Contents lists available at ScienceDirect

Journal of Information Security and Applications

journal homepage: www.elsevier.com/locate/jisa

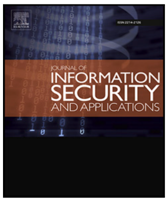

Exploiting attribute correlation for reconstruction attacks on differentially private multi-attributed data

Yanna Jiang [a] [1], Baihe Ma [a] [1], Xu Wang [a], Guangsheng Yu [a], Caijun Sun [b] [*], Wei Ni [c], Ren Ping Liu [a]

[a] University of Technology Sydney, Sydney, 2007, NSW, Australia
[b] Zhejiang Lab, Hangzhou, 311121, Zhejiang, China
[c] Data 61, Eveleigh, 2015, NSW, Australia

ARTICLE INFO

Keywords:
Differential privacy
Data attribute correlation
Reconstruction attack
Machine learning

ABSTRACT

Differential Privacy (DP) is a widely used data privacy-preserving technique with single-attribute DP being a common approach, in which manipulated noise is applied to each data attribute individually. However, data in practical scenarios often contains multiple data attributes, and the correlations between these attributes, which are often overlooked, introduce vulnerabilities to single-attribute DP schemes. In this paper, we present a rigorous analysis demonstrating that these correlations can undermine the protection offered by single-attribute DP schemes, with the risk of compromise increasing as the correlation between attributes becomes more pronounced. We propose a novel attack framework to evade the single-attribute DP protection on multi-attributed data by exploiting the overlooked data attribute correlations. We further implement the attack by developing Machine Learning (ML) algorithms to uncover the straightforward and hidden attribute correlations. Extensive experiments with various ML algorithms are conducted to corroborate our analysis, demonstrating the existence of privacy leakage caused by data attribute correlations and the effectiveness of the proposed attack with significantly enhanced reconstruction accuracy. In one of our experiments, the proposed attack method mitigated over 50% of the DP noise, significantly enhancing the accuracy of reconstruction attacks.

1. Introduction

Data privacy has taken center stage in the digital age [1], critical for maintaining the confidentiality of sensitive information [2] and preserving trust in the systems that rely heavily on big data [3]. Differential Privacy (DP) [4] stands out as a popular approach to protecting data privacy, drawing widespread attention due to its mathematical rigor and effectiveness [5]. Sophisticated DP mechanisms can introduce manipulated noises for data privacy protection [6] in data-sharing scenarios, widely applied across various domains, especially for vehicular networks [7,8] and finance [9,10].

However, DP still faces challenges [11] when protecting multi-attributed data against reconstruction attacks and inference attacks where attackers aim to recover original data or deduce sensitive information. Existing single-attribute DP researches [12,13] predominantly apply DP to perturb each data attribute individually, failing to capture complex correlations in multi-attributed data. Some studies [14,15] have explored the temporal dynamics of individual attributes and calibrated noise accordingly in user-level DP. These approaches rely on strong and explicit patterns (e.g., temporal ordering [14] or Markovian dependencies [15]) across a substantial number of records, but overlook correlations across attributes. This line of work implicitly treats each attribute as an independent unit, applying DP to each attribute in isolation without accounting for their underlying correlations. However, attributes within a single record often exhibit dependencies due to their shared origin from the same entity. Ignoring such intra-record correlations can disrupt the structural consistency of the data, leading to compounded privacy leakage when perturbed attributes are jointly analyzed or reconstructed.

Reconstruction attackers could potentially remove single-attribute DP noises on sensitive attributes, by leveraging the correlations between DP-obfuscated attributes and unprotected attributes. For instance, in vehicular networks, strong correlations exist between vehicle speed, fuel consumption, GPS location, and traffic conditions. Some sensitive attributes (e.g., location) can be obfuscated by DP to protect

[*] Corresponding author.
E-mail addresses: Yanna.Jiang@student.uts.edu.au (Y. Jiang), Baihe.Ma@uts.edu.au (B. Ma), Xu.Wang@uts.edu.au (X. Wang), guangsheng.yu@uts.edu.au (G. Yu), sun.cj@zhejianglab.com (C. Sun), Wei.Ni@data61.csiro.au (W. Ni), RenPing.Liu@uts.edu.au (R.P. Liu).
[1] Contributed equally to this research.

https://doi.org/10.1016/j.jisa.2025.104224

Available online 13 September 2025

privacy, while some driving data (e.g., speed and turns) may remain unobfuscated for system functionality. Attackers can infer DP-protected attributes (i.e., location) from plain-text attributes (i.e., speed and turns) [16].

This paper investigates the impact of attribute correlations on event-level DP protection [17], where each event consists of multiple correlated and heterogeneous attributes. We prove that such correlations compromise DP protection if DP is applied independently to individual attributes in isolation, leaving multi-attributed data vulnerable to reconstruction attacks. We propose a novel reconstruction attack that exploits attribute correlations to bypass single-attribute DP protection and recreate target data. Implemented using various Machine Learning (ML) models, our attack demonstrates the privacy leakage caused by attribute correlations and its effectiveness in breaching single-attribute DP safeguards. The attack shows significant potential in practical applications like in-vehicle records and financial data sharing. It highlights the inadequacy of current single-attribute DP protection against reconstruction attacks that leverage attribute correlations, addressing a critical gap in data privacy.

The key contributions of our work are listed as follows.

(1) We conduct a theoretical analysis of the privacy budget for single-attribute DP-protected multi-attributed data, revealing that the effectiveness of DP is significantly compromised by correlations among data attributes. We find that this compromise escalates with the number of attributes involved.
(2) We propose a new reconstruction attack framework that targets the univariate implementation of DP in multi-attributed data. This framework can capture both straightforward and hidden attribute correlations in the target data, thereby making the DP protections applied individually to each attribute collectively neutralized.
(3) We develop various ML algorithms to uncover the hidden correlations between data attributes. With the algorithms, we conduct new data reconstruction attacks on single-attribute DP-protected data, which challenges the conventional way of applying DP to individual data attributes.

Theoretical findings are validated through extensive experiments using real-world multi-attributed datasets and advanced ML models. Our results demonstrate a substantial reduction in DP noise, significantly improving the precision of reconstruction attacks. Furthermore, the proposed framework is versatile for various dataset types and DP mechanisms, effectively compromising the protection of numerical, non-numerical, and time-domain DP, as demonstrated by our experiments that successfully conducted reconstruction attacks across these single-attribute DP methods.

The rest of this paper is organized as follows. In Section 2, background and related works are reviewed. The proposed attack based on the data attributes correlation is presented in Section 3 and experimentally evaluated in Section 4, followed by conclusions in Section 5.

2. Background and related work

In this section, we first introduce the basic concepts and common mechanisms of DP for a better understanding of our work. Then, we discuss the threats to single-attribute DP posed by data correlation and provide an overview of the existing related work to highlight the innovation and necessity of our work.

2.1. Differential privacy

DP is designed to safeguard the confidentiality of individual data by providing a mathematically defined level of data protection [4]. It ensures that statistical analysis and data mining can be achieved while preserving the privacy of individuals' sensitive information [18].

Definition 1 (*ε-DP*)**.** A mechanism satisfying an ε-DP with privacy budget ε implies that if and only if, for any pair of datasets X and X' with only one data sample in difference, a possible obfuscated data y has

$$\frac{\Pr[y|X]}{\Pr[y|X']} \leq e^{\varepsilon}, \tag{1}$$

where a dataset refers to a finite collection of data with multiple attributes. As the privacy budget ε decreases, it becomes increasingly difficult to distinguish between X and X', resulting in higher data security.

Building upon this stringent definition, ε-DP can be transformed into a relaxation (ε, δ)-DP by introducing a certain probability threshold for potential privacy breaches.

Definition 2 (*(ε, δ)-DP [19]*)**.** A mechanism satisfying an (ε, δ)-DP with privacy budget ε and given δ implies that if and only if, for any pair of datasets X and X' with only one data inconsistency, a possible obfuscated data y has

$$\Pr[y|X] \leq \Pr[y|X'] \times e^{\varepsilon} + \delta. \tag{2}$$

The values of ε and δ inversely affect the level of security, thus indicating that a smaller value results in a higher degree of security.

As the privacy guarantee of DP relies on the idea that the inclusion or exclusion of any single data point should not significantly change the outcome of any analysis [20], it is crucial to introduce the concept of global sensitivity, which plays a fundamental role in calibrating the noise added to data for ensuring DP privacy.

Definition 3 (*Global Sensitivity*)**.** The global sensitivity measures the maximum possible difference by which a single data point can change the outcome of any query. Consider the complete data repository **B** that may include one or more datasets X. For a randomized query function f in **B**, the global sensitivity Δf satisfies

$$\Delta f = \max \|f(\mathbf{B}) - f(\mathbf{B}')\|_2, \tag{3}$$

where **B** and **B**′ differ by only one data point, and $\|\cdot\|_2$ represents the L_2 norm.

We categorize the popular DP mechanism used for data protection in ML into the following three types based on the data processing methods:

Numerical DP on data attribute. Adding noise to data is the widely adopted and fundamental approach to achieve DP protection [4]. By adjusting the magnitude of the added noise, one can strike a balance between data security and data utility [21]. Common methods include the Laplace and the Gaussian mechanism [22], especially applied in the fields of vehicular networks [7,23] and finance [24]. Here, the Laplace noise $L \sim \mathrm{Lap}(0, \Delta f/\epsilon)$ is computed using the Laplacian function [25] in the Laplace mechanism to achieve ϵ-DP and the Gaussian noise $G \sim \mathrm{Lap}(0, \sigma^2)$ with the σ satisfying for $\forall\delta \in (0,1), \sigma \geq \sqrt{2\ln(1.25/\delta)}\Delta f/\epsilon$ is applied based on Gaussian distribution [26] in Gaussian mechanism to achieve (ε, δ)-DP.

Non-numerical DP on data attribute. DP can also be attained by transforming the data into an interval candidate set and subsequently perturbing it [12]. This technique provides an alternative method to achieve DP by manipulating the data representation and introducing perturbations within the defined intervals [27], commonly used in applications like censuses [28] where non-numerical data is prevalent.

DP on time domain. Introducing random perturbations to modify the original time series of the data, thereby obfuscating the link between specific data points and individuals or sensitive information, represents another approach for DP [29]. This method effectively safeguards privacy by rendering it challenging to discern personal or sensitive details from the perturbed time series data, commonly used in the vehicular networks [30].

Each type of DP mechanism addresses the unique privacy concerns and data characteristics of different datasets effectively [31]. Numerical data often requires direct modification to obscure individual contributions, hence noise addition. Non-numerical data, in contrast, does not lend itself easily to noise addition, so transformation into intervals or categories followed by perturbation is a more effective approach. Lastly, time-series data presents unique challenges as the sequence of data points can be revealing, so modifying the temporal order or the specific timestamps can help mitigate privacy risks. This classification allows for targeted privacy protections that are specifically tailored to the nature of the data, enhancing both the effectiveness and applicability of DP across various data types and scenarios.

2.2. Threats to DP posed by data correlation

Although DP has emerged as a popular approach to ensuring rigorous privacy guarantees across diverse datasets and application domains [32,33], researchers have pointed out that DP may not offer privacy guarantees as expected [34]. There are concessions to the DP in models driven by the aim of accuracy, and it is hard to achieve the desired level of privacy protection. A recent study [35] demonstrated that DP controlled alone by the DP parameter ϵ might not provide effective protection against reconstruction attacks and more factors, such as the batch size hyperparameter and the amount of background knowledge, needed to be taken into account. The Pufferfish privacy framework [36] further supports this concern by showing that correlations in the data can fundamentally undermine DP guarantees, unless explicitly modeled and protected. However, existing Pufferfish implementations often rely on strong assumptions about the data distribution or correlation structure [37]. This differs from our setting, where attribute correlations are heterogeneous and less structured, often emerging implicitly in multi-attributed data without predefined statistical models.

The "No Free Lunch" theorem proposed by Kifer et al. [38] demonstrates that for any useful privacy mechanism, there exists a data distribution where the mechanism fails to ensure privacy, revealing a fundamental vulnerability in DP when correlations are present. Yang et al. [39] investigated how data correlations affect privacy and proposed a new privacy definition, by which the privacy level can be evaluated with related data. Liu et al. [40] conducted experiments on inference attacks using real-world datasets to further demonstrate that adversaries could exploit probabilistic dependencies between data to violate the DP guarantees. Wang et al. [41] performed a linkage attack to show the vulnerability of DP when the data are correlated. These existing studies have primarily focused on correlations between data points within databases, not the correlation between different attributes of the same multi-attributed data samples. Such works cannot readily extend to attribute-level correlations, as they typically assume homogeneous structures across data points, which enables direct modeling of inter-point dependencies, such as time-series or Markovian structures. In contrast, attributes within a single data point are often heterogeneous in type, format, and semantics, and their correlations are more complex, less explicit, and structurally constrained by the underlying entity.

Due to the high-dimensional curse in DP [42], directly applying DP to highly correlated attributes or dealing with multiple related queries can result in a degradation of privacy guarantees [43]. To achieve practical DP for high-dimensional and multi-attributed data, cutting-edge techniques such as compressed sensing mechanisms [44] are being integrated, mitigating the impact of dimensionality by efficiently reducing data complexity while preserving essential privacy characteristics. Zhang et al. [45] considered the influence of data attribute correlation on DP protection and managed the extent of data correlation with feature selection to achieve an improved data utility. However, these studies have only recognized the impact of data attribute correlations on DP without quantifying the extent of the impact.

Our research bridges this gap by focusing on the attribute-level correlations on multi-attributed data. Specifically, we introduce formally a mathematical formulation to estimate the upper bound of the DP privacy budget, thereby accounting for the inherent correlations among data attributes. Our findings indicate that these correlations can be exploited to compromise DP, as they provide additional information that is not altered by DP. By harnessing these correlations, adversaries can more effectively deduce and reconstruct the original data, thus challenging the privacy assurances provided by DP. This study explores the level of protection DP offers to multi-attributed data, demonstrating that overlooking the correlation between data attributes can lead to a false sense of privacy security.

3. Proposed attack model and theoretical analysis

This section proposes a novel reconstruction attack model against single-attribute DP protection using a new attack vector through data attribute correlation. The proposed attack effectively removes the DP noise applied to multi-attributed data and achieves a highly accurate reconstruction of the original data. With rigorous analysis, the proposed attack shows that the correlation between different data attributes can breach the privacy protection of multi-attributed data.

3.1. DP-protected multi-attributed data

In a database, the attributes within each data point often exhibit strong correlations, rather than being independent from each other. Taking Morningstar database [46], a financial database, for example, there is a functional correlation between attributes of revenue, total revenue, expense, and total expense. Similar correlations between data attributes are commonly observed in various databases, including vehicular [47] and finance databases [48]. These correlations reflect the intrinsic relationships and dependencies among different attributes within the data. We consider the scenario where the multiple attributes of data are correlated with each other. The attributes of data points can be classified into two categories [49]:

- **Confidential Attributes (CAs).** CAs are used to identify information that requires high protection to ensure privacy security, such as identification numbers and names.
- **Quasi-identifier Attributes (QAs).** The leakage of QAs information, like nationality and gender, does not pose an immediate threat to privacy.

Here, CAs and QAs are artificially defined; in reality, a CA can be any attribute within the data, making it a general process. For instance, the location in the vehicular dataset and the attribute of total current asset in the finance dataset are commonly considered as CAs [50,51] due to their sensitive nature, while other attributes like speed in vehicular networks and cash in finance are more often treated as QAs. Given that CAs require higher privacy protection, they are usually secured with stronger DP measures. Successfully reconstructing CAs implies that other attributes can also be reconstructed. Therefore, we focus on reconstructing CAs as the primary object of our attack.

Let $\mathbb{C} = \{CA_1, CA_2, \ldots, CA_m\}$ be the set of CAs and $\mathbb{Q} = \{QA_1, QA_2, \ldots, QA_n\}$ be the set of QAs. Assume that there exists a correlation function $g(\cdot)$ between $\mathbb{C}$ and $\mathbb{Q}$, such that: $\forall CA_i \in \mathbb{C}$, $i \in [1, m]$,

$$CA_i = g(\mathbb{Q}_i), \tag{4}$$

where $\mathbb{Q}_i$ is a subset of QAs that is related to CA_i, i.e., $\mathbb{Q}_i \subseteq \mathbb{Q}$, and $\mathbb{Q}_i = \{QA_{i_1}, QA_{i_2} \cdots, QA_{i_x}\}$. Here, QA_{i_x} is the xth QA of $\mathbb{Q}_i$. We note that the form of the correlation function $g(\cdot)$ may vary for different CA_i. More precisely, each CA_i may correspond to a distinct function $g_i(\cdot)$. Since this paper focuses on the reconstruction of a specific target attribute CA_i, for conciseness, we suppress the subscript "$_i$" in function $g_i(\cdot)$.

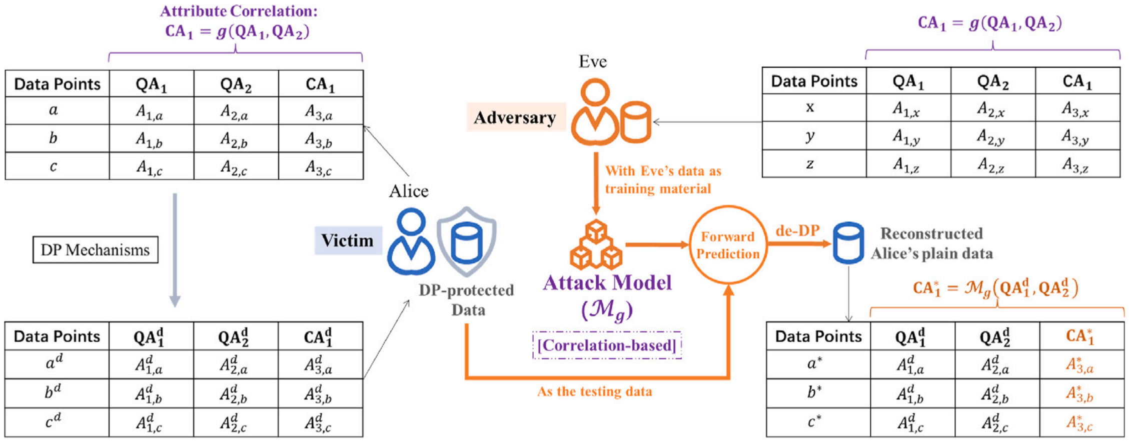
Attribute Correlation:
CA1 = g(QA1, QA2)
Data Points | QA1 | QA2 | CA1
a | A1,a | A2,a | A3,a
b | A1,b | A2,b | A3,b
c | A1,c | A2,c | A3,c
DP Mechanisms
Data Points | QA1d | QA2d | CA1d
ad | A1,ad | A2,ad | A3,ad
bd | A1,bd | A2,bd | A3,bd
cd | A1,cd | A2,cd | A3,cd
Alice
Victim
DP-protected Data
Eve
Adversary
With Eve's data as training material
Attack Model (Mg)
[Correlation-based]
As the testing data
Forward Prediction
de-DP
Reconstructed Alice's plain data
CA1 = g(QA1, QA2)
Data Points | QA1 | QA2 | CA1
x | A1,x | A2,x | A3,x
y | A1,y | A2,y | A3,y
z | A1,z | A2,z | A3,z
CA1* = Mg(QA1d, QA2d)
Data Points | QA1d | QA2d | CA1*
a* | A1,ad | A2,ad | A*3,a
b* | A1,bd | A2,bd | A*3,b
c* | A1,cd | A2,cd | A*3,c

Fig. 1. The proposed reconstruction attack for multi-attributed data. Data owners use DP mechanisms to protect their data, while the adversary can exploit the data attribute correlation to launch the attack effectively. Here, CA is artificially defined, and any attribute can be selected as a CA for reconstruction.

To prevent unauthorized access or inference from curious or malicious external parties, the data owner protects the data with DP. Based on the data characteristics and privacy requirements, the owner selects a single-attribute DP mechanism, such as adding Laplace [25] or Gaussian [26] noise to numerical attributes, and then individually perturbs each CA_i and QA_j within the data, resulting in DP-protected CA_i^d and QA_j^d with

$$\begin{aligned} CA_i^d &= \mathcal{DP}(CA_i, \varepsilon_{CA_i}, \delta_{CA_i}), \\ QA_j^d &= \mathcal{DP}(QA_j, \varepsilon_{QA_j}, \delta_{QA_j}), \end{aligned} \tag{5}$$

where CA_i^d satisfies $(\varepsilon_{CA_i}, \delta_{CA_i})$-DP and QA_j^d satisfies $(\varepsilon_{QA_j}, \delta_{QA_j})$-DP. The privacy budgets ε_{CA_i} and ε_{QA_j} are determined by the owner based on their needs. Since CAs require higher privacy protection, they are typically assigned a smaller ε as their privacy budget.

Considering the attribute correlations between CAs and QAs, the current practice of applying DP protection to individually perturb data attributes overlooks potential security vulnerabilities. Attackers can exploit less protected or unprotected CAs to infer heavily protected QAs, thereby breaching defenses and achieving highly accurate data reconstruction attacks.

3.2. Proposed attack model

The proposed attack employs the data attribute correlation to launch a reconstruction attack, as depicted in Fig. 1. We consider scenarios where the target multi-attributed data is shared publicly or illegally accessed by the attacker. In this case, the data owner employs DP methods to protect the CAs and QAs of their original data before release. The attacker, on the other hand, may possess prior knowledge about the CAs and QAs of the target data and the existence of correlations between the CAs and QAs. Current data privacy protections focus primarily on shielding the values and the distribution of attributes [52] rather than the correlations between attributes, making it relatively easy for the attacker to acquire this knowledge. We assume that the attacker has access to a dataset with the same data structure and attributes as the target. Such access can arise through centralized authorities with population-level data (e.g., tax offices), the use of standardized data formats across institutions (as in healthcare or finance), or through private datasets in schema-aligned settings such as Federated Learning (FL) [53], or inter-organizational collaborations. This allows the attacker to exploit not only obvious and well-known attribute correlations, but also those subtle and less-known.

More specifically, Alice, the owner of multi-attributed data and the target of our attack, employs a DP mechanism individually on each attribute to obtain the DP-protected CA_i^d, QA_i^d and the corresponding $\mathbb{Q}_i^d$. However, the proposed attack exploits a vulnerability inherent in the correlation of data attributes, i.e., the additional information leaked by the function $g(\cdot)$. Eve, the attacker, exploits her own dataset, similar to Alice's, as training material to develop a specialized attack model $\mathscr{M}_g$, which is designed to exploit these data attribute correlations and reconstruct the real CAs, i.e.,

$$CA_i^* = \mathscr{M}_g(\mathbb{Q}_i^d), \tag{6}$$

where, $\mathbb{Q}_i^d$ represents the DP protected $\mathbb{Q}_i$, and CA_i^* is the reconstructed result for CA_i. Compared to CA_i^d, which is the CA_i after DP mechanism processing, CA_i^* closely approximates the original CA_i, as CA_i^* is evaded the typical DP configuration during the prediction in $\mathscr{M}_g$.

This attack model $\mathscr{M}_g$ analyzes and identifies patterns in Alice's data that are not sufficiently masked by the DP mechanisms. Using this trained attack model, Eve processes Alice's shared or published data to evade the typical DP configuration focusing on individual attributes. The attack model $\mathscr{M}_g$ effectively strips away the noise added by DP, revealing underlying data patterns and, ultimately, the original, unobfuscated data. This breach compromises the DP privacy guarantee and allows Eve to achieve her objective of accessing and reconstructing Alice's plain data (focused on CAs).

The proposed attack exploits a significant limitation of the current DP application, particularly when dealing with multi-attributed data that have inherent correlations between attributes. In the rest of this section, we then provide rigorous theoretical proofs to analyze the privacy protection degradation of DP caused by data attribute correlation.

3.3. Attack impact analysis

In this section, we delve into the impact of data attribute correlation on DP attacks. We estimate the upper bound for the corresponding privacy budget as measures of attack impact. A larger privacy budget indicates weaker privacy-preserving capabilities of DP, making it more susceptible to attacks.

To comprehensively analyze the effects of attribute correlations, we consider two realistic scenarios under different assumptions about the QAs. First, we analyze the case where $QA_i, \forall i$ are independent of each other, and then extend our analysis to the more complex case where $QA_i, \forall i$ exhibit interdependence, which often occurs in real-world datasets, such as financial records.

3.3.1. Privacy budget with independent QAs

When the QAs are independent, our focus narrows to the correlation between CAs and QAs. Assuming that both the CAs and QAs satisfy (ε, δ)-DP, we present the following theorem:

Theorem 1. *Suppose that CAs and independent QAs are protected by* $(\varepsilon_{CA_i}, \delta_{CA_i})$*-DP and* $(\varepsilon_{QA_j}, \delta_{QA_j})$*-DP schemes, respectively, and there exists a correlation function* $CA_i = g(QA_{i_1}, \dots, QA_{i_x})$*. Then, the effective privacy guarantee on* CA_i *is limited to* (ε', δ')*-DP, where* $\varepsilon' = \max(\varepsilon_{CA_i}, \sum_{j=1}^{x} \varepsilon_{QA_{i_j}})$ *and* $\delta' = \max(\delta_{CA_i}, \sum_{j=1}^{x} \delta_{QA_{i_j}})$*, which could be weaker than the intended* $(\varepsilon_{CA_i}, \delta_{CA_i})$*-DP.*

Proof. This theorem can be proved by first applying the DP definition to each QA, then summing their privacy budgets under the assumption of independence, and finally extending this bound to the CA via the post-processing property. Comparing this bound with the CA's declared DP guarantee shows that the true privacy loss is the larger of the two, indicating that aggregating results from multiple QAs can weaken overall privacy. For the detailed proof, please refer to Appendix A. □

Theorem 1 shows the limitation in applying DP independently to each attribute in multi-attributed data, as correlations among attributes can weaken the overall privacy guarantee. The theorem reveals that the effective privacy budget on a CA is constrained not only by its own DP parameters but also by the cumulative DP budgets of the correlated QAs. The theorem reveals that even with strict DP protections on individual attributes, unaddressed correlations may lead to higher-than-anticipated privacy risks.

Most existing DP-based methods [12,54] suggested that insensitive data (i.e., QA) can use a large privacy budget to obtain high data utility, which means ε_{CA_i} with sensitive information should not be larger than any $\varepsilon_{QA_{i_j}}$. In this case, we have:

$$\sum_{j=1}^{x} \varepsilon_{QA_{i_j}} \geq \varepsilon_{CA_i}. \tag{7}$$

Hence, the actual privacy-preserving capability of CA_i takes $\sum_{j=1}^{x} \varepsilon_{QA_{i_j}}$, which is larger than the expected ε_{CA_i}. In other words, the privacy of ε_{CA_i} will be exhausted faster than expected. This vulnerability is related to the number of QA_{i_j} associated with CA_i in the correlation function $CA_i = g(QA_{i_1}, \dots, QA_{i_j}, \dots, QA_{i_x})$. The greater the number of QA_{i_j} linked to CA_i, the higher the sum of $\varepsilon_{QA_{i_j}}$ on the left side of (7), increasing the privacy budget and making the privacy of CA_i more susceptible to leakage.

3.3.2. Privacy budget with associated QAs

Considering that the QAs are interdependent, where the value of a QA is influenced by the values of others, we then prove that the privacy protection of a CA is also degraded by its related QAs if a correlation exists between CAs and QAs. For illustration convenience, in this section, we assume all QAs use the same DP parameter, i.e., all QAs satisfy (ε, δ)-DP.

Theorem 2. *Suppose that CAs and interdependent QAs are protected by* $(\varepsilon_{CA_i}, \delta_{CA_i})$*-DP and* (ε, δ)*-DP schemes, respectively, and there exists a correlation function* $CA_i = g(QA_{i_1}, \dots, QA_{i_x})$*. Then, the effective privacy guarantee on* CA_i *is limited to* (ε', δ')*-DP, where* $\varepsilon' = \max(\varepsilon_{CA_i}, \sqrt{2x\ln(1/\delta')}\varepsilon + x\varepsilon(e^{\varepsilon} - 1))$ *for* $0 < \delta' < 1$*, which could be weaker than the intended* $(\varepsilon_{CA_i}, \delta_{CA_i})$*-DP.*

Proof. Theorem 2 shows that when QAs are interdependent, privacy leakage can accumulate through their correlations. The theorem can be proved by modeling the accumulation via a martingale sequence of log-likelihood ratios, then applying Azuma's inequality to bound the combined leakage, and finally using the post-processing property to transfer this bound to the CA. Comparison with the CA's nominal DP guarantee shows that the effective privacy bound is the larger of the two, indicating that correlation-induced leakage can degrade the intended level of protection. For the detailed proof, please refer to Appendix B. □

Theorem 2 confirms that the actual privacy guarantee can fall short of the intended protection level in the case of interdependent QAs due to attribute correlations, highlighting the limitations of single-attribute DP protection on multi-attributed data.

Consistent with the analysis in Theorem 1, when $\varepsilon' > \varepsilon_{CA_i}$, the privacy protection for CA_i is weaker than expected; i.e., the correlation of data attributes can significantly degrade the effectiveness of DP protection. Following the setting of Theorem 2 that all QAs satisfy (ε, δ)-DP, i.e. $\forall i \in n$, $\varepsilon_{QA_i} = \varepsilon$, Theorem 1 can be rewritten as that the privacy budget of CA_i is $\max(\varepsilon_{CA_i}, x\varepsilon)$ when the QAs are independent. Comparing the two theorems, we conclude that $\sqrt{2x\ln(1/\hat{\delta})}\varepsilon + x\varepsilon(e^{\varepsilon} - 1) \geq x\varepsilon$ when $\varepsilon \geq \ln(2)$. To this end, we can conclude that the interdependent QAs decrease the privacy protection of CA_i when $\varepsilon \geq \ln(2)$. It is corroborated that attribute correlations, especially in real-world datasets like financial records, can deteriorate privacy leakage beyond the expectation set by traditional DP assumptions.

4. Experimental results

In this section, we validate Theorems 1 and 2 through experiments and demonstrate that the proposed attack can effectively reconstruct multi-attributed data protected by single-attribute DP. We first conduct experiments using various ML models as attack models for the proposed reconstruction attack to assess the impact of the attack models. Then, we target data protected by different DP methods in the proposed reconstruction attack to demonstrate that the proposed attack is universally effective to any DP methods and shows significant potential for widespread adoption and application. To align with real-world application scenarios, we select two multi-attributed datasets (see Table 1) from different domains to conduct the proposed reconstruction attack, as follows:

Vehicular Dataset. We use the GeoLife dataset, which consists of GPS recordings obtained from over 180 drivers in Beijing, China [47]. This dataset captures real multi-attributed data that is jointly recorded by in-vehicle GPS loggers and smartphones, illustrating a key trend in the vehicular network toward enhanced connectivity and real-time data sharing for improved traffic management and user services [56]. Moreover, these records trace the daily routes of users and commonly employ DP for privacy protection, as attacks on such data could pose significant threats to user privacy. Therefore, choosing this dataset to validate the proposed attack in this paper is both logical and of practical value.

We preprocess the GeoLife dataset, resulting in a dataset where each data sample contains eight attributes: vehicle number, longitude and latitude of the starting position, longitude and latitude of the ending position, travel time, travel speed, and travel direction. We designate the longitude and latitude of the ending position as CAs, which require the highest level of privacy protection, while treating the other attributes as QAs. According to the relationship between displacement, time, and speed, the CAs, i.e., the longitude and latitude of the ending position, can be deduced from QAs such as the longitude and latitude of the starting position, travel time, travel speed, and travel direction, indicating the existence of attribute correlation $CA = g(QA)$. Since the data in this dataset are recorded through onboard or mobile devices, there is some numerical inaccuracy compared to theoretical models. Therefore, the function $g(\cdot)$ in this case represents a more relaxed and implicit association rather than a strict mathematical equation.

Finance Dataset. We employ a widely-used finance dataset sourced from the open-access SIRCA dataset[2], which includes data from 1596 providers across 48 distinct attribute labels [55]. This dataset is integral for real-time financial analytics and decision-making processes in practical applications. We extract 9 attributes related to current

[2] https://datalibrary.sirca.org.au/index ASX Dataset, membership required.

Table 1
Dataset overview.

Dataset	No. records	No. CA	CA	No. QA	QA
Vehicular Dataset [47]	44 593	2	Ending longitude; Ending latitude.	6	Vehicle number; Starting longitude; Starting latitude; Travel time; Travel speed; Travel direction.
Finance Dataset [55]	4788	1	Total current asset.	8	Company name; Cash; Receivables; Prepaid expense; Inventories; Investment; NCA held for sale; Other.

assets: company name, cash, receivables, prepaid expense, inventories, investment, NCA held for sale, other, and total current assets, and select data from all suppliers from 2018 to 2020 for our experiments. These attributes are crucial as they provide a comprehensive view of a company's short-term financial state and ability to manage and utilize its resources effectively, influencing decisions made by management, investors, and creditors.

Since "Total Current Asset" is a key indicator of a company's financial health and operational efficiency, we designate the Total Current Asset as CA with the highest privacy requirement and the other attributes as QAs. According to the fundamental principle in finance, the Total Current Asset is equal to the total of all the other numeric attributes [57], demonstrating attribute correlation $CA = g(QA)$. In this case, the function $g(\cdot)$ is a strict mathematical equation.

According to the definition of CAs and QAs, we focus on the attack against CAs, which is highly correlated with QAs and demands stringent data protection. We randomly shuffle the data and evenly distribute it between the attacker and the target, ensuring that the attacker possesses sufficient prior knowledge to learn and analyze the data attribute correlations in the local dataset of the target. After DP protection measures, the local dataset of the target is made public, while the attacker trains an attack model with his local dataset to evade the DP configuration on the publicly released dataset from the target and execute the proposed reconstruction attack.

We assess the effectiveness of the reconstruction attack by measuring the distance between the reconstructed results and the original target data. A smaller average distance indicates that the reconstructed results closely resemble the original target data, leading to highly accurate reconstructions and effective attacks. Additionally, we analyze the distribution of the distance with the value of the Probability Density Function (PDF) and the Cumulative Distribution Function (CDF) [58], and a high concentration of values around zero indicates an accurate inference and effectiveness of the reconstruction attack. The distance before and after the application of DP protection on the original target data is utilized as a benchmark for comparison, highlighting the effectiveness of the attack.

Our experiments are run on a computer with AMD Ryzen 9 5900HX and 16G of RAM, and we use Python 3.9.12 and TensorFlow 2.9.1 to build and train the ML models for reconstruction attacks.

4.1. Proposed attack with different ML models

The proposed reconstruction attack trains an ML model with datasets similar to the target to learn the correlations between data attributes. This learned correlation is then exploited to evade the single-attribute DP configuration and reconstruct the target dataset. To investigate the effect of different ML models, we incorporate five prevalent models into the proposed architecture as the attack model:

- Multilayer Perceptron (MLP): A fully connected feedforward neural network model consisting of multiple layers of interconnected nodes, commonly used for various tasks such as classification and regression in ML [59]. Here, we utilize five fully-connected layers for the experiments.
- Convolutional Neural Network (CNN): A deep learning model designed to automatically extract relevant features from input data, especially suited for image and video analysis tasks [60]. Here, we apply two convolutional layers connected to a fully-connected layer for the experiments.
- Recurrent Neural Network (RNN): A type of neural network that can process sequential data by maintaining internal memory, making it well-suited for tasks involving time series data and natural language processing [61]. Here, we chose a SimpleRNN layer and a fully-connected layer for the experiments.
- Long Short-Term Memory (LSTM): A specific type of RNN architecture that addresses the vanishing gradient problem by incorporating memory cells, allowing it to capture long-term dependencies and effectively process sequential data [62]. Here, we utilize a classic LSTM layer with a dropout connected to a fully-connected layer for the experiments.
- Attention Model: A neural network architecture that learns to focus on specific parts of the input sequence, dynamically weighting the importance of different elements, commonly used in tasks involving sequential or textual data [63]. Here, we use an Attention module with weight calculation, followed by two fully-connected layers for the experiments.

We evaluate the effectiveness of the proposed attack model with the above ML model structures. In this experiment, the dataset owner employs the traditional Laplace Mechanism [25] to protect their original data, as it is one of the most commonly used numerical DP mechanisms. To further demonstrate the effectiveness of our proposed attack, we compare it against two state-of-the-art reconstruction attacks: the Random FP attack [64], which targets implementation-level weaknesses in DP mechanisms by exploiting floating-point side channels and applying Maximum Likelihood Estimation (MLE), and the ImpI attack [65], which reconstructs target attributes using a trained model and auxiliary dataset with a similar structure.

As shown in Figs. 2(a) and 2(d), the proposed attack with all the mentioned ML model structures effectively removes the noise introduced by DP protection, thus reconstructing a dataset with a smaller average distance from the original dataset. The rapid convergence of all models indicates their ability to quickly learn and capture the strong correlation between data attributes that all models make accurate reconstructions. MLP and CNN outperform RNN, LSTM, and attention models, especially on the finance dataset. This is because RNN, LSTM, and attention models are suitable for sequential data, but the sequential characteristic is weak in the used datasets.

In contrast, polynomial fitting with MLP and feature extraction with CNN are suited to launch the proposed attack model in such a scenario, with CNN performing slightly better. Compared to the proposed approach, FP attack shows minimal reconstruction ability. Its results are close to those of the original DP-protected data, indicating that it is ineffective in attribute-level reconstruction. The ImpI attack achieves reconstruction performance comparable to the proposed attack with LSTM structure, but falls short of MLP and CNN. This is because ImpI exploits distributional disparities between groups rather than modeling fine-grained attribute-level correlations, making it less effective in the targeted reconstruction setting we study.

The reconstruction results using the CNN model show that the distance between the reconstructed data and the original data is reduced by approximately 50%, compared to the distance between the DP-protected results and the original data (specifically 46.84% for the vehicular dataset and 50.50% for the finance dataset). This demonstrates that the protective capability of DP is conditional. Due to the differences in data distribution between the vehicular and finance

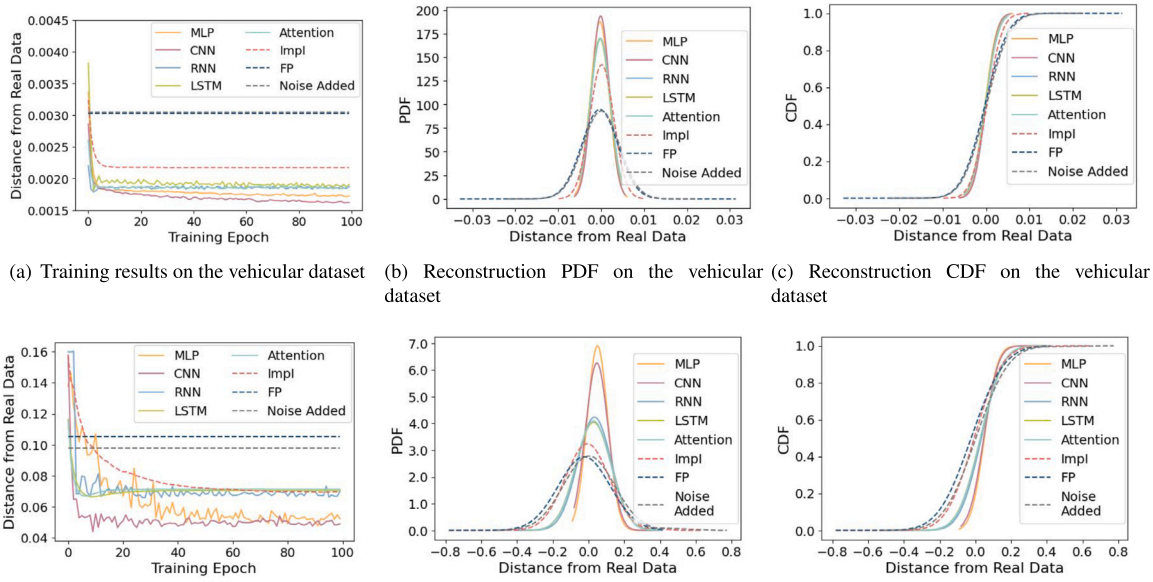

(a) Training results on the vehicular dataset (b) Reconstruction PDF on the vehicular dataset (c) Reconstruction CDF on the vehicular dataset

(d) Training results on the finance dataset (e) Reconstruction PDF on the finance dataset (f) Reconstruction CDF on the finance dataset

Fig. 2. The effect of reconstruction attack with different model structures on numerical DP. In Figs. 2(a) and 2(d), the effectiveness of these attacks is evaluated based on the proximity of the reconstructed data to the original target data, as reflected by the average distance. In Figs. 2(b) and 2(e), it is observed that by launching the proposed inference model, the inferred data approaches its original version with a high probability, i.e., the distance between the inferred data and the original data is approximately 0. Figs. 2(c) and 2(f) demonstrate that the CDF value of the inferred results grows rapidly when the distance is close to 0. In our evaluated setting, FP attack performs comparably to baseline DP results, while ImpI achieves moderate improvement but remains inferior to our proposed method.

datasets, it is reasonable that the magnitudes of distances between the reconstructed results and the original datasets vary in Figs. 2(a) and 2(d).

We evaluate the statistics of the reconstruction results of the attack using PDF and CDF. In Figs. 2(b) and 2(e), it is observed that by launching the proposed model, the reconstructed data approaches its original data with a high probability, i.e., the distance between the reconstructed data and the original data is close to zero. Due to the uneven distribution and significant variation of data samples in the finance dataset, the reconstructed results are slightly larger than the original data with an offset of about 6% in Fig. 2(e). Since the data in the vehicular dataset are relatively close to each other, the PDF of the distance between the reconstructed results and the original data in Fig. 2(b) is symmetrical around zero. Figs. 2(c) and 2(f) demonstrate that the CDF value of the reconstructed results grows rapidly when the distance is close to zero. These illustrate that the proposed attack model can effectively perform reconstruction attacks and reconstructed results close to the original data, highlighting the validity of Theorems 1 and 2.

4.2. Proposed attack against various DP methods

We further investigate the performance of the proposed attack on target data protected by different DP methods. Following previous categorizations, we consider three types of DP: numerical DP, non-numerical DP, and time-domain DP. We conducted experiments using the best-performing MLP and CNN as attack models.

It is worth noting that the baseline attack methods used earlier, i.e., the Random FP attack [64] and the ImpI attack [65], are inherently restricted to numerical DP settings. For example, Random FP attacks rely on floating-point leakage specific to numeric mechanisms, and ImpI attacks assume numerical consistency to perform accurate inference. These constraints render them unsuitable for non-numerical DP, and time-domain DP applications. In contrast, our attack leverages the inherent correlations between attributes and can be applied across all types of single-attribute DP mechanisms, highlighting its broader applicability and practical significance.

4.2.1. Analysis under numerical DP

We evaluate the proposed inference model under the DP mechanism, which directly applies noise addition to the numerical data point. As shown in Fig. 2, for numerical DP, the proposed attack effectively utilizes MLP and CNN to learn data attribute correlations and evade the DP configuration.

To further compare the impact of different numerical DP methods, we conduct experiments using the classic Laplace and Gaussian mechanisms. We analyze the impact of Laplace and Gaussian mechanisms with varying privacy budgets ϵ on the effectiveness of the proposed reconstruction attack across different model structures. It is evident in Fig. 3 that both the Laplace and Gaussian mechanisms exhibit reduced noising effects and weak security as the ϵ increases. Due to variations in data distribution, we apply different normalization techniques to the vehicular and finance datasets, resulting in slight differences in the impact of the Laplace and Gaussian mechanisms on the level of noise addition.

As shown in Fig. 3, the proposed attack, employing both MLP and CNN models, effectively breaches DP protection across various values of ϵ. As ϵ increases, the reconstructed results increasingly approximate the original data. When ϵ is small, and the added noise is significant, the proposed attack can eliminate most of the noise and evade the DP configuration. Comparing Figs. 3(a) and 3(b), it is observed that the Laplace and Gaussian mechanisms do not show significant differences on the vehicular dataset, whereas notable disparities exist in the finance dataset. The proposed attack can effectively reconstruct data on both the Laplace and Gaussian mechanisms, yielding reconstruction results very close to the original. The Gaussian mechanism is more effective on the finance dataset, suggesting that using Gaussian noise can more naturally integrate into datasets with large variations in data samples to provide better protection and impede the proposed attack.

4.2.2. Analysis under non-numerical DP

We evaluate the effectiveness of the proposed inference model on non-numerical DP. In the vehicular dataset, we partition the geographical region encompassing all vehicle trajectory points (i.e., Beijing,

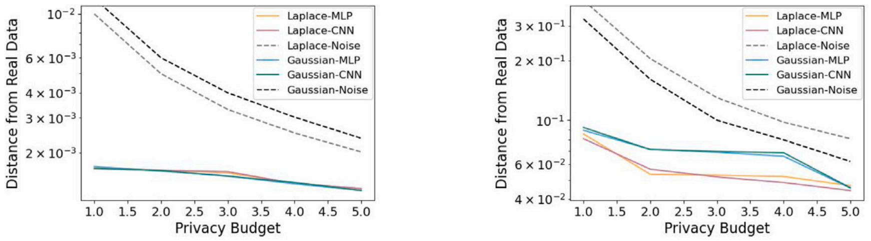

(a) Numerical DP mechanism on the vehicular dataset

(b) Numerical DP mechanism on the finance dataset

Fig. 3. The impact of Laplace and Gaussian mechanisms with varying ϵ on the effectiveness of the proposed reconstruction attack. The proposed attack, employing both MLP and CNN models, effectively breaches DP protection across various values of ϵ. As ϵ increases, the reconstructed results increasingly approximate the original data.

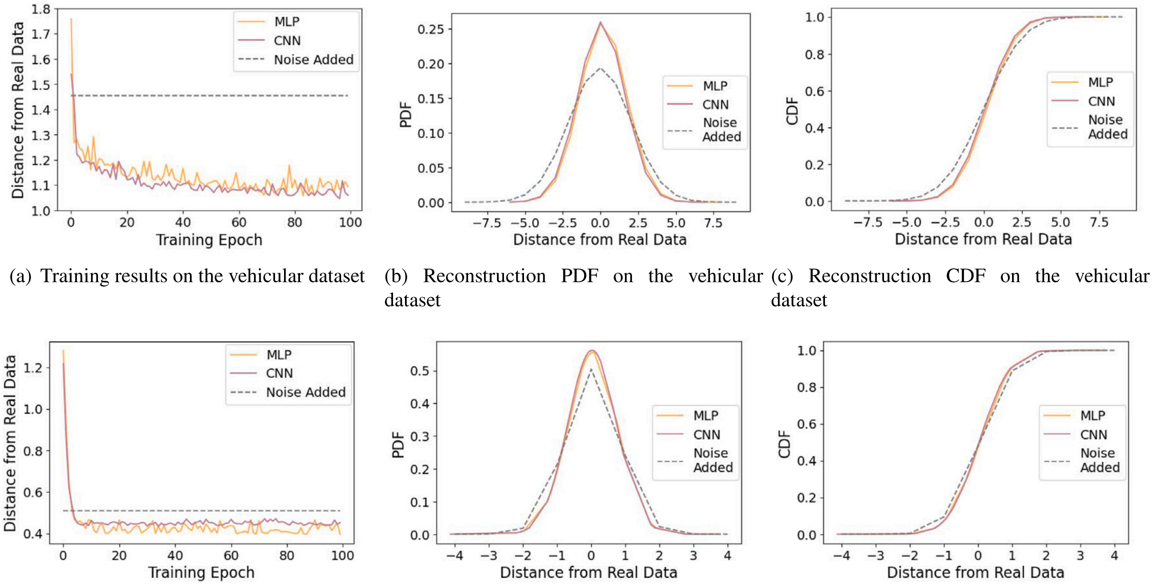

(a) Training results on the vehicular dataset

(b) Reconstruction PDF on the vehicular dataset

(c) Reconstruction CDF on the vehicular dataset

(d) Training results on the finance dataset

(e) Reconstruction PDF on the finance dataset

(f) Reconstruction CDF on the finance dataset

Fig. 4. The effect of reconstruction attack on non-numerical DP. On the vehicular dataset, the proposed attack effectively breaches the non-numerical DP security protection, while on the finance dataset, its performance is not as strong due to significant disruption of the data attribute correlation during the conversion to a non-numerical format that poses challenges for the proposed attack model.

China) into a series of small intervals. Each location information is clustered into the corresponding interval, converting the numerical data into non-numerical candidates. Similarly, we divided the finance data into a series of ranges based on the value size, replacing all numbers within the same range with the same interval code, thereby converting numerical data into non-numerical data. After converting to the non-numerical data, all the original data in both the vehicular and finance datasets are perturbed based on the distance to the obfuscated candidate with a crafted probability for DP protection.

Fig. 4(a) illustrates the impact of the proposed attack with MLP and CNN structures on the vehicular dataset protected by non-numerical DP, where both MLP and CNN models effectively breach the DP security threshold (eliminating about 38.58% of the DP protection), satisfying Theorems 1 and 2. Compared with the results obtained under numerical DP protection in Fig. 2(a), MLP and CNN models require more training epochs to achieve convergence under non-numerical DP protection. The reason is that the non-numerical approach introduces a distortion in the functional relationship between data attributes, leading to a reduction in data attribute correlation. Since the results are presented in discrete settings, slight variations can lead to abrupt changes in the outcomes, making the oscillation of the model performance curve more pronounced. The PDF in Fig. 4(b) and the CDF in Fig. 4(c) exhibit similar trends to those under numerical DP. This suggests that the reconstructed results align statistically with the original data, indicating a high level of accuracy in the reconstruction attack. The PDF in Fig. 4(b) and the CDF in Fig. 4(c) curves also exhibit irregular folds caused by discretization, leading to non-smoothness in their shapes.

However, it is shown in Figs. 4(d), 4(e) and 4(f) that the proposed attack does not perform as well on the finance dataset protected by non-numerical DP as it does on the vehicular dataset, eliminating only about 18.71% of the DP protection. This is because, in the finance dataset, the data attribute correlation is represented by a simple and strict mathematical expression. Once converted into a non-numerical format, this attribute correlation is significantly disrupted, limiting the ability of the proposed attack model to utilize the available attribute correlations for evading the DP configuration. This indicates that the effectiveness of the proposed attack depends on the form of data attribute correlation. The stricter the mathematical expression of data attribute correlation, the more likely it is to be disrupted by DP mechanisms, especially non-numerical DP, posing challenges to the proposed attack.

4.2.3. Analysis under time-domain DP

In this experiment, we evaluate the effectiveness of the proposed inference model on DP protection methods that apply random substitution to the time-domain data points. Here, on the vehicular dataset,

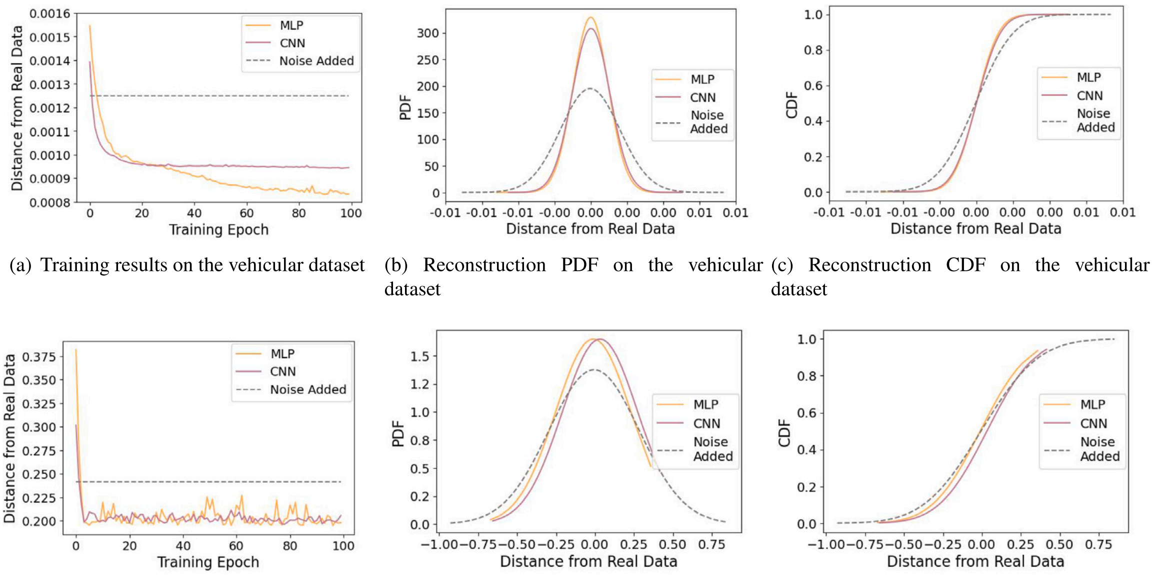

(a) Training results on the vehicular dataset (b) Reconstruction PDF on the vehicular dataset (c) Reconstruction CDF on the vehicular dataset

(d) Training results on the finance dataset (e) Reconstruction PDF on the finance dataset (f) Reconstruction CDF on the finance dataset

Fig. 5. The effect of reconstruction attack on time-domain DP. The proposed attack model performs high destruction of DP on both the vehicular and the finance datasets. In particular, under DP protection based on time-domain perturbation, the performance of MLP surpasses that of CNN significantly on the vehicular dataset as shown in Fig. 5(a), because the ability of MLP to capture associations between data attributes is less constrained by time-domain DP.

time-domain DP involves the exchange of points within the vehicle trajectory, resulting in the transformation of originally sequential position information into randomly ordered points. On the finance dataset, we randomly swap the corresponding data attributes of the same provider across different years to implement time-domain DP.

The proposed attack model employing MLP and CNN structures under time-domain DP is depicted in Fig. 5. The proposed attack model performs high destruction of DP and gets the reconstructed results close to the original data on both the vehicular and the finance datasets. Especially on the vehicular dataset, the proposed attack model reduces the effect of DP protection by around 33.28%. Overall, the proposed reconstruction attack effectively evades the typical DP configuration focusing on individual attributes under numerical, non-numerical, and time-domain DPs, validating Theorems 1 and 2 that state DP provides an overestimated protection. The attack is most effective against numerical DP because this method causes the least disruption to data attribute correlations.

In our previous experiment (Figs. 2 and 4), MLP and CNN models showed similar performance under numerical and non-numerical DP. However, with time-domain DP perturbation, MLP significantly outperformed CNN on the vehicular dataset (Fig. 5(a)). This suggests that time-domain perturbation affects CNN's feature extraction capabilities while limiting MLP's ability to capture data attribute associations. Interestingly, this characteristic was not prominent in the finance dataset (Fig. 5(d)). Similar to non-numerical DP, time-domain DP disrupts the strict data attribute correlation in financial data. This limits the attack's effectiveness to the attribute correlations the model can learn to evade the DP configuration, rather than the model structure itself.

5. Conclusion

In this paper, we presented a novel privacy attack approach that exploits attribute correlations in multi-attributed data to undermine DP protections designed for single attributes. Our theoretical analysis revealed that single-attribute DP effectiveness was significantly compromised by attribute correlations, with the vulnerability increasing alongside the number of attributes. We introduced a reconstruction attack framework that targeted single-attribute DP implementations, capturing both obvious and hidden correlations. Using various ML algorithms, we uncovered these correlations and conducted precise reconstruction attacks. Experiments with real-world datasets validated our findings, demonstrating the framework's effectiveness across different dataset types and DP mechanisms, including numerical, non-numerical, and time-domain implementations.

Both theoretical insights and practical applications suggest that the privacy leakage issues arising from attribute correlations warrant serious consideration. To address these challenges, privacy budgets should be carefully set considering the impact of QAs on CAs, potentially adjusting the privacy budget based on the number of associated QAs, employing dynamic privacy budgets [66], or integrating with emerging privacy protection technologies [67,68]. In the future, we plan to develop a practical and effective privacy protection scheme for multi-attributed data to mitigate the vulnerabilities identified in this study. One potential approach involves developing anonymity technologies to disrupt data attribute correlations.

CRediT authorship contribution statement

Yanna Jiang: Writing – original draft, Visualization, Validation, Methodology, Conceptualization. **Baihe Ma:** Investigation, Conceptualization. **Xu Wang:** Writing – review & editing, Supervision, Conceptualization. **Guangsheng Yu:** Writing – review & editing, Supervision, Methodology. **Caijun Sun:** Visualization, Validation. **Wei Ni:** Writing – review & editing, Supervision. **Ren Ping Liu:** Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Proof of Theorem 1

Proof. Given two possible values of CA_i, i.e., CA_i' and CA_i'', there exist two groups of QA, i.e., $\mathbb{Q}_i' = \{QA_{i_1}', \dots, QA_{i_x}'\}$ and $\mathbb{Q}_i'' = \{QA_{i_1}'', \dots, QA_{i_x}''\}$, which are related to CA_i' and CA_i'', respectively. $\mathbb{Q}_i'$ and $\mathbb{Q}_i''$ are two possible values of $\mathbb{Q}_i$ that is related to CA_i. According to (4), there is a correlation function $g(\cdot)$, yielding $CA_i' = g(\mathbb{Q}_i')$ and $CA_i'' = g(\mathbb{Q}_i'')$.

Let $f(\cdot)$ be the DP-based obfuscated query function. As CA_i and QA_{i_j} satisfy $(\varepsilon_{CA_i}, \delta_{CA_i})$-DP and $(\varepsilon_{QA_{i_j}}, \delta_{QA_{i_j}})$-DP respectively, the following holds according to the definition of (ε, δ)-DP:

$$\begin{cases} \Pr[f(CA_i')=t] \leq e^{\varepsilon_{CA_i}} \Pr[f(CA_i'')=t]+\delta_{CA_i}, \\ \Pr[f(QA_{i_j}')=t_j] \leq e^{\varepsilon_{QA_{i_j}}} \Pr[f(QA_{i_j}'')=t_j]+\delta_{QA_{i_j}}, \end{cases} \tag{8}$$

where t and t_j are any possible outputs of $f(\cdot)$ based on CA_i and QA_{i_j}, respectively. $\Pr[f(CA_i') = t]$ is the probability that $f(CA_i')$ outputs result as t.

By using the correlation between CA_i and $\mathbb{Q}_i$ in (4), we can calculate two values of CA_i, i.e., $\overline{CA}_i'$ and $\overline{CA}_i''$, with the obfuscated results $f(\mathbb{Q}_i')$ and $f(\mathbb{Q}_i'')$, respectively. Assuming there exists a function $h(\cdot)$ that satisfies $h(CA_i') = g(f(\mathbb{Q}_i')) = \overline{CA}_i'$ and $h(CA_i'') = g(f(\mathbb{Q}_i'')) = \overline{CA}_i''$. Let T represent a possible output of $h(\cdot)$. Then, we have

$$\begin{cases} \Pr[h(CA_i') = T] = \Pr[g(f(\mathbb{Q}_i')) = T] \\ \Pr[h(CA_i'') = T] = \Pr[g(f(\mathbb{Q}_i'')) = T]. \end{cases} \tag{9}$$

As the considered QAs in $\mathbb{Q}_i$ are independent of each other, We can use the multiplication theorem of probability [69] and express the left-hand side of the (9) as follows:

$$\begin{aligned} \Pr[h(CA_i') = T] &= \Pr[g(f(\mathbb{Q}_i')) = T] \\ &= \Pr[f(QA_{i_1}')=t_1,\cdots,f(QA_{i_x}')=t_x] \\ &= \prod_{j=1}^{x} \Pr[f(QA_{i_j}') = t_j], \end{aligned} \tag{10}$$

where t_j presents the value of QA_{i_j} and $g(t_1, \ldots, t_x) = T$. Similarly, we have

$$\Pr[h(CA_i'') = T] = \prod_{j=1}^{x} \Pr[f(QA_{i_j}'') = t_j]. \tag{11}$$

Following the probability axioms [70] and integrating (8), we can derive:

$$\Pr[f(QA_{i_j}')=t_j] \leq \min\left(1, e^{\varepsilon_{QA_{i_j}}} \Pr[f(QA_{i_j}'')=t_j]\right)+\delta_{QA_{i_j}}. \tag{12}$$

By substituting (12) into the first term on the right-hand side of (10), we have:

$$\begin{aligned} &\Pr[h(CA_i') = T] \\ &\leq \left(\min\left(1, e^{\varepsilon_{QA_{i_1}}} \Pr[f(QA_{i_1}'')=t_1]\right)+\delta_{QA_{i_1}}\right) \times \prod_{j=2}^{x} \Pr[f(QA_{i_j}')=t_j] \\ &\leq \min\left(1, e^{\varepsilon_{QA_{i_1}}} \Pr[f(QA_{i_1}'')=t_1]\right) \prod_{j=2}^{x} \Pr[f(QA_{i_j}')=t_j]+\delta_{QA_{i_1}}. \end{aligned} \tag{13}$$

Here, following the probability axioms [70], $\forall \Pr[f(\cdot)] \leq 1$, so $\prod_{j=2}^{x} \Pr[f(QA_{i_j}') = t_j] \leq 1$ holds.

By repeating the expansion steps of (13) recursively, we can simplify (10) as follows

$$\begin{aligned} \Pr[h(CA_i')=T] &\leq \prod_{j=1}^{x} \min\left(1, e^{\varepsilon_{QA_{i_j}}} \Pr[f(QA_{i_j}'')=t_j]\right)+\sum_{j=1}^{x} \delta_{QA_{i_j}} \\ &\leq \prod_{j=1}^{x} e^{\varepsilon_{QA_{i_j}}} \Pr[f(QA_{i_j}'')=t_j]+\sum_{j=1}^{x} \delta_{QA_{i_j}} \\ &= e^{\sum_{j=1}^{x} \varepsilon_{QA_{i_j}}} \prod_{j=1}^{x} \Pr[f(QA_{i_j}'')=t_j]+\sum_{j=1}^{x} \delta_{QA_{i_j}}. \end{aligned} \tag{14}$$

By substituting (11) into (14), it readily follows that:

$$\Pr[h(CA_i')=T] \leq e^{\sum_{j=1}^{x} \varepsilon_{QA_{i_j}}} \Pr[h(CA_i'')=T]+\sum_{j=1}^{x} \delta_{QA_{i_j}}. \tag{15}$$

By considering the correlation between CA_i and $\mathbb{Q}_i$, i.e., $g(\cdot)$, CA_i satisfies $(\sum_{j=1}^{x} \varepsilon_{QA_{i_j}}, \sum_{j=1}^{x} \delta_{QA_{i_j}})$-DP based on (15). By comparing (8) and (15), it is evident that CA_i would follow the (ε', δ')-DP rather than $(\varepsilon_{CA_i}, \delta')$-DP, where $\varepsilon' = \max(\varepsilon_{CA_i}, \sum_{j=1}^{x} \varepsilon_{QA_{i_j}})$ and $\delta' = \max(\delta_{CA_i}, \sum_{j=1}^{x} \delta_{QA_{i_j}})$. Theorem 1 is proved. □

Appendix B. Proof of Theorem 2

Proof. Since the QAs are interdependent, their values can influence each other. We can use conditional probabilities and define the xth variable Z_x:

$$Z_x = \ln \frac{\Pr[f(QA_{i_x}')=t_x | f(QA_{i_1}')=t_1,\cdots,f(QA_{i_{x-1}}')=t_{x-1}]}{\Pr[f(QA_{i_x}'')=t_x | f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_{x-1}}'')=t_{x-1}]}, \tag{16}$$

which is a martingale [71] representing the difference of the probabilities that $f(QA_{i_x}')$ and $f(QA_{i_x}'')$ output the same value t_x under the influence by the others.

Building upon the assumptions and notation introduced in the proof of Theorem 1 and the chain rule of joint probability distribution [72], we can have

$$\begin{aligned} &\sum_{i=1}^{n} Z_x \\ &= \sum_{i=1}^{n} \ln \frac{\Pr[f(QA_{i_x}')=t_x | f(QA_{i_1}')=t_1,\cdots,f(QA_{i_{x-1}}')=t_{x-1}]}{\Pr[f(QA_{i_x}'')=t_x | f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_{x-1}}'')=t_{x-1}]} \\ &= \ln \frac{\prod_{i=1}^{x} \Pr[f(QA_{i_x}')=t_x | f(QA_{i_1}')=t_1,\cdots,f(QA_{i_{x-1}}')=t_{x-1}]}{\prod_{i=1}^{x} \Pr[f(QA_{i_x}'')=t_x | f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_{x-1}}'')=t_{x-1}]}. \end{aligned} \tag{17}$$

According to the multiplication theorem of probability [69] and the assumption of $h(\cdot)$, (17) can be rewritten as

$$\begin{aligned} \sum_{i=1}^{n} Z_x &= \ln \frac{\Pr[f(QA_{i_1}')=t_1,\cdots,f(QA_{i_x}')=t_x]}{\Pr[f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_x}'')=t_x]} \\ &= \ln \frac{\Pr[h(CA_i')=T]}{\Pr[h(CA_i'')=T]}. \end{aligned} \tag{18}$$

As $f(\cdot)$ satisfies (ε, δ)-DP for QAs, the following holds:

$$\begin{aligned} |Z_x| &= \left| \ln \frac{\Pr[f(QA_{i_x}')=t_x | f(QA_{i_1}')=t_1,\cdots,f(QA_{i_{x-1}}')=t_{x-1}]}{\Pr[f(QA_{i_x}'')=t_x | f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_{x-1}}'')=t_{x-1}]} \right| \\ &\leq \varepsilon. \end{aligned} \tag{19}$$

Based on the feature of martingale [71], we have

$$\begin{aligned} &\mathbb{E}[Z_x | Z_1 = z_1,\cdots,Z_{x-1} = Z_{x-1}] \\ &= \mathbb{E}\left[\ln \frac{\Pr[f(QA_{i_x}')=t_x | f(QA_{i_1}')=t_1,\cdots,f(QA_{i_{x-1}}')=t_{x-1}]}{\Pr[f(QA_{i_x}'')=t_x | f(QA_{i_1}'')=t_1,\cdots,f(QA_{i_{x-1}}'')=t_{x-1}]} \right] \\ &\leq \varepsilon(e^{\varepsilon} - 1), \end{aligned} \tag{20}$$

which is proved in Appendix B.1.

According to (19) and (20), Z_x satisfies the conditions of Azuma's inequality [73]. Using Azuma's inequality, we have: for $\forall y \geq 0$ the following inequality holds,

$$\Pr\left[\sum_{i=1}^{n} Z_x > x\varepsilon(e^{\varepsilon} - 1) + y\varepsilon\sqrt{x} \right] \leq e^{-y^2/2}, \tag{21}$$

where y is an arbitrary positive real number.

Let $\hat{\delta} = e^{-y^2/2}$ and $\hat{\varepsilon} = \sqrt{2x \ln(1/\hat{\delta})}\varepsilon + x\varepsilon(e^{\varepsilon} - 1)$. Then, (21) can be rewritten as

$$\Pr\left[\sum_{i=1}^{n} Z_x > \hat{\varepsilon} \right] \leq \hat{\delta}, \tag{22}$$

which can be rewritten as

$$\Pr\left[\frac{\Pr[h(CA_i') = T]}{\Pr[h(CA_i'') = T]} > e^{\hat{\varepsilon}} \right] \leq \hat{\delta}, \tag{23}$$

which means, considering the correlation between CA_i and $\mathbb{Q}_i$, CA_i satisfies $(\hat{\varepsilon}, \hat{\delta})$-DP, where $\hat{\delta} = e^{-y^2/2}$ and $\hat{\varepsilon} = \sqrt{2x \ln(1/\hat{\delta})}\varepsilon + x\varepsilon(e^{\varepsilon} - 1)$.

Since $-y^2 < 0$, $\hat{\delta} = e^{-y^2/2} \in (0,1)$. In contrast, CA_i is designed to satisfy $(\varepsilon_{CA_i}, \delta_{CA_i})$-DP, as shown in (8). Therefore, CA_i would follow the (ε', δ')-DP, rather than $(\varepsilon_{CA_i}, \delta_{CA_i})$-DP, where $\varepsilon' = \max(\varepsilon_{CA_i}, \hat{\varepsilon})$ and $\delta' = \max(\delta_{CA_i}, \hat{\delta})$. Theorem 2 is proved. □

B.1. The upper boundary of expected value

Proof. To find the upper bound of the following formula

$$\mathbb{E}[\ln \frac{\Pr[f(QA'_{i_x})=t_x | f(QA'_{i_1})=t_1,\cdots,f(QA'_{i_{x-1}})=t_{x-1}]}{\Pr[f(QA''_{i_x})=t_x | f(QA''_{i_1})=t_1,\cdots,f(QA''_{i_{x-1}})=t_{x-1}]}], \tag{24}$$

we simplify the numerator and denominator of the above equation by

$$\begin{cases} A := \left(f(QA'_{i_x})=t_x | f(QA'_{i_1})=t_1,\cdots,f(QA'_{i_{x-1}})=t_{x-1} \right), \\ B := \left(f(QA''_{i_x})=t_x | f(QA''_{i_1})=t_1,\cdots,f(QA''_{i_{x-1}})=t_{x-1} \right). \end{cases} \tag{25}$$

Then, all we ask for is the upper bound of $\mathbb{E}[\ln \frac{\Pr[A]}{\Pr[B]}]$.

According to the law of total expectation and log sum inequality, we have

$$\begin{aligned} &\mathbb{E}[\ln \frac{\Pr[A]}{\Pr[B]}] = \sum_y (\Pr[B=y] \cdot \ln \frac{\Pr[A=y]}{\Pr[B=y]}) \\ &\geq (\sum_y \Pr[B=y]) \cdot \ln \frac{\sum_y \Pr[A=y]}{\sum_y \Pr[B=y]} = 1 \cdot \ln \frac{1}{1} = 0, \end{aligned} \tag{26}$$

where y is any possible term that has the same expression of A and B. Similarly, we have $\mathbb{E}[\ln \frac{\Pr[B]}{\Pr[A]}] \geq 0$.

Therefore, through the operation of combining like terms, the following inequality holds

$$\begin{aligned} \mathbb{E}[\ln \frac{\Pr[A]}{\Pr[B]}] \leq & \mathbb{E}[\ln \frac{\Pr[A]}{\Pr[B]}] + \mathbb{E}[\ln \frac{\Pr[B]}{\Pr[A]}] \\ =& \sum_y (\Pr[B=y] \cdot \ln \frac{\Pr[A=y]}{\Pr[B=y]}) \\ &+ \sum_y (\Pr[A=y] \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]}) \\ =& \sum_y (\Pr[B=y] \cdot \ln \frac{\Pr[A=y]}{\Pr[B=y]} \\ &+ \Pr[A=y] \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]}) \\ =& \sum_y (\Pr[B=y] \cdot (\ln \frac{\Pr[A=y]}{\Pr[B=y]} + \ln \frac{\Pr[B=y]}{\Pr[A=y]})) \\ &+ \sum_y (\Pr[A=y] - \Pr[B=y]) \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]} \\ =& \sum_y (\Pr[B=y] \cdot \ln \frac{\Pr[A=y] \cdot \Pr[B=y]}{\Pr[B=y] \cdot \Pr[A=y]}) \\ &+ \sum_y (\Pr[A=y] - \Pr[B=y]) \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]} \\ =& 0 + \sum_y (\Pr[A=y] - \Pr[B=y]) \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]} \\ =& \sum_y (\Pr[A=y] - \Pr[B=y]) \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]}. \end{aligned} \tag{27}$$

According to the definition of DP, we have

$$\begin{cases} \ln \frac{\Pr[B=y]}{\Pr[A=y]} \leq \varepsilon, \\ \ln \frac{\Pr[A=y]}{\Pr[B=y]} \leq \varepsilon, \end{cases} \tag{28}$$

with which the following holds

$$\Pr[A=y] \leq e^{\varepsilon} \cdot \Pr[B=y]. \tag{29}$$

By substituting (28) and (29) into the (27), it readily follows that

$$\begin{aligned} \mathbb{E}[\ln \frac{\Pr[A]}{\Pr[B]}] \leq & \sum_y (\Pr[A=y] - \Pr[B=y]) \cdot \ln \frac{\Pr[B=y]}{\Pr[A=y]} \\ \leq & \sum_y (e^{\varepsilon} \cdot \Pr[B=y] - \Pr[B=y]) \cdot \varepsilon \\ =& \sum_y \varepsilon (e^{\varepsilon} - 1) \Pr[B=y] \\ =& \varepsilon (e^{\varepsilon} - 1) \sum_y \Pr[B=y] \\ =& \varepsilon (e^{\varepsilon} - 1). \end{aligned} \tag{30}$$

Hence, (20) holds. This proof is completed. □

Data availability

No data was used for the research described in the article.

References

[1] Li Y, Liu Y, Li B, Wang W, Liu N. Towards practical differential privacy in data analysis: Understanding the effect of epsilon on utility in private erm. Comput Secur 2023;128:103147.

[2] Chamikara MAP, Bertok P, Khalil I, Liu D, Camtepe S. Privacy preserving face recognition utilizing differential privacy. Comput Secur 2020;97:101951.

[3] Liu B, Ding M, Shaham S, Rahayu W, Farokhi F, Lin Z. When machine learning meets privacy: A survey and outlook. ACM Comput Surv 2021;54(2):1–36.

[4] Dwork C. Differential privacy. In: Automata, languages and programming: 33rd international colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, proceedings, part II 33. Springer; 2006, p. 1–12.

[5] Zhao Y, Chen J. A survey on differential privacy for unstructured data content. ACM Comput Surv 2022;54(10s):1–28.

[6] Yu G, Wang X, Sun C, Yu P, Ni W, Liu RP. Obfuscating the dataset: Impacts and applications. ACM Trans Intell Syst Technol 2023. Just Accepted.

[7] Kim JW, Edemacu K, Kim JS, Chung YD, Jang B. A survey of differential privacy-based techniques and their applicability to location-based services. Comput Secur 2021;111:102464.

[8] Yu P, Ni W, Yu G, Zhang H, Liu RP, Wen Q. Efficient anonymous data authentication for vehicular ad hoc networks. Secur Commun Netw 2021;2021(1):6638453.

[9] Wen C, Yang J, Gan L, Pan Y. Big data driven Internet of Things for credit evaluation and early warning in finance. Future Gener Comput Syst 2021;124:295–307.

[10] Shao C, Yang Y, Juneja S, GSeetharam T. IoT data visualization for business intelligence in corporate finance. Inf Process Manage 2022;59(1):102736.

[11] Zhao J, Chen Y, Zhang W. Differential privacy preservation in deep learning: Challenges, opportunities and solutions. IEEE Access 2019;7:48901–11.

[12] Ma B, Wang X, Ni W, Liu RP. Personalized location privacy with road network-indistinguishability. IEEE Trans Intell Transp Syst 2022;23(11):20860–72.

[13] Johnson N, Near JP, Song D. Towards practical differential privacy for SQL queries. Proc VLDB Endow 2018;11(5):526–39.

[14] Mao Y, Ye Q, Wang Q, Hu H. Differential privacy for time series: A survey. IEEE Data Eng Bull 2024;47(2):67–92.

[15] Li H, Wang Y, Guo F, Wang J, Wang B, Wu C. Differential privacy location protection method based on the Markov model. Wirel Commun Mob Comput 2021;2021(1):4696455.

[16] Li Z, Pei Q, et al. Location privacy violation via GPS-agnostic smart phone car tracking. IEEE Trans Veh Technol 2018;67(6):5042–53.

[17] Gu H, Plagemann T, Benndorf M, Goebel V, Koldehofe B. Differential privacy for protecting private patterns in data streams. In: 2023 IEEE 39th international conference on data engineering workshops. ICDEW, IEEE; 2023, p. 118–24.

[18] Yao L, Chen Z, Hu H, Wu G, Wu B. Privacy preservation for trajectory publication based on differential privacy. ACM Trans Intell Syst Technol (TIST) 2022;13(3):1–21.

[19] Dwork C, Kenthapadi K, McSherry F, Mironov I, Naor M. Our data, ourselves: Privacy via distributed noise generation. In: Adv. in cryptology-EUROCRYPT 2006: 24th annu. international conference on the theory and applications of cryptographic techniques, st. Petersburg, Russia, May 28-June 1, 2006. proceedings 25. Springer; 2006, p. 486–503.

[20] Friha O, Ferrag MA, Benbouzid M, Berghout T, Kantarci B, Choo K-KR. 2DF-IDS: Decentralized and differentially private federated learning-based intrusion detection system for industrial IoT. Comput Secur 2023;127:103097.

[21] Sarathy R, Muralidhar K. Evaluating Laplace noise addition to satisfy differential privacy for numeric data. Trans Data Priv 2011;4(1):1–17.
[22] Hassan MU, Rehmani MH, Chen J. Differential privacy techniques for cyber physical systems: A survey. IEEE Commun Surv Tutor 2020;22(1):746–89.
[23] Zhao P, Zhang G, Wan S, Liu G, Umer T. A survey of local differential privacy for securing internet of vehicles. J Supercomput 2020;76:8391–412.
[24] Kumar GS, Premalatha K, Maheshwari GU, Kanna PR, Vijaya G, Nivaashini M. Differential privacy scheme using Laplace mechanism and statistical method computation in deep neural network for privacy preservation. Eng Appl Artif Intell 2024;128:107399.
[25] Dwork C, Roth A, et al. The algorithmic foundations of differential privacy. Found Trends Theor Comput Sci 2014;9(3–4):211–407.
[26] Liu F. Generalized Gaussian mechanism for differential privacy. IEEE Trans Knowl Data Eng 2018;31(4):747–56.
[27] Ma B, Lin X, Wang X, et al. New Cloaking Region obfuscation for road network-indistinguishability and location privacy. In: Proceedings of the 25th international symposium on research in attacks, intrusions and defenses. 2022, p. 160–70.
[28] Qi X, Ma X, Bai X, Li W. Differential privacy preserving data publishing based on Bayesian network. In: 2020 IEEE 19th international conference on trust, security and privacy in computing and communications. TrustCom, IEEE; 2020, p. 1718–26.
[29] Wang N, Xiao X, Yang Y, Zhao J, Hui SC, Shin H, Shin J, Yu G. Collecting and analyzing multidimensional data with local differential privacy. In: 2019 IEEE 35th international conference on data engineering. ICDE, IEEE; 2019, p. 638–49.
[30] Cai S, Lyu X, Li X, Ban D, Zeng T. A trajectory released scheme for the internet of vehicles based on differential privacy. IEEE Trans Intell Transp Syst 2021;23(9):16534–47.
[31] Wang H, Wang H. Differentially private publication for correlated non-numerical data. Comput J 2022;65(7):1726–39.
[32] Qu Y, Yu S, Zhou W, Chen S, Wu J. Customizable reliable privacy-preserving data sharing in cyber-physical social networks. IEEE Trans Netw Sci Eng 2020;8(1):269–81.
[33] Tschantz MC, Sen S, Datta A. SoK: Differential privacy as a causal property. In: 2020 IEEE symposium on security and privacy. SP, IEEE; 2020, p. 354–71.
[34] Blanco-Justicia A, Sánchez D, Domingo-Ferrer J, Muralidhar K. A critical review on the use (and misuse) of differential privacy in machine learning. ACM Comput Surv 2022;55(8):1–16.
[35] Hayes J, Balle B, Mahloujifar S. Bounding training data reconstruction in DP-SGD. Adv Neural Inf Process Syst 2024;36.
[36] Kifer D, Machanavajjhala A. Pufferfish: A framework for mathematical privacy definitions. ACM Trans Database Syst 2014;39(1):1–36.
[37] Song S, Wang Y, Chaudhuri K. Pufferfish privacy mechanisms for correlated data. In: Proceedings of the 2017 ACM international conference on management of data. 2017, p. 1291–306.
[38] Kifer D, Machanavajjhala A. No free lunch in data privacy. In: Proceedings of the 2011 ACM SIGMOD international conference on management of data. 2011, p. 193–204.
[39] Yang B, Sato I, Nakagawa H. Bayesian differential privacy on correlated data. In: Proceedings of the 2015 ACM SIGMOD international conference on management of data. 2015, p. 747–62.
[40] Liu C, Chakraborty S, Mittal P. Dependence makes you vulnerable: Differential privacy under dependent tuples. In: NDSS. vol. 16, 2016, p. 21–4.
[41] Wang J, Li Z, Lui JC, Sun M. Topology-theoretic approach to address attribute linkage attacks in differential privacy. Comput Secur 2022;113:102552.
[42] He Y, Strohmer T, Vershynin R, Zhu Y. Differentially private low-dimensional representation of high-dimensional data. 2023, arXiv preprint arXiv:2305.17148.
[43] Ren X, Yu C-M, Yu W, Yang S, Yang X, McCann JA, Philip SY. LoPub: High-dimensional crowdsourced data publication with local differential privacy. IEEE Trans Inf Forensics Secur 2018;13(9):2151–66.
[44] Zheng Z, Wang T, Wen J, Mumtaz S, Bashir AK, Chauhdary SH. Differentially private high-dimensional data publication in internet of things. IEEE Internet Things J 2020;7(4):2640–50.
[45] Zhang T, Zhu T, Xiong P, Huo H, Tari Z, Zhou W. Correlated differential privacy: Feature selection in machine learning. IEEE Trans Ind Inf 2019;16(3):2115–24.
[46] Elton EJ, Gruber MJ, Blake CR. A first look at the accuracy of the CRSP mutual fund database and a comparison of the CRSP and morningstar mutual fund databases. J Financ 2001;56(6):2415–30.
[47] Zheng Y, Li Q, Chen Y, Xie X, Ma W-Y. Understanding mobility based on GPS data. In: Proc. of the 10th int. conf. ubi. comput.. 2008, p. 312–21.
[48] Wen C, Yang J, Gan L, Pan Y. Big data driven Internet of Things for credit evaluation and early warning in finance. Future Gener Comput Syst 2021;124:295–307.
[49] Liu F, Li T. A clustering K-anonymity privacy-preserving method for wearable IoT devices. Secur Commun Netw 2018;2018(1):4945152.
[50] Gupta M, Awaysheh FM, Benson J, Alazab M, Patwa F, Sandhu R. An attribute-based access control for cloud enabled industrial smart vehicles. IEEE Trans Ind Inform 2020;17(6):4288–97.
[51] Yan Y, Wei C, Guo X, Lu X, Zheng X, Liu Q, Zhou C, Song X, Zhao B, Zhang H, et al. Confidentiality support over financial grade consortium blockchain. In: Proceedings of the 2020 ACM SIGMOD international conference on management of data. 2020, p. 2227–40.
[52] Chen S, Fu A, Yu S, Ke H, Su M. DP-QIC: A differential privacy scheme based on quasi-identifier classification for big data publication. Soft Comput 2021;25:7325–39.
[53] Wei K, Li J, Ding M, Ma C, Yang HH, Farokhi F, Jin S, Quek TQ, Poor HV. Federated learning with differential privacy: Algorithms and performance analysis. IEEE Trans Inf Forensics Secur 2020;15:3454–69.
[54] Niu B, Chen Y, Wang B, Wang Z, Li F, Cao J. AdaPDP: Adaptive personalized differential privacy. In: IEEE INFOCOM 2021-IEEE conference on computer communications. IEEE; 2021, p. 1–10.
[55] Ahmed M, Choudhury N, Uddin S. Anomaly detection on big data in financial markets. In: Proceedings of the 2017 IEEE/ACM international conference on advances in social networks analysis and mining 2017. 2017, p. 998–1001.
[56] Qureshi KN, Din S, Jeon G, Piccialli F. Internet of vehicles: Key technologies, network model, solutions and challenges with future aspects. IEEE Trans Intell Transp Syst 2020;22(3):1777–86.
[57] McLaney E, Atrill P. Accounting and finance: An introduction. Pearson; 2020.
[58] Pan T, Pedrycz W, Cui J, Yang J, Wu W. Interpretability of neural networks with probability density functions. Adv Theory Simul 2022;5(3):2100459.
[59] Karlik B, Olgac AV. Performance analysis of various activation functions in generalized MLP architectures of neural networks. Int J Artif Intell Expert Syst 2011;1(4):111–22.
[60] Albawi S, Mohammed TA, Al-Zawi S. Understanding of a convolutional neural network. In: 2017 international conference on engineering and technology. ICET, IEEE; 2017, p. 1–6.
[61] Pascanu R, Mikolov T, Bengio Y. On the difficulty of training recurrent neural networks. In: International conference on machine learning. ICML, PMLR; 2013, p. 1310–8.
[62] Sainath TN, Vinyals O, Senior A, Sak H. Convolutional, long short-term memory, fully connected deep neural networks. In: 2015 IEEE international conference on acoustics, speech and signal processing. ICASSP, IEEE; 2015, p. 4580–4.
[63] Chaudhari S, Mithal V, Polatkan G, Ramanath R. An attentive survey of attention models. ACM Trans Intell Syst Technol 2021;12(5):1–32.
[64] Jin J, McMurtry E, Rubinstein BI, Ohrimenko O. Are we there yet? timing and floating-point attacks on differential privacy systems. In: 2022 IEEE symposium on security and privacy. SP, IEEE; 2022, p. 473–88.
[65] Kabir E, Craig L, Mehnaz S. Disparate privacy vulnerability: Targeted attribute inference attacks and defenses. In: Proceedings of the 34th USeNIX security symposium. USENIX Security 2025, 2025.
[66] Chen L, Yue D, Ding X, Wang Z, Choo K-KR, Jin H. Differentially private deep learning with dynamic privacy budget allocation and adaptive optimization. IEEE Trans Inf Forensics Secur 2023.
[67] Zhang K, Tian J, Xiao H, Zhao Y, Zhao W, Chen J. A numerical splitting and adaptive privacy budget-allocation-based LDP mechanism for privacy preservation in blockchain-powered IoT. IEEE Internet Things J 2022;10(8):6733–41.
[68] Liu W, Cao B, Peng M. Web3 technologies: Challenges and opportunities. IEEE Netw 2023.
[69] Bertsekas D, Tsitsiklis JN. Introduction to probability, vol. 1, Athena Scientific; 2008.
[70] Kairouz P, Oh S, Viswanath P. The composition theorem for differential privacy. In: International conference on machine learning. PMLR; 2015, p. 1376–85.
[71] Hall P, Heyde CC. Martingale limit theory and its application. Academic Press; 2014.
[72] Wang L, Zhao H. Learning a flexible K-dependence Bayesian dlassifier from the chain rule of joint probability distribution. Entropy 2015;17(6):3766–86.
[73] Chung F, Lu L. Concentration inequalities and martingale inequalities: A survey. Internet Math 2006;3(1):79–127.