  and  $\mathcal{B}_1 \cup \dots \cup \mathcal{B}_m = \mathcal{X}$ . We denote the sum over  $\mathcal{B}_i$  as  $S_i = \sum_{x' \in \mathcal{B}_i} x'$  and its perturbed output as  $\tilde{S}_i = S_i + z_i$ . We can sample  $z_i$  with the Laplace mechanism  $\text{Lap}(1/\epsilon)$  and craft the  $\epsilon$ -differentially private sum in  $\mathcal{B}_i$ .

For any count  $x$  in the block  $\mathcal{B}_i$ , we have two types of errors: *Perturbation Error* (PE) and *Aggregation Error* (AE). Assuming that we replace any count  $x \in \mathcal{B}_i$  with  $\tilde{x}_i = (S_i + z_i)/|\mathcal{B}_i|$ , the absolute error between  $x$  and  $\tilde{x}_i$  can be computed as

$$|x - \tilde{x}_i| = \left| \left( x - \frac{S_i}{|\mathcal{B}_i|} \right) - \frac{z_i}{|\mathcal{B}_i|} \right| \leq \left| x - \frac{S_i}{|\mathcal{B}_i|} \right| + \left| \frac{z_i}{|\mathcal{B}_i|} \right|. \quad (3)$$

Therefore, the total error over block  $\mathcal{B}_i$ , namely, the segmentation error (SE), can be given by:

$$\text{SE}(\mathcal{B}_i) = \sum_{x \in \mathcal{B}_i} |x - \tilde{x}_i| \leq \text{AE}(\mathcal{B}_i) + \text{PE}(\mathcal{B}_i) \quad (4)$$



**Figure 2: HDPVIEW efficiently discovers blocks (i.e., groups of count cells) with smaller AEs (black arrow) and averages over each block with injected noise (red arrow). The p-view stores the randomized counts in a blockwise way.**

where

$$\text{AE}(\mathcal{B}_i) := \sum_{x \in \mathcal{B}_i} \left| x - \frac{S_i}{|\mathcal{B}_i|} \right|, \quad (5)$$

$$\text{PE}(\mathcal{B}_i) := |z_i|. \quad (6)$$

(5) and (6) represent the AE and the PE, respectively.

**Problem.** The partitioning makes the PE of each block  $\frac{1}{|\mathcal{B}_i|}$  times smaller than those of the original counts with Laplace noise. Furthermore, we consider the expectation of the SE

$$\begin{aligned} \mathbb{E} \left[ \sum_{i \in [m]} \text{SE}(\mathcal{B}_i) \right] &\leq \mathbb{E} \left[ \sum_{i \in [m]} \text{AE}(\mathcal{B}_i) \right] + \mathbb{E} \left[ \sum_{i \in [m]} \text{PE}(\mathcal{B}_i) \right] \\ &= \sum_{i \in [m]} \text{AE}(\mathcal{B}_i) + \sum_{i \in [m]} \mathbb{E}[\text{PE}(\mathcal{B}_i)] \\ &= \sum_{i \in [m]} \text{AE}(\mathcal{B}_i) + m \cdot \frac{1}{\epsilon}. \end{aligned} \quad (7)$$

Thus, to discover the optimal partition  $\pi$ , we need to minimize Eq. (7). The optimization problem is denoted as follows:

$$\begin{aligned} &\underset{\pi}{\text{minimize}} \quad \sum_{\mathcal{B}_i \in \pi} \left( \text{AE}(\mathcal{B}_i) + \frac{1}{\epsilon} \right) \\ &\text{subject to} \quad \mathcal{B}_i \cap \mathcal{B}_{j \neq i} = \emptyset, \quad \mathcal{B}_i, \mathcal{B}_j \in \pi \\ &\quad \bigcup_{\mathcal{B}_i \in \pi} \mathcal{B}_i = \mathcal{X} \end{aligned} \quad (8)$$

**Challenges.** It is not easy to discover the optimal partition  $\pi$ . This problem is an instance of the set partitioning problem [10], which is known to be NP-complete, where the objective function is computed by brute-force searching for every combination of candidate blocks. It is hard to solve since the search space is basically a very large scale due to large  $|\text{dom}(A)|$ . Therefore, this paper seeks an efficient heuristic solution with a good balance between utility (i.e., smaller errors) and privacy.

## 4.2 P-view Definition

Our proposed p-view has a simple structure. The p-view consists of a set of blocks, each of which has a range for each attribute and an appropriately randomized count value, as shown in Figure 1. Formally, we define the p-view as follows.

$$\begin{aligned} \text{p-view } \tilde{\mathcal{X}} &= \{\mathcal{B}_1, \dots, \mathcal{B}_m\}, \\ \text{for } i \in [m], \mathcal{B}_i &= (\{r[s_{a_1}^{(i)}, e_{a_1}^{(i)}], \dots, r[s_{a_d}^{(i)}, e_{a_d}^{(i)}]\}, \tilde{S}_i) \end{aligned} \quad (9)$$

Thus, each block  $\mathcal{B}_i$  has this  $d$ -dimensional domain and the sanitized sum of count values  $\tilde{S}_i$ .

In the range counting query processing, a counting query  $q$  needs to have the range condition  $c_q = \{r[s_{a_1}^{(q)}, e_{a_1}^{(q)}], \dots, r[s_{a_d}^{(q)}, e_{a_d}^{(q)}]\}$ . Let the ranges of  $\mathcal{B}_i$  be  $\{r[s_{a_1}^{(i)}, e_{a_1}^{(i)}], \dots, r[s_{a_d}^{(i)}, e_{a_d}^{(i)}]\}$ , and we calculate the intersection of  $c_q$  and the block and add the count value according to the size of the intersection. Hence, the result can be calculated as follows.

$$q(\tilde{\mathcal{X}}) = \sum_{i=1, \dots, m} \left( \prod_{l=1, \dots, d} \left( |r[s_{a_l}^{(q)}, e_{a_l}^{(q)}] \cap r[s_{a_l}^{(i)}, e_{a_l}^{(i)}]| \right) * \frac{\tilde{S}_i}{|\mathcal{B}_i|} \right) \quad (10)$$

The number of intersection calculations is proportional to the number of blocks, and the complexity of the query processing is  $O(md)$ .

## 5 PROPOSED ALGORITHM

This section introduces our proposed solution. Our solution constructs a p-view of the input relational data while preserving utility and privacy with analytical reliability to estimate errors in the arbitrary counting queries against the p-view (Eq.(10)).

### 5.1 Overview

Our challenge is to devise a simple yet effective algorithm that enables us to efficiently search a block partitioning with small total errors and DP guarantees. As a realization of the algorithm, we propose HDPVIEW.

Figure 2 illustrates an overview of our proposed algorithm. First, HDPVIEW creates the initial block  $\mathcal{B}^{(0)}$  that covers the whole count tensor  $\mathcal{X}$ . Second, we recursively bisect a block  $\mathcal{B}$  (initially  $\mathcal{B} = \mathcal{B}^{(0)}$ ) into two disjoint blocks  $\mathcal{B}_L$  and  $\mathcal{B}_R$ . Before bisecting  $\mathcal{B}$ , we check whether the AE over  $\mathcal{B}$  is sufficiently small. If the result of the check is positive, we stop the recursive bisection for  $\mathcal{B}$ . Otherwise, we continue to split  $\mathcal{B}$ . We pick a splitting point  $p \in \text{dom}(a)$  ( $a \in A$ ) for splitting  $\mathcal{B}$  into  $\mathcal{B}_L$  and  $\mathcal{B}_R$  which have smaller AEs. Although splitting does not always result in smaller total AEs, proper cut point obviously makes AEs much smaller. Third, HDPVIEW recursively executes these steps separately for  $\mathcal{B}_L$  and  $\mathcal{B}_R$ . After convergence is met for all blocks, HDPVIEW generates a randomized aggregate by  $S_i + z_i$  where  $z_i \sim \text{Lap}(1/\epsilon)$  for each block  $\mathcal{B}_i$ . Finally, for all  $x \in \mathcal{B}_i$ , we obtain the randomized count  $\tilde{x} = (S_i + z_i)/|\mathcal{B}_i|$ .

The abovementioned algorithm can discover blocks that heuristically reduce the AEs, and is efficient due to its simplicity. However, the question is *how can we make the above algorithm differentially private?* To solve this question, we introduce two mechanisms, *random converge* (Section 5.2) and *random cut* (Section 5.3). Random converge determines the convergence of the recursive bisection, and random cut determines the effective single cutting point. These provide reasonable partitioning strategy to reduce the total errors with small privacy budget consumption.

The overall algorithm of HDPVIEW are described in Algorithm 1. Let  $\epsilon_b = \epsilon_r + \epsilon_p$  be the total privacy budget for HDPVIEW, where  $\epsilon_r$  is the budget for the recursive bisection and  $\epsilon_p$  is the budget for the perturbation. HDPVIEW utilizes  $\gamma\epsilon_r$  for random converge and  $(1 - \gamma)\epsilon_r$  for random cut ( $0 \leq \gamma \leq 1$ ).  $\alpha$  is a hyperparameter that determines the size of  $\lambda$  and  $\delta$ , where  $\lambda$  corresponds to the Laplace noise scale of random converge (Lines 8, 16) and  $\delta$  is a bias term for AE (Lines 9, 15). These are, sketchily, tricks for performing random converge with depth-independent scales, which are explained in Section 5.2 and a detailed proof of DP is given in Section 5.4. The algorithm runs recursively (Lines 10, 28, 29), alternating between random converge (Lines 14-18) and random cut (Lines 19-26). The random converge stops when the AE becomes small enough, consuming a total budget of  $\gamma\epsilon_r$  independent of the number of depth. The random cut consumes a budget of  $(1 - \gamma)\epsilon_r/\kappa$  for each cutting point selection until the depth exceeds  $\kappa$ .  $\kappa$  is set as  $\kappa = \beta \log_2 \bar{n}$ , where  $\beta > 0$  is hyperparameter and  $\bar{n}$  is the total domain size of the data. As we see later in Theorem 4, AE is not increased by splitting, so if the depth is greater than  $\kappa$ , we split randomly without any privacy consumption until convergence. After the recursive bisection converges, HDPVIEW perturbs the count by adding the Laplace noise while consuming  $\epsilon_p$ .

## 5.2 Random Converge

AE decreases by properly splitting the blocks, however unnecessary block splitting leads to an increase in PE as mentioned above. To stop the recursive bisection at the appropriate depth, we need to obtain the exact AE of the block, which is a data-dependent output, therefore we need to provide a DP guarantee. One approach is to publish differential private AE so that making the decision for the stop is also DP by the post-processing property. In other words, the stop is determined by  $AE(\mathcal{B}) + Lap(\lambda) \leq \theta$  where  $\theta$  is a threshold indicating AE is small enough. However, this method consumes privacy budget every time the AE is published, and the budget cannot be allocated unless the depth of the partition is decided in advance. Therefore, we utilize the observation for the *privacy loss* of Laplace mechanism-based threshold query [46] and design the biased AE (BAE) of the block  $\mathcal{B}$  instead of  $AE(\mathcal{B})$ :  $BAE(\mathcal{B}) = \max(\theta + 2 - \delta, AE(\mathcal{B}) - k\delta)$ , where  $k$  is the current depth of bisection,  $\delta (> 0)$  is a bias parameter, i.e., we determine the convergence by  $BAE(\mathcal{B}) + Lap(\lambda) \leq \theta$ . Intuitively, the BAE is designed to tightly bound the privacy loss of the any number of Laplace mechanism-based threshold queries with constant noise scale  $\lambda$ . When the value is sufficiently larger than the threshold, this privacy loss decreases exponentially [46]. Then, it can be easily bounded by an infinite series regardless of the number of queries. Conversely, when the value is small compared to the threshold, each threshold query consumes a constant budget. To limit the number of such budget consumptions, a bias  $\delta$  is used to force a decrease in the value for each threshold query (i.e., each depth) because BAE has a minimum and if the value is guaranteed to be less than the minimum for adjacent databases, the privacy loss is zero. The design of our BAE allows for two constant budget consumptions at most, with the remainder being bounded by an infinite series. We give a detailed proof in Section 5.4. As a whole, since BAE is basically close to AE, AEs are expected to become sufficiently small overall.

---

### Algorithm 1 HDPVIEW

---

**Input:** initial block  $\mathcal{B}^{(0)}$ , privacy budget  $\epsilon_b$ , recursive bisection budget ratio  $\epsilon_r/\epsilon_b$ , hyperparameters  $\alpha, \beta, \gamma$

**Output:** p-view  $\tilde{\mathcal{X}}$

```

1: procedure HDPVIEW( $\mathcal{B}^{(0)}, \epsilon_b, \epsilon_r/\epsilon_b, \alpha, \beta, \gamma$ )
2:    $\epsilon_r \leftarrow \epsilon_b \cdot (\epsilon_r/\epsilon_b)$ ;  $\epsilon_p \leftarrow \epsilon_b \cdot (1 - \epsilon_r/\epsilon_b)$ 
3:    $\bar{n} \leftarrow \text{TOTALDOMAINSIZEOF}(\mathcal{X})$ 
4:    $\kappa \leftarrow \beta \log_2 \bar{n}$  // maximum depth of random cut
5:    $\pi \leftarrow \{\}$ ;  $k \leftarrow 1$  // converged blocks; current depth
6:    $\theta \leftarrow 1/\epsilon_p$  // threshold
7:    $\epsilon_{cut} \leftarrow (1 - \gamma)\epsilon_r/\kappa$  // privacy budget for random cut
8:    $\lambda \leftarrow \left(\frac{2\alpha-1}{\alpha-1} + 1\right) \cdot \left(\frac{2}{\gamma\epsilon_r}\right)$  // noise scale for random converge
9:    $\delta \leftarrow \lambda \log \alpha$  // bias parameter
10:  RECURSIVEBISECTION( $\mathcal{B}^{(0)}, \pi, \epsilon_{cut}, k, \kappa, \theta, \lambda, \delta$ )
11:   $\tilde{\mathcal{X}} \leftarrow \text{PERTURBATION}(\pi, \epsilon_p)$ 
12:  return  $\tilde{\mathcal{X}}$ 
13: procedure RECURSIVEBISECTION( $\mathcal{B}, \pi, \epsilon_{cut}, k, \kappa, \theta, \lambda, \delta$ )
14:  /* Random Converge */
15:   $BAE(\mathcal{B}) \leftarrow \max(\theta + 2 - \delta, AE(\mathcal{B}) - k\delta)$ 
16:  if  $BAE(\mathcal{B}) + Lap(\lambda) \leq \theta$  then
17:     $\pi \leftarrow \pi \cup \mathcal{B}$ 
18:    return
19:  /* Random Cut */
20:  if  $k \leq \kappa$  then
21:    for all  $i \in [d], j \in [|\text{dom}(a_i)|]$  do
22:       $quality[i, j] \leftarrow Q(\mathcal{B}, a_{ij})$ 
23:       $(i^*, j^*) \leftarrow \text{WEIGHTEDSAMPLING}(\epsilon_{cut}, quality)$ 
24:    else
25:       $(i^*, j^*) \leftarrow \text{RANDOMSAMPLING}([d], [|\text{dom}(a_i)|])$ 
26:       $(\mathcal{B}_L, \mathcal{B}_R) \leftarrow \text{SPLIT}(i^*, j^*)$ 
27:    /* Repeat Recursively */
28:    RECURSIVEBISECTION( $\mathcal{B}_L, \pi, \epsilon_{cut}, k + 1, \kappa, \theta, \lambda, \delta$ )
29:    RECURSIVEBISECTION( $\mathcal{B}_R, \pi, \epsilon_{cut}, k + 1, \kappa, \theta, \lambda, \delta$ )
30:  return

```

---

Then, we consider about  $\theta$  where if  $\theta$  is too large, block partitioning will not sufficiently proceed, causing large AEs, and if it is too small, more blocks will be generated, leading to increase in total PEs. To prevent unwanted splitting, it is appropriate to stop when the increase in PE is greater than the decrease in AE. We design the threshold  $\theta$  as  $1/\epsilon_p$  which is the standard deviation of the Laplace noise to be perturbed. Considering the each bisection increases the total PE by  $1/\epsilon_p$ , when the AE becomes less than the PE, the division will increase the error at least. Hence, it is reasonable to stop under this condition.

## 5.3 Random Cut

Here, the primary question is how to pick a reasonable cutting point from all attribute values in a block  $\mathcal{B}$  under DP. Our intuition is that a good cutting point results in smaller AEs in the two split blocks. We design random cut by combining an exponential mechanism with scoring based on the total AE after splitting.

Let  $\mathcal{B}_L^{(p)}$  and  $\mathcal{B}_R^{(p)}$  be the blocks split from  $\mathcal{B}$  by the cutting point  $p$ , and the quality function  $Q$  of  $p$  in  $\mathcal{B}$  is defined as follows:

$$Q(\mathcal{B}, p) = -(\text{AE}(\mathcal{B}_L^{(p)}) + \text{AE}(\mathcal{B}_R^{(p)})). \quad (11)$$

Then, we compute the score for all attribute values  $p \in \text{dom}(a)$ ,  $a \in A$ , and satisfies  $|\mathcal{B}_L^{(p)}| \geq 1$  and  $|\mathcal{B}_R^{(p)}| \geq 1$ . Note that the number of candidates for  $p$  is proportional to the sum of the domains for each attribute  $\sum_{i \in [d]} |\text{dom}(a_i)|$ , not to the total domains  $\prod_{i \in [d]} |\text{dom}(a_i)|$ . We employ weighted sampling via an exponential mechanism to choose one cutting point  $p^*$ . The sampling probability of  $p$  is proportional to

$$\Pr[p^* = p] \sim \exp\left(\frac{\epsilon Q(\mathcal{B}, p)}{2\Delta_Q}\right) \quad (12)$$

where  $\Delta_Q$  is the L1-sensitivity of the quality metric  $Q$ . We denote the L1-sensitivity of AE as  $\Delta_{AE}$ , and we can easily find  $\Delta_Q = 2\Delta_{AE}$  because  $Q$  is the sum of two AEs. Thus, each time a cut point  $p$  is published according to such weighted sampling, a privacy budget of  $\epsilon$  is consumed. We set  $\epsilon$  as the budget allocated to random cut (i.e.,  $(1 - \gamma)\epsilon_r$ ) divided by  $\kappa$ . If the cutting depth exceeds  $\kappa$ , we switch to random sampling (Line 25 in Algorithm 1). Hence, cutting will not stop regardless of the depth or budget.

Compared to Privtree [46], for a  $d$ -dimensional block, at each cut, HDPVIEW generates just 2 blocks with this random cut while Privtree generates  $2^d$  blocks with fixed cutting points. Privtree's heuristics prioritizes finer partitioning, which sufficiently works in low-dimensional data because AEs become very small and the total PEs is not so large. In high-dimensional data, however, it causes unnecessary block splitting resulting in too much PEs. HDPVIEW carefully splits blocks one by one, thus suppressing unnecessary block partitioning and reducing the number of blocks i.e., smaller PEs. It also enables flexibly shaped multidimensional block partitioning. Moreover, while whole design of HDPVIEW including convergence decision logic and cutting strategy are based on an error optimization problem as described in Section 4, Privtree has no such background. This allows HDPVIEW to provide effective block partitioning rather than simply fewer blocks, which we empirically confirm in Section 6.2.

## 5.4 Privacy Accounting

For privacy consumption accounting, since HDPVIEW recursively splits a block into two disjoint blocks, we only have to trace a path toward convergence. In other words, because HDPVIEW manipulates all the blocks separately, we can track the total privacy consumption by the parallel composition for each converged block. The information published by the recursive bisection is the result of segmentation; however, note that since there is a constraint on the cutting method for the block, it must be divided into two parts; in the worst case, the published blocks may expose all the cutting points. For a given converged block  $\mathcal{B}$ , we denote the series of cutting points by  $S_{\mathcal{B}} = [p_1, \dots, p_k]$ , and  $\mathcal{B}_{p_i}$  as the block after being divided into two parts at cutting point  $p_i$ . To show the DP guarantee, let  $D$  and  $D'$  be the neighboring databases, and let  $\Pr[S_{\mathcal{B}}|D]$  be the probability that  $S_{\mathcal{B}}$  is generated from  $D$ . We need to show that for any  $D, D'$ , and  $S_{\mathcal{B}}$  that

$$\left| \frac{\Pr[S_{\mathcal{B}}|D]}{\Pr[S_{\mathcal{B}}|D']} \right| \leq e^\epsilon \quad (13)$$

to show that the recursive bisection satisfies  $\epsilon$ -DP.

The block with the largest  $\left| \frac{\Pr[S_{\mathcal{B}}|D]}{\Pr[S_{\mathcal{B}}|D']} \right|$  of the converged disjoint blocks is  $\mathcal{B}^*$ , which has the longest  $S_{\mathcal{B}^*}$  and contains different data between  $D$  and  $D'$ . Random converge and random cut are represented as follows.

$$\begin{aligned} \left| \frac{\Pr[S_{\mathcal{B}^*}|D]}{\Pr[S_{\mathcal{B}^*}|D']} \right| &= \frac{\Pr[\text{BAE}(\mathcal{B}_{p_0}) + \text{Lap}(\lambda) > \theta]}{\Pr[\text{BAE}(\mathcal{B}'_{p_0}) + \text{Lap}(\lambda) > \theta]} \\ &\cdot \frac{\Pr[p^* = p_1|D]}{\Pr[p^* = p_1|D']} \cdot \frac{\Pr[\text{BAE}(\mathcal{B}_{p_1}) + \text{Lap}(\lambda) > \theta]}{\Pr[\text{BAE}(\mathcal{B}'_{p_1}) + \text{Lap}(\lambda) > \theta]} \\ &\cdots \frac{\Pr[p^* = p_k|D]}{\Pr[p^* = p_k|D']} \cdot \frac{\Pr[\text{BAE}(\mathcal{B}_{p_k}) + \text{Lap}(\lambda) \leq \theta]}{\Pr[\text{BAE}(\mathcal{B}'_{p_k}) + \text{Lap}(\lambda) \leq \theta]} \end{aligned} \quad (14)$$

where  $\mathcal{B}_{p_0}$  is the initial count tensor and for all  $i$ ,  $\mathcal{B}'_{p_i}$  indicates a neighboring block for  $\mathcal{B}_{p_i}$ . Taking the logarithm,

$$\begin{aligned} \ln\left(\frac{\Pr[S_{\mathcal{B}^*}|D]}{\Pr[S_{\mathcal{B}^*}|D']}\right) &= \sum_{i=1}^k \underbrace{\ln\left(\frac{\Pr[p^* = p_i|D]}{\Pr[p^* = p_i|D']}\right)}_{(*1) : \text{for random cut}} \\ &+ \underbrace{\sum_{i=0}^k \ln\left(\frac{\Pr[\text{BAE}(\mathcal{B}_{p_i}) + \text{Lap}(\lambda) > \theta]}{\Pr[\text{BAE}(\mathcal{B}'_{p_i}) + \text{Lap}(\lambda) > \theta]}\right) + \ln\left(\frac{\Pr[\text{BAE}(\mathcal{B}_{p_k}) + \text{Lap}(\lambda) \leq \theta]}{\Pr[\text{BAE}(\mathcal{B}'_{p_k}) + \text{Lap}(\lambda) \leq \theta]}\right)}_{(*2) : \text{for random converge}} \end{aligned} \quad (15)$$

and let the first item of the right-hand of Eq.(15) be (\*1), and the other items be (\*2).

(\*1) corresponds to the privacy of the random cut, with each probability following Eq.(12). Given  $\epsilon = \epsilon_{cut}$ , for any  $k$ , the following holds from sequential composition.

$$\left| \sum_{i=1}^k \ln\left(\frac{\Pr[p^* = p_i|D]}{\Pr[p^* = p_i|D']}\right) \right| \leq \kappa \epsilon_{cut} = (1 - \gamma)\epsilon_r. \quad (16)$$

The following are privacy guarantees for the other part, (\*2), based on the observations presented in [46]. First, we consider the sensitivity of AE  $\Delta_{AE}$ .

**THEOREM 3.** *The L1-sensitivity of the AE is  $2(1 - 1/|\mathcal{B}|)$ .*

**PROOF.** Let  $\mathcal{B}'$  be the block that differs by only one count from  $\mathcal{B}$ . The AE( $\mathcal{B}'$ ) can be computed as follows:

$$\text{AE}(\mathcal{B}') = \sum_{i \neq j \in [|\mathcal{B}|]} \left| x_i - \frac{S+1}{|\mathcal{B}|} \right| + \left| x_j + 1 - \frac{S+1}{|\mathcal{B}|} \right|.$$

Finally, the L1-sensitivity of AE can be derived as:

$$\Delta_{AE} = (|\mathcal{B}| - 1) \frac{1}{|\mathcal{B}|} + 1 - \frac{1}{|\mathcal{B}|} = 2(1 - 1/n) \quad \square$$

Thus, we also obtain  $|\text{BAE}(\mathcal{B}) - \text{BAE}(\mathcal{B}')| \leq 2$ , and

$$\begin{aligned} (*2) &\leq \sum_{i=0}^{k-1} \ln\left(\frac{\Pr[\text{BAE}(\mathcal{B}_{p_i}) + \text{Lap}(\lambda) > \theta]}{\Pr[\text{BAE}(\mathcal{B}_{p_i}) - 2 + \text{Lap}(\lambda) > \theta]}\right) \\ &+ \ln\left(\frac{\Pr[\text{BAE}(\mathcal{B}_{p_k}) + \text{Lap}(\lambda) \leq \theta]}{\Pr[\text{BAE}(\mathcal{B}_{p_k}) + 2 + \text{Lap}(\lambda) \leq \theta]}\right). \end{aligned} \quad (17)$$

Furthermore, from the proof in the Appendix in [46], when we have  $f(x) = \ln\left(\frac{\Pr[x+Lap(\lambda) > \theta]}{\Pr[\bar{x}-2+Lap(\lambda) > \theta]}\right)$ , then

$$\begin{cases} f(x) \leq \frac{2}{\lambda}, & (\theta - x + 2 > 0) \\ f(x) \leq \frac{2}{\lambda} \exp\left(\frac{\theta - x + 2}{\lambda}\right), & (\theta - x + 2 \leq 0) \end{cases} \quad (18)$$

Next, we show the monotonic decreasing property of AE for block partitioning.

**THEOREM 4.** *For any  $i = 0, \dots, k-1$ ,  $AE(\mathcal{B}_{p_i}) \geq AE(\mathcal{B}_{p_{i+1}})$ .*

**PROOF.** We show that when  $\mathcal{B}^+$  is an arbitrary block  $\mathcal{B}$  with an arbitrary element  $x (> 0)$  added to it, the AEs always satisfy  $AE(\mathcal{B}) \leq AE(\mathcal{B}^+)$ . Let the elements in  $\mathcal{B}$  be  $x_1, \dots, x_k$ , and let  $\mathcal{B}^+$  be the block with  $x_{k+1}$  added. The mean values in each block are  $\bar{x} = \frac{1}{k}(x_1 + \dots + x_k)$  and  $\bar{x}^+ = \frac{1}{k+1}(x_1 + \dots + x_{k+1})$  and  $AE(\mathcal{B}) = \sum_{i=1}^k |x_i - \bar{x}|$  and  $AE(\mathcal{B}^+) = \sum_{i=1}^{k+1} |x_i - \bar{x}^+|$ . Considering how much the AE can be reduced with the addition of  $x_{k+1}$  to  $\mathcal{B}$ ,  $|x_i - \bar{x}| - |x_i - \bar{x}^+| \leq |\bar{x} - \bar{x}^+|$  holds for each  $i (= 1, \dots, k)$ , so  $AE(\mathcal{B}) - AE(\mathcal{B}^+)$  is at most  $k \cdot |\bar{x} - \bar{x}^+|$ . On the other hand, with the addition of  $x_{k+1}$ , AE increases by at least  $|x_{k+1} - \bar{x}^+|$  because this is a new item. Since  $x_{k+1} = (k+1)\bar{x}^+ - (x_1 + \dots + x_k) = (k+1)\bar{x}^+ - k\bar{x}$ , then  $|x_{k+1} - \bar{x}^+| = |k \cdot (\bar{x} - \bar{x}^+)| = k \cdot |\bar{x} - \bar{x}^+|$ . Hence,  $AE(\mathcal{B}^+) - AE(\mathcal{B}) \geq k \cdot |\bar{x} - \bar{x}^+| - k \cdot |\bar{x} - \bar{x}^+| = 0$  always holds. Therefore, since  $\mathcal{B}_{p_i}$  always has more elements than  $\mathcal{B}_{p_{i+1}}$ ,  $AE(\mathcal{B}_{p_i}) \geq AE(\mathcal{B}_{p_{i+1}})$ .  $\square$

Considering  $BAE(\mathcal{B})$ , there exists a natural number  $m (1 \leq m \leq k)$  where if  $i < m$ ,  $BAE(\mathcal{B}_{p_i}) \geq BAE(\mathcal{B}_{p_{i+1}}) + \delta \geq \theta + 2 - \delta$ , if  $i = m$ ,  $\theta + 2 \geq BAE(\mathcal{B}_{p_i}) \geq \theta + 2 - \delta$ , and if  $m < i$ ,  $BAE(\mathcal{B}_{p_i}) = \theta + 2 - \delta$ . Therefore, using Eqs.(17, 18),

$$\begin{aligned} (*) &\leq \frac{2}{\lambda} + \sum_{i=1}^{m-1} \frac{2}{\lambda} \exp\left(\frac{\theta - BAE(\mathcal{B}_{p_i}) + 2}{\lambda}\right) + \frac{2}{\lambda} \\ &\leq \frac{4}{\lambda} + \frac{2}{\lambda} \cdot \frac{1}{1 - \exp\left(-\frac{\delta}{\lambda}\right)} \\ &= \frac{2}{\lambda} \cdot \frac{3 \exp\left(\frac{\delta}{\lambda}\right) - 2}{\exp\left(\frac{\delta}{\lambda}\right) - 1}. \end{aligned} \quad (19)$$

Thus, to make  $(*)$  satisfy  $\gamma \epsilon_r$ -DP,  $\frac{2}{\lambda} \cdot \frac{3 \exp\left(\frac{\delta}{\lambda}\right) - 2}{\exp\left(\frac{\delta}{\lambda}\right) - 1} \leq \gamma \epsilon_r$  should hold.

Since the  $\lambda$  and  $\delta$  that satisfy these conditions are not uniquely determined, these values are determined by giving  $\exp(\delta/\lambda)$  as a hyperparameter  $\alpha$ . Then, we can always calculate  $\lambda = \left(\frac{3\alpha-2}{\alpha-1}\right) \cdot \left(\frac{2}{\gamma\epsilon_r}\right)$  and  $\delta = \lambda \log \alpha$ , in turn, which satisfies  $(*) \leq \gamma \epsilon_r$ .  $\alpha$  is valid for  $\alpha > 1$ . If  $\alpha$  is extremely close to 1,  $\lambda$  diverges and *random convergence* is too inaccurate. As  $\alpha$  increases,  $\lambda$  decreases, but  $\delta$  increases. Thus  $\lambda$  and  $\delta$  are trade-offs, and independently of the dataset, there exists a point at which both values are reasonably small. Around  $\alpha = 1.4 \sim 1.8$  works well empirically. We show a specific analytical result in Appendix of the full version [26].

Finally, together with  $(*)$ , the recursive bisection by random converge and random cut satisfies  $\epsilon_r$ -DP. In addition, the perturbation consumes  $\epsilon_p$  for each block to add Laplace noise, so together with this, HDPVIEW satisfies  $\epsilon_p + \epsilon_r = \epsilon_b$ -DP.

## 5.5 Error Analysis

When a p-view created by HDPVIEW publishes a counting query answer, we can dynamically estimate an upper bound distribution of the error included in the noisy answer. The upper bound of the error can be computed from the number of blocks used to answer the query and the distribution of the perturbation. Note that this can be computed without consuming any extra privacy budget because, as shown in 5.4, in addition to the count values, block partitioning results are released in a DP manner.

As a count query on the p-view is processed as Eq. (10), the answer consists of the sum of the query results for each block, and from Eq. (4), each block contains two types of errors: AE and PE. Let the error of a counting query  $q$  be  $Error(q, \mathcal{X}, \mathcal{X}') := \|q(\mathcal{X}) - q(\mathcal{X}')\|_1$  where  $\mathcal{X}$  and  $\mathcal{X}'$  are the original and noisy data, respectively, and we define the error by the L1-norm. First, since the AE depends on the concrete count values of each block involved in each query condition, we characterise the block distribution by defining an  $\xi$ -uniformly scattered block.

**DEFINITION 5 ( $\xi$ -UNIFORMLY SCATTERED).** *A block  $\mathcal{B}$  is  $\xi$ -uniformly scattered if for any subblock  $\mathcal{B}' \subset \mathcal{B}$ ,*

$$AE(\mathcal{B}')/|\mathcal{B}'| \leq \xi \cdot AE(\mathcal{B})/|\mathcal{B}|. \quad (20)$$

While  $\xi$  depends on the actual data, it is expected to decrease with each step by random cut.

Then, we have the following theorem for the error.

**THEOREM 5.** *If for all  $i$ , block  $\mathcal{B}_i$  is  $\xi_i$ -uniformly scattered, any  $\mu$  satisfying  $0 < \mu < 1$ , and any  $t$  satisfying  $|t| < \epsilon_p$  and  $|t| < \frac{1}{\lambda}$ , the error of a counting query satisfies  $Error(q, \mathcal{X}, \mathcal{X}') \geq \Theta_{min}(\mu)$  and  $Error(q, \mathcal{X}, \mathcal{X}') \leq \Theta_{max}(\mu)$  with probability of at least  $1 - \mu$ , respectively, with*

$$\begin{aligned} \Theta_{min}(\mu) &= \frac{1}{t} \left( \log \mu + \sum_{i=1, \dots, m} \log \left( 1 - \left( \frac{w_i}{\epsilon_p} \right)^2 t^2 \right) \right) \\ \Theta_{max}(\mu) &= \sum_{i=1, \dots, m} \xi_i w_i (k_i \delta + \theta) \\ &\quad - \frac{1}{t} \left( \log \mu + \sum_{i=1, \dots, m} \log \left( 1 - \left( \frac{w_i}{\epsilon_p} \right)^2 t^2 \right) + \log \left( 1 - (\xi_i w_i \lambda)^2 t^2 \right) \right) \end{aligned} \quad (21)$$

where  $w_i = \frac{|\mathcal{B}_i \cap c_q|}{|\mathcal{B}_i|}$ ,  $k_i$  is depth of  $\mathcal{B}_i$  that can be public information.

**PROOF.** The errors included in  $Error(q, \mathcal{X}, \mathcal{X}')$  are PEs and AEs. Both of them follow independent probability distributions for each block, and we first show the PE. For each  $\mathcal{B}_i$ , perturbation noise is uniformly divided inside  $\mathcal{B}_i$ . Hence, the total PE in the query  $q$  is represented by  $\sum_{i=1, \dots, m} w_i * PE(\mathcal{B}_i)$  where  $w_i = \frac{|\mathcal{B}_i \cap c_q|}{|\mathcal{B}_i|}$  and  $PE(\mathcal{B}_i)$  is Laplace random variable following  $Lap\left(\frac{1}{\epsilon_p}\right)$ .

Then, we consider the AE. From random converge, given a  $\mathcal{B}_i$ , then  $BAE(\mathcal{B}_i) + Lap(\lambda) \leq \theta$  holds. Considering  $BAE(\mathcal{B}) = \max(\theta + 2 - \delta, AE(\mathcal{B}) - k\delta)$ , when  $\theta + 2 - \delta \leq AE(\mathcal{B}_i) - k_i\delta$ ,

$$AE(\mathcal{B}_i) \leq Lap(\lambda) + k_i\delta + \theta. \quad (22)$$

And when  $\theta + 2 - \delta > AE(\mathcal{B}_i) - k_i\delta$ ,

$$AE(\mathcal{B}_i) - k_i\delta < \theta + 2 - \delta \leq Lap(\lambda) + \theta \quad (23)$$



Thus, the upper bound of  $AE(\mathcal{B}_i)$  is distributed under  $Lap(\lambda) + k_i\delta + \theta$ . In other words, AE cannot be observed directly, but the upper bound distribution is bounded by the Laplace distribution. Also note that the AE satisfies  $AE(\mathcal{B}_i) \geq 0$ .

Therefore, for the error lower bound, we only need to consider the  $m$  PEs,  $\sum_{i=1,\dots,m} w_i * PE(\mathcal{B}_i)$ .  $PE(\mathcal{B}_i)$  is independent random variable, respectively. We apply Chernoff bound to the sum, for any  $a$  and  $t$ ,

$$\Pr[Error(q, \mathcal{X}, \mathcal{X}') \leq a] \leq e^{(ta)} \prod_{i=1,\dots,m} E[e^{-tw_i PE(\mathcal{B}_i)}], \quad (24)$$

where  $|t| < \epsilon_p$  is required for existence of the moment generating function. By using  $PE(\mathcal{B}_i)$  follows  $Lap(\frac{1}{\epsilon_p})$ , we can derive

$$\Pr\left[Error(q, \mathcal{X}, \mathcal{X}') \leq \frac{1}{t} \left( \log \mu + \sum_{i=1,\dots,m} \log \left( 1 - \left( \frac{w_i}{\epsilon_p} \right)^2 t^2 \right) \right)\right] \leq \mu. \quad (25)$$

On the other hand, for the error upper bound, we need to consider AEs as well. Hence, we apply Chernoff bound to the sum of  $2m$  independent random variables following each Laplace distribution. Considering the upper bound distribution of  $AE(\mathcal{B}_i)$  has  $w_i(k_i\delta + \theta)$  for the mean and  $\lambda$  for the variance, let  $\bar{AE}(\mathcal{B}_i)$  be a Laplace random variable whose mean and variance are 0 and  $\lambda$ , respectively, then we have

$$\begin{aligned} & \Pr\left[Error(q, \mathcal{X}, \mathcal{X}') - \sum_{i=1,\dots,m} \xi_i w_i (k_i \delta + \theta) \geq a\right] \\ & \leq e^{(-ta)} \prod_{i=1,\dots,m} E[e^{(tw_i PE(\mathcal{B}_i))}] E[e^{(tw_i \xi_i \bar{AE}(\mathcal{B}_i))}], \end{aligned} \quad (26)$$

where since block  $\mathcal{B}_i$  is  $\xi_i$ -uniformly scattered, AE included in the query  $q$  and block  $\mathcal{B}_i$  is at most  $\xi_i w_i AE(\mathcal{B}_i)$ . Lastly, for any  $t$ , where  $|t| < \epsilon_p$  and  $|t| < \frac{1}{\lambda}$ , from the inequality, we can derive as follows:

$$\begin{aligned} & \Pr\left[Error(q, \mathcal{X}, \mathcal{X}') \geq \sum_{i=1,\dots,m} \xi_i w_i (k_i \delta + \theta) \right. \\ & \quad \left. - \frac{1}{t} \left( \log \mu + \sum_{i=1,\dots,m} \log \left( 1 - \left( \frac{w_i}{\epsilon_p} \right)^2 t^2 \right) + \log \left( 1 - (\xi_i w_i \lambda)^2 t^2 \right) \right) \right] \\ & \leq \mu. \end{aligned}$$

This completes the proof.  $\square$

Importantly, this can be dynamically computed for any counting queries, helping the analyst to perform a reliable exploration.

Similarly, since the HDMM [31] optimizes budget allocations for counting queries by the MM, we can statically calculate the error distributions for each query. However, this is workload-dependent. In data exploration, we consider predefined workload is strong assumption to be avoided.

## 6 EVALUATION

In this section, we report the results of the experimental evaluation of HDPVIEW. To evaluate our proposed method, we design experiments to answer the following questions:

- How effectively can the constructed p-views be used in data exploration via various range counting queries?
- How space-efficiently can the constructed p-views represent high-dimensional count tensors?

**Table 3: Datasets.**

Dataset	#Record	#Column (categorical)	#Domain	Variance
Adult [1]	48842	15 (9)	$9 \times 10^{19}$	0.0360
Small-adult	48842	4 (2)	$3 \times 10^5$	0.0237
Numerical-adult	48842	7 (1)	$2 \times 10^{11}$	0.0200
Traffic [5]	48204	8 (2)	$1 \times 10^{14}$	0.0484
Bitcoin [6]	500000	9 (1)	$4 \times 10^{12}$	0.0379
Electricity [2]	45312	8 (1)	$1 \times 10^{14}$	0.0407
Phoneme [4]	5404	6 (1)	$2 \times 10^6$	0.0304
Jm1 [3]	10885	22 (1)	$2 \times 10^{21}$	0.0027

We show the effectiveness of HDPVIEW via range counting queries in section 6.2, and section 6.3 reports the space efficiency.

### 6.1 Experimental Setup

We describe the experimental setups. In the following experiments, we run 10 trials with HDPVIEW and the competitors and report their averages to remove bias. Throughout the experiments, the hyperparameters of HDPVIEW are fixed as  $(\epsilon_r/\epsilon_b, \alpha, \beta, \gamma) = (0.9, 1.6, 1.2, 0.9)$ , and we basically use  $\epsilon = 1.0$  as a privacy budget. We provide observations and insights into all the hyperparameters of HDPVIEW in Appendix of the full version [26].

**Datasets.** We use several multidimensional datasets commonly used in the literature, as shown in Table 3. Adult [1] includes 6 numerical and 9 categorical attributes. We prepare Small-adult by extracting 4 attributes (age, workclass, race, and capital-gain) from Adult. Additionally, we form Numerical-adult by extracting only numerical attributes and a label. Traffic [5] is a traffic volume dataset. Bitcoin [6] is a Bitcoin transaction graph dataset. Electricity [2] is a dataset on changes in electricity prices. Phoneme [4] is a dataset for distinguishing between nasal and oral sounds. Jm1 [3] is a dataset of static source code analysis data for detecting defects with 22 attributes. HDPVIEW and most competitors require the binning of all numerical attribute values for each dataset. Basically, we set the number of bins to 100 or 10 when the attribute is a real number. We consider that the number of bins should be determined by the level of granularity that analysts want to explore, regardless of the distribution of the data. For categorical columns, we simply apply ordinal encoding. In Table 3, #Domain shows the total domain sizes after binning. Variance is the mean of the variance for each dimension of the binned and normalized dataset and gives an indication of how scattered the data is.

**Implementations of competitors.** We compare our proposed method HDPVIEW with Identity [15], Privtree [46], HDMM [31], PrivBayes [45], and DAWA partitioning mechanism [29]. For these methods, we perform the following pre- and postprocessing steps. For Identity, we estimate errors following [31], employing implicit matrix representations and workload-based estimation, because it is infeasible to add noises on a high-dimensional count tensor because of the huge space. For Privtree, as described in [46], we set the threshold to 0 and allocate half of the privacy budget to tree construction and half to perturbation. Using the same method as HDPVIEW, the blocks obtained by Privtree are used as the p-view. For the HDMM, we utilize p-Identity strategies as a template.

DAWA’s partitioning mechanism can be applied to multidimensional data by flattening data into a 1D vector. However, when the domain size becomes large, the optimization algorithm based on the v-optimal histogram for the count vector cannot be applied due to the computational complexity. Therefore, we apply DAWA to `Small-adult` and `Phoneme` because their domain sizes are relatively small. We perform only DAWA partitioning without workload optimization to compare the partitioning capability without a given workload to evaluate workload-independent p-view generation. For fairness, `PrivBayes` is trained on raw data<sup>8</sup>. `PrivBayes`, in counting queries, samples the exact number of original data points; therefore, it may consume extra privacy budget.

**Workloads.** We prepare different types of workloads. *k-way All Marginal* is all marginal counting queries using all combinations of *k* attributes. *k-way All Range* is the range version of the marginal queries. *Prefix-kD* consists of prefix queries using all combinations of *k* attributes. *Random-kD Range Query* is range counting queries for arbitrary *k* attributes and we randomly generate 3000 queries for a single workload.

**Reproducibility.** The experimental code is publicly available on the <https://github.com/FumiyukiKato/HDPVIEW>.

## 6.2 Effectiveness

We evaluate how effective p-views constructed by HDPVIEW are in data exploration by issuing various range counting queries.

**Evaluation metrics.** We evaluate HDPVIEW and other mechanisms by measuring the RMSE for all counting queries. Formally, given the count tensor  $\mathcal{X}$ , randomized view  $\mathcal{X}'$  and workload  $\mathbf{W}$ , the RMSE is defined as:  $\text{RMSE} = \sqrt{\frac{1}{|\mathbf{W}|} \sum_{q \in \mathbf{W}} (q(\mathcal{X}) - q(\mathcal{X}'))^2}$ . This metric is useful for showing the utility of the p-view. It corresponds to the objective function optimized by MM families [30, 31], where given a workload matrix  $\mathbf{W}$  and a query strategy  $\mathbf{A}$ , which is the optimized query set to answer the workload, the expected error of the workloads is  $\frac{2}{\epsilon^2} \|\mathbf{A}\|_1^2 \|\mathbf{W}\mathbf{A}^+\|_F^2 = \text{RMSE}^2$ . Thus, we can compare the measured errors with this optimized estimated errors. We also report the relative RMSE against HDPVIEW for comparison.

**High quality on average.** Figure 3 shows the relative RMSEs for all datasets and workloads and algorithms with privacy budget  $\epsilon=1.0$ . The relative RMSE (log-scale) is plotted on the vertical axis and the dataset on the horizontal axis where high-dimensional datasets (`Jm1` and `Adult`) are on the left, medium-dimensional datasets (`Traffic`, `Electricity`, `Bitcoin` and `Numerical-adult`) are in the middle, and low-dimensional datasets (`Small-adult` and `Phoneme`) are on the right. The errors with `Identity` for high-dimensional data are too large and are omitted for appearance. Overall, HDPVIEW works well. In Section 1, Table 2 shows the relative RMSE averaged over all workloads and all datasets in Figure 3, and HDPVIEW achieves the lowest error on average. In data exploration, we want to run a variety of queries, so the average accuracy is important. We believe HDPVIEW has such a desirable property. A detailed comparisons with the competitors are explained in the following paragraphs.

**Comparison with Identity, HDMM and DAWA.** `Identity`, which is the most basic Baseline, and `HDMM`, which performs workload optimization, cause more errors for high-dimensional datasets than

HDPVIEW. For `Identity`, the reason is that the accumulation of noise increases as the number of domains increases. HDPVIEW avoids the noise accumulation by grouping domains into large blocks. The results of `HDMM` show that the increasing dimension of the dataset and the dimension of the query can increase the error. This is because the matrix representing the counting queries to which the matrix mechanism is applied becomes complicated, making it hard to find efficient budget allocations. This is why the accuracy of the 3- or 4-dimensional queries for `Jm1` and `Adult` is poor with `HDMM`. In particular, the `HDMM`’s sensitivity to dimensionality increases can also be seen in Figure 7. DAWA’s partitioning leads more errors than the HDPVIEW and `Privtree`. When applied to multi-dimensional data, DAWA finds the optimal partitioning on a domain mapped in one-dimension, while HDPVIEW and `Privtree` finds more effective multi-dimensional data partitioning.

**Comparison with Privtree.** As a whole, HDPVIEW outperforms `Privtree`’s accuracy mainly for mid- to high-dimensional datasets. In particular, we can see `Privtree`’s performance drops drastically in high-dimensionality (i.e., `Jm1`). `Privtree` achieves higher accuracy than HDPVIEW for `Phoneme`. This is likely because `Privtree`’s strategy, which prioritize finer splitting, are sufficient for the small domain size rather than HDPVIEW’s heuristic algorithm. Even if the blocks is too fine, the accumulation of PEs is not so large in low-dimensionality, and AEs become smaller, which results in an accurate p-view. The reason why HDPVIEW is better for `Small-adult` despite the low-dimensionality may be that the sizes of the cardinality of attributes are uneven (`Small-adult`: {74, 9, 5, 100}, `Phoneme`: {10, 10, 10, 10, 10, 10, 2}), which may make `Privtree`’s fixed cutting strategy ineffective. To see the very low-dimensional case, Figure 4 shows the block partitioning for the 2D data with a popular `Gowalla` [28] check-in dataset. The table below shows the number of blocks and the RMSE for the 3000 `Random 2D` range query. HDPVIEW yields fewer blocks and `Privtree` generates a less noisy p-view for the abovementioned reason. The figure also confirms that HDPVIEW performs a flexible shape of block partitioning.

On the other hand, for high-dimensional dataset, this property can be avenged. In `Privtree`, a single cutting always generates  $2^d$  new blocks, which are too fine, resulting in very large PEs even though the AEs are smaller. Figure 5 shows the distribution of AEs for blocks on `Adult` for HDPVIEW and `Privtree`. HDPVIEW has slightly larger AE blocks, but `Privtree` has a large number of blocks and cause larger PEs. An extreme case is `Jm1` in which `Privtree` causes large errors. This is probably because `Jm1` actually requires fewer blocks since the distribution is highly concentrated (c.f., Table 3). Figure 8 shows that the number of blocks of generated p-view by HDPVIEW and `Privtree`. For `Jm1`, HDPVIEW generates very small number of blocks while `Privtree` does not. We can confirm that HDPVIEW avoids unnecessary splitting via *random cut* and suppresses the increase in PEs which causes in `Privtree`. This would be noticeable for datasets with concentrated distributions, where the required number of blocks is essentially small.

Figure 6 shows the results of reducing the number of cut attributes in `Privtree` and adjusting the number of blocks in p-view on `Numerical-adult` and `Jm1`. If the number of cut attributes is smaller than the dimension  $d$ , we choose target attributes in a round-robin way (Appendix of [46]). In the case of `Numerical-adult`, the

<sup>8</sup>`PrivBayes` shows worse performances with binned data in our prestudy.

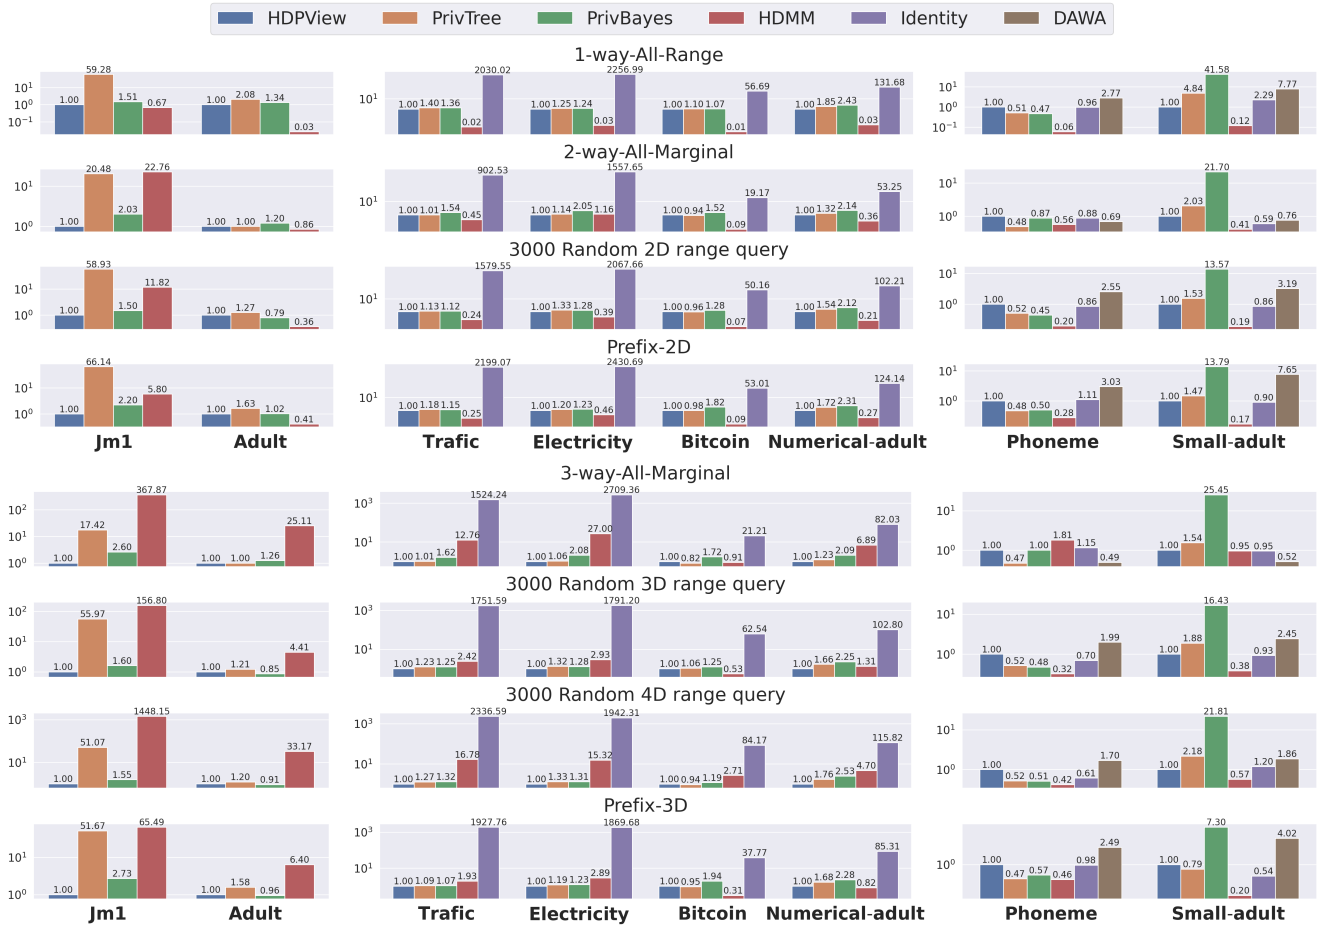


Figure 3: Relative RMSEs against HDPVIEW on all the datasets and workloads: HDPVIEW shows small errors for a wide variety of high-dimensional range counting queries.

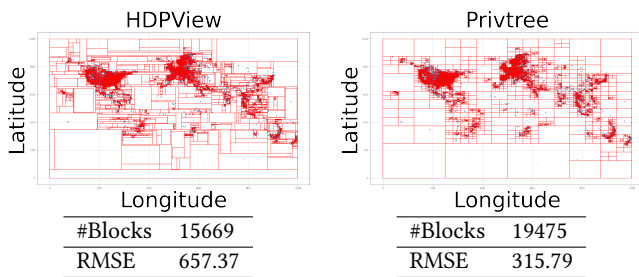


Figure 4: Examples of HDPVIEW (left) and Privtree (right) on 2D dataset (Gowalla): HDPVIEW has fewer blocks, leading to noisier results than Privtree for very low-dimensional data. Also, HDPVIEW provides flexible block partitioning.

error basically decreases as the number of cut attributes is increased, similar to the observation in Appendix of [46]. However, for high-dimensional data such as Jm1, the error increases rapidly as the number of cut attributes increases to some extent. This is consistent with the earlier observation that influence of PEs increases. Also,

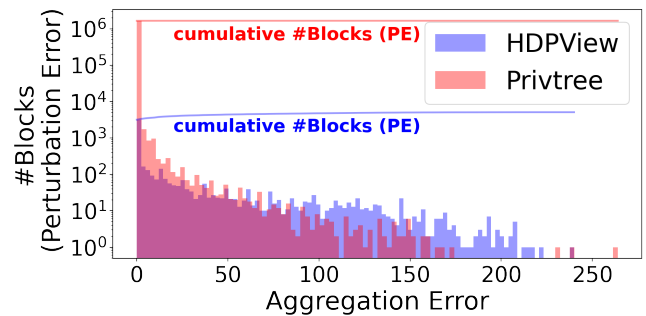


Figure 5: Number of blocks (log-scale) with various AEs for high-dimensional dataset (i.e., Adult) for HDPVIEW and Privtree. HDPVIEW has slightly larger AE blocks, but Privtree has a much more number of blocks i.e., much larger PEs.

in any cases, the error of HDPVIEW is smaller, indicating that HDPVIEW not only has a smaller number of blocks, but also performs effective block partitioning compared to Privtree on these datasets.




Figure 6: HDPVIEW is more effective than Privtree even with controlled number of cuttings on Numerical-adult (left) and Jm1 (log-scale) (right).

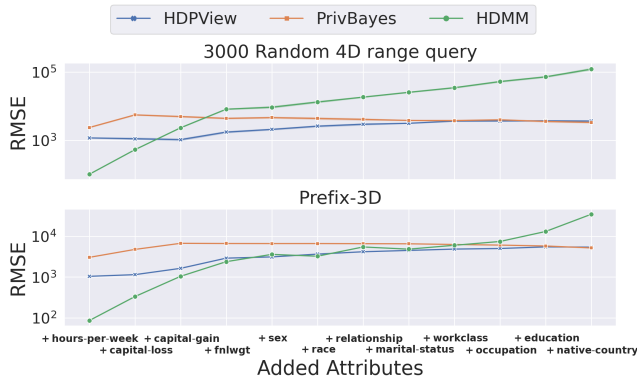


Figure 7: Changes in the performance when adding attributes to Adult one by one in HDPVIEW, PrivBayes, and HDMM.

**Comparison with PrivBayes.** We do not consider PrivBayes a direct competitor because it is a generative model approach that does not provide any analytical reliability as described in Section 2. However, PrivBayes is a state-of-the-art specialized for publishing differentially private marginal queries; therefore, we compared the accuracy to demonstrate the performance of HDPVIEW. As shown in Figure 3, HDPVIEW is a little more accurate than PrivBayes in many cases. However, in Adult, PrivBayes slightly outperforms HDPVIEW. Because PrivBayes uses Bayesian network to learn the data distribution, it can fit well even to high-dimensional data as long as the distribution of the data is easily modelable. In HDPVIEW, with larger dimensionality, the PEs grow slightly because the total number of blocks increases. The AEs also grow since more times of random converge result in larger errors. Thus, the total error is at least expected to increase, and the larger dimensionality may work to the advantage of PrivBayes. Still, HDPVIEW is advantageous, especially for concentrated data such as Jm1.

We consider the reason why on Numerical-adult, which has a smaller dimensionality than Adult, PrivBayes is less accurate than HDPVIEW is because the effective attributes for capturing the accurate marginal distributions with Bayesian network are removed. We can see the same behavior for Small-adult. The following experimental results can support this. Figure 7 describes the changes in the RMSE with attributes added to Adult one by one in two workloads, where the added attributes are shown on the horizontal axis. Initially, HDPVIEW is more accurate than PrivBayes. As




Figure 8: The number of blocks generated by HDPVIEW is much lower than that generated by Privtree.

Table 4: HDPVIEW’s p-view is space efficient (up to  $10^{13}\times$ ).

Dataset	Identity-based	HDPVIEW
Adult	30.99 EB	3.61 MB
Bitcoin	1.27 TB	6.77 MB
Electricity	1.11 TB	2.19 MB
Phoneme	781.34 KB	273.59 KB

attributes are added, HDPVIEW is basically robust with increasing dimensionality, but the error increases slightly. On the other hand, interestingly, the error in PrivBayes becomes slightly smaller.

Lastly, considering HDPVIEW is better in Numerical-adult and worse in Adult, one of the advantages of PrivBayes may be due to the increase in categorical attributes. Because HDPVIEW bisects the ordered domain space, it may be hard to effectively divide categorical attributes, which possibly worsens the accuracy in HDPVIEW.

### 6.3 Space Efficiency

Our proposed p-view stores each block in a single record. This method avoids redundancy in recording all cells that belong to the same block. The p-view consists of blocks and values, and basically, the space complexity follows the number of blocks. Figure 8 shows a comparison between the numbers of blocks of HDPVIEW and Privtree. While the accuracy of the counting queries of HDPVIEW is higher than that of Privtree, the number of blocks generated by HDPVIEW is much lower than that of Privtree, indicating that the strategy of HDPVIEW avoids unnecessary splitting. In particular, on Jm1, HDPVIEW is  $4 \times 10^4$  more efficient than Privtree. Table 4 shows the size of the randomized views, Identity-based noisy count vector (not p-view) and p-view generated by HDPVIEW at  $\epsilon=1.0$ . Since HDPVIEW constructs the p-view by a compact representation, it results in up to  $10^{13}$  times smaller space on Adult.

## 7 CONCLUSION

We addressed the following research question: How can we construct a privacy-preserving materialized view to explore the unknown properties of the high-dimensional sensitive data? To practically construct the p-view, we proposed a data-aware segmentation method, HDPVIEW. In our experiments, we confirmed the following desirable properties, (1) Effectiveness: HDPVIEW demonstrated smaller errors for various range counting queries in multidimensional queries. (2) Space efficiency: HDPVIEW generates a compact representation of the p-view. We believe that our method helps us explore sensitive data in the early stages of data mining while preserving data utility and privacy.

## REFERENCES

- [1] 1996. Adult Data Set - UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml/datasets/Adult>. accessed on 2022-05-10.
- [2] 2014. electricity - OpenML. <https://www.openml.org/d/151>. accessed on 2022-05-10.
- [3] 2014. jm1 - OpenML. <https://www.openml.org/d/1053>. accessed on 2022-05-10.
- [4] 2015. phoneme - OpenML. <https://www.openml.org/d/1489>. accessed on 2022-05-10.
- [5] 2019. Metro Interstate Traffic Volume Data Set - UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml/datasets/Metro+Interstate+Traffic+Volume>. accessed on 2022-05-10.
- [6] 2020. Bitcoin Data Set - UCI Machine Learning Repository. <https://archive.ics.uci.edu/ml/datasets/BitcoinHeistRansomwareAddressDataset>. accessed on 2022-05-10.
- [7] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 308–318.
- [8] John M Abowd. 2018. The us census bureau adopts differential privacy. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2867–2867.
- [9] Gergely Acs, Luca Melis, Claude Castelluccia, and Emiliano De Cristofaro. 2018. Differentially private mixture of generative neural networks. *IEEE Transactions on Knowledge and Data Engineering* 31, 6 (2018), 1109–1121.
- [10] Egon Balas and Manfred W Padberg. 1976. Set partitioning: A survey. *SIAM review* 18, 4 (1976), 710–760.
- [11] Vincent Bindschaedler, Reza Shokri, and Carl A Gunter. 2017. Plausible deniability for privacy-preserving data synthesis. *Proceedings of the VLDB Endowment* 10, 5 (2017), 481–492.
- [12] Kamalika Chaudhuri, Jacob Imola, and Ashwin Machanavajjhala. 2019. Capacity bounded differential privacy. In *Advances in Neural Information Processing Systems*. 3469–3478.
- [13] Rui Chen, Qian Xiao, Yu Zhang, and Jianliang Xu. 2015. Differentially private high-dimensional data publication via sampling-based inference. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 129–138.
- [14] Graham Cormode, Cecilia Procopiuc, Divesh Srivastava, Entong Shen, and Ting Yu. 2012. Differentially private spatial decompositions. In *2012 IEEE 28th International Conference on Data Engineering*. IEEE, 20–31.
- [15] Cynthia Dwork. 2006. Differential privacy. In *Proceedings of the 33rd international conference on Automata, Languages and Programming-Volume Part II*. Springer-Verlag, 1–12.
- [16] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3-4 (2014), 211–407.
- [17] Ju Fan, Junyou Chen, Tongyu Liu, Yuwei Shen, Guoliang Li, and Xiaoyong Du. 2020. Relational Data Synthesis Using Generative Adversarial Networks: A Design Space Exploration. 13, 12 (2020).
- [18] Chang Ge, Xi He, Ihab F Ilyas, and Ashwin Machanavajjhala. 2019. Apex: Accuracy-aware differentially private data exploration. In *Proceedings of the 2019 International Conference on Management of Data*. 177–194.
- [19] Chang Ge, Shubhankar Mohapatra, Xi He, and Ihab F Ilyas. 2020. Kamino: Constraint-Aware Differentially Private Data Synthesis. *arXiv preprint arXiv:2012.15713* (2020).
- [20] Frederik Harder, Kamil Adamczewski, and Mijung Park. 2021. DP-MERF: Differentially Private Mean Embeddings with RandomFeatures for Practical Privacy-preserving Data Generation. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 1819–1827.
- [21] Michael Hay, Vibhor Rastogi, Gerome Miklau, and Dan Suciu. 2009. Boosting the accuracy of differentially-private histograms through consistency. *arXiv preprint arXiv:0904.0942* (2009).
- [22] Noah Johnson, Joseph P Near, Joseph M Hellerstein, and Dawn Song. 2018. Chorus: Differential privacy via query rewriting. *arXiv preprint arXiv:1809.07750* (2018).
- [23] Noah Johnson, Joseph P Near, and Dawn Song. 2018. Towards practical differential privacy for SQL queries. *Proceedings of the VLDB Endowment* 11, 5 (2018), 526–539.
- [24] William B Johnson and Joram Lindenstrauss. 1984. Extensions of Lipschitz mappings into a Hilbert space 26. *Contemporary mathematics* 26 (1984).
- [25] James Jordon, Jinsung Yoon, and Mihaela van der Schaar. 2018. PATE-GAN: generating synthetic data with differential privacy guarantees. In *International Conference on Learning Representations*.
- [26] Fumiyuki Kato, Tsubasa Takahashi, Shun Takagi, Yang Cao, Seng Pei Liew, and Masatoshi Yoshikawa. 2022. HDPView: Differentially Private Materialized View for Exploring High Dimensional Relational Data. *arXiv preprint arXiv:2203.06791* (2022).
- [27] Ios Kotsogiannis, Yuchao Tao, Xi He, Maryam Fanaeepour, Ashwin Machanavajjhala, Michael Hay, and Gerome Miklau. 2019. PrivateSQL: a differentially private SQL query engine. *Proceedings of the VLDB Endowment* 12, 11 (2019), 1371–1384.
- [28] Jure Leskovec. 2011. Gowalla Dataset. <http://snap.stanford.edu/data/loc-Gowalla.html>. accessed on 2022-05-10.
- [29] Chao Li, Michael Hay, Gerome Miklau, and Yue Wang. 2014. A Data- and Workload-Aware Algorithm for Range Queries under Differential Privacy. 7, 5 (2014).
- [30] Chao Li, Gerome Miklau, Michael Hay, Andrew McGregor, and Vibhor Rastogi. 2015. The matrix mechanism: optimizing linear counting queries under differential privacy. *The VLDB journal* 24, 6 (2015), 757–781.
- [31] Ryan McKenna, Gerome Miklau, Michael Hay, and Ashwin Machanavajjhala. 2018. Optimizing Error of High-Dimensional Statistical Queries under Differential Privacy. 11, 10 (2018).
- [32] Frank D McSherry. 2009. Privacy integrated queries: an extensible platform for privacy-preserving data analysis. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of data*. 19–30.
- [33] Nicolas Papernot, Martin Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. 2016. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755* (2016).
- [34] Wahbeh Qardaji, Weining Yang, and Ninghui Li. 2013. Differentially private grids for geospatial data. In *2013 IEEE 29th international conference on data engineering (ICDE)*. IEEE, 757–768.
- [35] Wahbeh Qardaji, Weining Yang, and Ninghui Li. 2014. Privview: practical differentially private release of marginal contingency tables. In *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. 1435–1446.
- [36] Xuebin Ren, Chia-Mu Yu, Weiren Yu, Shusen Yang, Xinyu Yang, Julie A McCann, and S Yu Philip. 2018. LoPub: High-Dimensional Crowdsourced Data Publication with Local Differential Privacy. *IEEE Transactions on Information Forensics and Security* 13, 9 (2018), 2151–2166.
- [37] Ryan Rogers, Subbu Subramaniam, Sean Peng, David Durfee, Seunghyun Lee, Santosh Kumar Kancha, Shraddha Sahay, and Parvez Ahammad. 2020. LinkedIn’s Audience Engagements API: A privacy preserving data analytics system at scale. *arXiv preprint arXiv:2002.05839* (2020).
- [38] Du Su, Hieu Tri Huynh, Ziao Chen, Yi Lu, and Wenmiao Lu. 2020. Re-identification Attack to Privacy-Preserving Data Analysis with Noisy Sample-Mean. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1045–1053.
- [39] Shun Takagi, Tsubasa Takahashi, Yang Cao, and Masatoshi Yoshikawa. 2021. P3GM: Private high-dimensional data release via privacy preserving phased generative model. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*. IEEE, 169–180.
- [40] Zhen Wang, Xiang Yue, Soheil Moosavinasab, Yungui Huang, Simon Lin, and Huan Sun. 2019. Surfcon: Synonym discovery on privacy-aware clinical data. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1578–1586.
- [41] Royce J Wilson, Celia Yuxin Zhang, William Lam, Damien Desfontaines, Daniel Simmons-Marengo, and Bryant Gipson. 2019. Differentially private sql with bounded user contribution. *arXiv preprint arXiv:1909.01917* (2019).
- [42] Yonghui Xiao, Li Xiong, Liyue Fan, and Slawomir Goryczka. 2012. DPCube: Differentially private histogram release through multidimensional partitioning. *arXiv preprint arXiv:1202.5358* (2012).
- [43] Chugui Xu, Ju Ren, Yaoxue Zhang, Zhan Qin, and Kui Ren. 2017. DPPro: Differentially Private High-Dimensional Data Release via Random Projection. *IEEE Transactions on Information Forensics and Security* 12, 12 (2017), 3081–3093.
- [44] Grigory Yaroslavtsev, Graham Cormode, Cecilia M Procopiuc, and Divesh Srivastava. 2013. Accurate and efficient private release of datacubes and contingency tables. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*. IEEE, 745–756.
- [45] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. 2017. Privbays: Private data release via bayesian networks. *ACM Transactions on Database Systems (TODS)* 42, 4 (2017), 25.
- [46] Jun Zhang, Xiaokui Xiao, and Xing Xie. 2016. Privtree: A differentially private algorithm for hierarchical decompositions. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 155–170.
- [47] 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, Philadelphia, Pennsylvania, USA, April 24-26, 2014*, Mohammed Javeed Zaki, Zoran Obradovic, Pang-Ning Tan, Arindam Banerjee, Chandrika Kamath, and Srinivasan Parthasarathy (Eds.). SIAM, 587–595.
- [48] Zhikun Zhang, Tianhao Wang, Ninghui Li, Jean Honorio, Michael Backes, Shibo He, Jiming Chen, and Yang Zhang. 2021. PrivSyn: Differentially Private Data Synthesis. In *USENIX Security Symposium*.
- [49] Zhigao Zheng, Tao Wang, Jiming Wen, Shahid Mumtaz, Ali Kashif Bashir, and Sajjad Hussain Chaudhary. 2019. Differentially private high-dimensional data publication in internet of things. *IEEE Internet of Things Journal* 7, 4 (2019), 2640–2650.