electronics

Article

DPShield: Optimizing Differential Privacy for High-Utility Data Analysis in Sensitive Domains

Pratik Thantharate [1], Shyam Bhojwani [1] and Anurag Thantharate [2,*]

[1] SUNY Binghamton, Jersey City, NJ 07304, USA
[2] School of Computing and Engineering, University of Missouri, Kansas City, MO 64112, USA
* Correspondence: adtmv7@mail.umkc.edu

Abstract: The proliferation of cloud computing has amplified the need for robust privacy-preserving technologies, particularly when dealing with sensitive financial and human resources (HR) data. However, traditional differential privacy methods often struggle to balance rigorous privacy protections with maintaining data utility. This study introduces DPShield, an optimized adaptive framework that enhances the trade-off between privacy guarantees and data utility in cloud environments. DPShield leverages advanced differential privacy techniques, including dynamic noise-injection mechanisms tailored to data sensitivity, cumulative privacy loss tracking, and domain-specific optimizations. Through comprehensive evaluations on synthetic financial and real-world HR datasets, DPShield demonstrated a remarkable 21.7% improvement in aggregate query accuracy over existing differential privacy approaches. Moreover, it maintained machine learning model accuracy within 5% of non-private benchmarks, ensuring high utility for predictive analytics. These achievements signify a major advancement in differential privacy, offering a scalable solution that harmonizes robust privacy assurances with practical data analysis needs. DPShield's domain adaptability and seamless integration with cloud architectures underscore its potential as a versatile privacy-enhancing tool. This work bridges the gap between theoretical privacy guarantees and practical implementation demands, paving the way for more secure, ethical, and insightful data usage in cloud computing environments.

Keywords: differential privacy; privacy optimization; data utility; sensitive data analysis; privacy parameter tuning; privacy guarantee; predictive analytics; machine learning

Citation: Thantharate, P.; Bhojwani, S.; Thantharate, A. DPShield: Optimizing Differential Privacy for High-Utility Data Analysis in Sensitive Domains. *Electronics* **2024**, *13*, 2333. https://doi.org/10.3390/electronics13122333

Academic Editors: Jianlong Zhou and Niusha Shafiabady

Received: 3 April 2024
Revised: 25 May 2024
Accepted: 12 June 2024
Published: 14 June 2024

1. Introduction

The proliferation of cloud computing has fundamentally altered the landscape of data analytics, ushering in an era where vast amounts of financial and human resources (HR) data can be processed to glean transformative insights. However, this capability introduces significant privacy challenges, particularly when it involves handling sensitive information. Traditional privacy-preserving methods often find themselves at a crossroads, struggling to strike a balance between ensuring robust privacy protections and maintaining the utility of data. This dilemma is further compounded in cloud environments, where data storage, access, and processing dynamics amplify privacy concerns.

Differential privacy stands as a cornerstone in modern data protection, offering a mathematically sound framework that ensures the confidentiality of individual data within large datasets. This framework distinguishes itself from ad hoc or heuristic approaches by providing explicit privacy guarantees. It meticulously reduces the risk that comes with data sharing, ensuring that the inclusion of any individual's information does not substantially elevate the likelihood of privacy breaches. At its core, differential privacy allows for disseminating aggregate data insights while safeguarding sensitive individual details. This is achieved through the strategic injection of noise into query responses or machine learning models before their release, thereby masking the influence of any single data point. Central to the concept of differential privacy are principles such as the computation of total privacy loss using specific metrics like the privacy budget (ϵ) and the fine-tuning of noise addition to balance model accuracy with privacy needs. These principles are grounded in rigorous mathematical proofs, positioning differential privacy as a leading method for conducting data analysis that is both ethical and secure. The framework's effectiveness is quantified by parameters including the privacy budget (ϵ) and the probability of privacy guarantee failure (δ). These parameters facilitate a nuanced negotiation between preserving privacy and maintaining the utility of the data analysis. However, applying these principles in real-world scenarios presents challenges, especially in calibrating these parameters to achieve the desired balance without compromising the integrity of privacy assurances or diluting the insights derived from data analysis.

In the face of these challenges, this paper introduces DPShield, a novel framework designed to optimize differential privacy for enhanced data utility in cloud-based analytics, specifically targeting financial and HR datasets. DPShield innovatively leverages advanced differential privacy techniques and domain-specific optimizations to improve data analysis outcomes' accuracy significantly. Our comprehensive evaluation demonstrates that DPShield not only achieves a notable improvement in query accuracy by 21.7% over existing differential privacy mechanisms, but also ensures that machine learning model accuracy remains commendably close, within 5% of non-private benchmarks. These advancements herald DPShield as a pivotal solution capable of enabling secure, efficient, and ethical data analysis practices in cloud computing settings.

The motivation behind DPShield stems from a critical examination of the existing differential privacy landscape, where a gap between theoretical models and their practical applicability persists. Many current frameworks offer robust privacy guarantees, but at the cost of significantly diminished data utility, rendering the outcomes less valuable for meaningful analysis. Conversely, efforts to enhance data utility often inadvertently compromise privacy protections, especially in scenarios involving multiple queries or interactive data analysis sessions. DPShield's development was driven by the need to address these challenges, aiming to provide a balanced, flexible solution that adapts to the nuanced requirements of different data types and analytical contexts.

Moreover, the relevance of differential privacy is increasingly magnified in the era of cloud computing, where data are abundant, more fluid, and more interconnected. The cloud environment introduces unique challenges in data privacy and security management, necessitating solutions that are robust and scalable and adaptable to the evolving nature of cloud architectures and services. In this light, DPShield represents a significant step forward, offering a comprehensive framework that integrates seamlessly with cloud-based data analytics pipelines, ensuring privacy without compromising on the depth and quality of insights derived.

DPShield introduces several key innovations that distinguish it from existing differential privacy frameworks. First, it employs an Adaptive Laplace Mechanism that dynamically adjusts noise levels based on query sensitivity and the remaining privacy budget, optimizing accuracy while maintaining privacy guarantees. Second, it incorporates advanced techniques like the Moment Accountant for cumulative privacy loss tracking across multiple queries. Third, DPShield features domain-specific customization's such as a Markov Quilt Mechanism to handle correlated data attributes common in financial and HR domains. Finally, it offers a flexible, modular architecture that can selectively activate privacy-enhancing components based on computational budgets, enabling seamless integration into diverse cloud analytics pipelines.

The structure of this paper is as follows: Section 2 delves into the related work, providing a critical overview of the differential privacy landscape and identifying the specific challenges and gaps that DPShield addresses. Section 3 outlines the methodology behind DPShield, detailing its architectural innovations and the theoretical underpinnings of its differential privacy optimizations. Section 4 presents an in-depth evaluation of DPShield, showcasing its performance across various datasets and analytical scenarios and highlighting its advantages over traditional differential privacy approaches. Finally, Section 5 discusses the broader implications of our findings, explores potential avenues for future research, and concludes the paper with reflections on DPShield's contributions to the field of data privacy and cloud computing.

Through this study, we contribute a novel perspective and a robust solution to the ongoing discourse on balancing privacy and utility in cloud-based data analytics. DPShield bridges the gap between theoretical privacy protections and practical analytical needs and sets a new benchmark for applying differential privacy in real-world scenarios, facilitating more secure, responsible, and effective use of sensitive data in the digital age.

2. Literature Review

The literature on differentially private data frameworks spans a broad spectrum of design principles, techniques, and applications, reflecting the growing importance of privacy-preserving data analysis in various domains. This review synthesizes key contributions, evaluates the effectiveness of these frameworks, and identifies areas requiring further investigation.

The field of differential privacy has seen rapid progress in recent years, driven by the increasing need to extract insights from large datasets while ensuring robust privacy protection. Various frameworks and techniques have been proposed, each aiming to address the fundamental trade-off between data utility and individual privacy. However, the current landscape is marked by several open challenges. Many existing solutions struggle to maintain high accuracy for complex analysis tasks involving multiple queries or high-dimensional data. There is also a lack of comprehensive frameworks that can adapt to different data domains and types without extensive manual configuration. Additionally, the long-term management of privacy budgets and the development of intuitive privacy metrics remain areas in need of further research.

2.1. Design Principles and Techniques

Differential privacy has evolved significantly since its formal introduction by Dwork et al. [1]. The foundational principle of differential privacy involves adding noise to the output of queries or algorithms to mask the contribution of individual data points. Various techniques have been developed to implement this principle, each with its own approach to balancing privacy and utility. One common technique is the Laplace mechanism, which adds noise drawn from a Laplace distribution to the query results [2]. Another approach is the Exponential mechanism, suitable for non-numeric queries, which selects outputs based on a probability distribution over the possible outcomes, ensuring that the likelihood of any given outcome is relatively insensitive to changes in any single individual's data [3]. More recently, advanced techniques such as the Gaussian mechanism have been proposed, offering different trade-offs regarding privacy guarantees, and the applicability to complex data types were discussed by Koskela, Antti, et al. [4]. Additionally, the concept of local differential privacy, where noise is added to the data before they are collected, providing privacy guarantees at the individual data point level, has gained traction [5].

2.2. Measuring Privacy Loss

The measurement of privacy loss in differential privacy frameworks is a critical area of research, focusing on developing and refining privacy loss budgets and constraints. The privacy loss budget, denoted by ϵ, quantifies the acceptable level of privacy risk, while δ allows for a small probability of exceeding this budget [6]. Researchers have proposed various methods for calculating and managing these parameters to optimize the trade-off between privacy protection and data utility. For instance, adaptive mechanisms dynamically adjust the noise amount based on the query sensitivity and the remaining privacy budget [7]. This adaptability is crucial for complex analyses involving multiple queries or when working with high-dimensional data.

While measuring and managing privacy loss budgets are crucial aspects of differential privacy, several challenges persist. Existing methods often struggle to provide intuitive mappings between abstract privacy parameters like epsilon and delta and the actual privacy risks perceived by users. There is a need for more interpretable privacy metrics that can effectively communicate the trade-offs between privacy and utility to practitioners. Additionally, many current techniques for privacy budgeting assume a fixed analysis workflow, but in real-world scenarios, analysis tasks are often dynamic and interactive, requiring more adaptive privacy budget management approaches.

2.3. Effectiveness of Existing Frameworks

The effectiveness of differentially private data frameworks is often assessed based on their ability to provide strong privacy guarantees while maintaining a high level of data utility. Studies have shown that, while differential privacy offers robust protection against a wide range of privacy attacks, introducing noise can significantly impact the accuracy of query results [8]. Frameworks such as Google's Differential Privacy Library and the OpenDP initiative by Harvard's Privacy Tools Project exemplify efforts to provide practical, open-source tools for implementing differential privacy and federated learning for non-centralized data learning, which have been discussed by researchers [9–11]. These frameworks have been applied in diverse fields, from healthcare to social science, demonstrating the versatility of differential privacy. However, challenges remain regarding usability, scalability, and handling complex analytical tasks without substantial utility loss.

2.4. Local Differential Privacy and Potential Adaptations

Central to the discussion on differential privacy is the concept of local differential privacy (LDP), which offers an alternative approach to ensuring individual data privacy. In LDP, noise is added directly to individual data points before they are collected or shared with a central server. This decentralized approach alleviates the need to trust a central entity with sensitive data, as the privacy protection is applied at the source. However, LDP also introduces its own challenges, as the increased noise required for individual data points can lead to significant utility loss, particularly for high-dimensional or complex data. Nonetheless, LDP presents an intriguing avenue for exploration, particularly in scenarios where data remain decentralized or hosted locally, such as in edge computing or Internet of Things (IoT) environments.

In the context of DPShield, adapting its mechanisms to operate in an LDP setting could unlock new possibilities for privacy-preserving data analysis in decentralized environments. Techniques like the Adaptive Laplace Mechanism and the Markov Quilt Mechanism could be extended to calibrate noise injection and privacy budget management at the individual data point level. Additionally, the modular architecture of DPShield could facilitate the integration of LDP components, allowing for seamless transitions between central and local differential privacy approaches based on specific deployment requirements. To further illustrate the potential applications of LDP, the authors in [5,12] applied LDP in the healthcare domain for privacy-preserving data collection and analysis. Their approach demonstrated the feasibility of LDP in sensitive domains while highlighting the trade-offs between privacy protection and utility loss. Such insights can inform the adaptation of DPShield to operate in an LDP context, leveraging its innovative techniques to optimize this trade-off dynamically.

2.5. Gaps and Future Research Directions

In previous research [13], we proposed two innovative solutions for cloud-based software testing: a distributed testing framework and a realistic environment simulation framework. These approaches significantly improve testing efficiency, effectiveness, and accuracy, which are essential for enhancing the privacy and reliability of cloud applications, thereby ensuring a secure and seamless user experience in cloud computing environments. Without embedding differential privacy, there remains a risk of exposing sensitive information, undermining the privacy assurances these cloud technologies aim to provide.

Despite the advancements in differential privacy research, several gaps remain. One notable area is the lack of comprehensive frameworks that can easily adapt to different domains and data types without requiring extensive privacy expertise. Furthermore, there is a need for more research on the long-term management of privacy budgets, especially in environments with frequent data queries and updates. Future research should also explore the integration of differential privacy with emerging technologies such as blockchain and federated learning, which present new opportunities and challenges for privacy-preserving data analysis. Additionally, developing more intuitive metrics for privacy loss that practitioners can easily understand and apply is a critical need.

The literature review highlights several key gaps that DPShield aims to address. First, there is a lack of comprehensive frameworks that can easily adapt to diverse data domains and types, often requiring manual configuration by privacy experts. DPShield's modular architecture and domain-specific customizations tackle this issue, enabling non-experts to leverage differential privacy effectively. Second, existing solutions struggle to maintain high utility for complex analysis over prolonged periods with multiple queries. DPShield's adaptive mechanisms, such as dynamic noise injection and cumulative privacy tracking, are designed to optimize this crucial utility–privacy trade-off dynamically. Third, the management of privacy budgets in interactive analysis scenarios is an underexplored area that DPShield's budget-tracking approaches aim to advance. Finally, DPShield introduces innovative techniques like the Markov Quilt Mechanism to handle correlated data attributes common in fields like finance and HR, addressing a gap in prior work.

3. Methodology and Proposed Framework

This section delineates the methodology adopted to assess the privacy guarantees offered by the proposed differentially private data frameworks. Our approach is twofold, encompassing both theoretical analysis and empirical validation, with a particular emphasis on heuristic modeling to simulate real-world data analysis scenarios under differential privacy constraints.

3.1. Proposed Framework: DPShield

We introduce DPShield—an optimized adaptive framework as shown in Figure 1 for differential privacy tailored to financial services and HR data analysis within cloud environments. DPShield employs a modular architecture, leveraging advanced composition theorems and noise-injection techniques while tracking privacy loss quantification. The core of DPShield features an Adaptive Laplace Mechanism, which dynamically adjusts noise levels based on query sensitivity and the remaining privacy budget. Additional optimizations include a Moment Accountant for cumulative privacy loss tracking, exponential smoothing for handling time series data, and a Bayesian inference model for adjusting query priorities based on their sensitivity and utility. The privacy-enhancing components within DPShield can be selectively activated, allowing the framework to align with varying computational budgets. Optional extensions support secure multi-party computation facilitated through homomorphic encryption techniques. Homomorphic encryption allows the computations to be carried out on encrypted data without requiring decryption first.

DPShield employs the Laplace mechanism for adding noise to query outputs. The amount of noise is determined by the sensitivity of the query, which measures the maximum possible change in the query output when a single individual's data are added or removed from the dataset. Formally, the sensitivity of a query function f is defined as:

$$\Delta f = \max_{D_1, D_2} \|f(D_1) - f(D_2)\|_1 \tag{1}$$

where D_1 and D_2 are neighboring datasets differing by a single individual's data and $\| \cdot \|_1$ denotes the L_1-norm [14]. The sensitivity calculation ensures that the noise added is proportional to the potential impact of individual data points on the query output, thus preserving differential privacy.

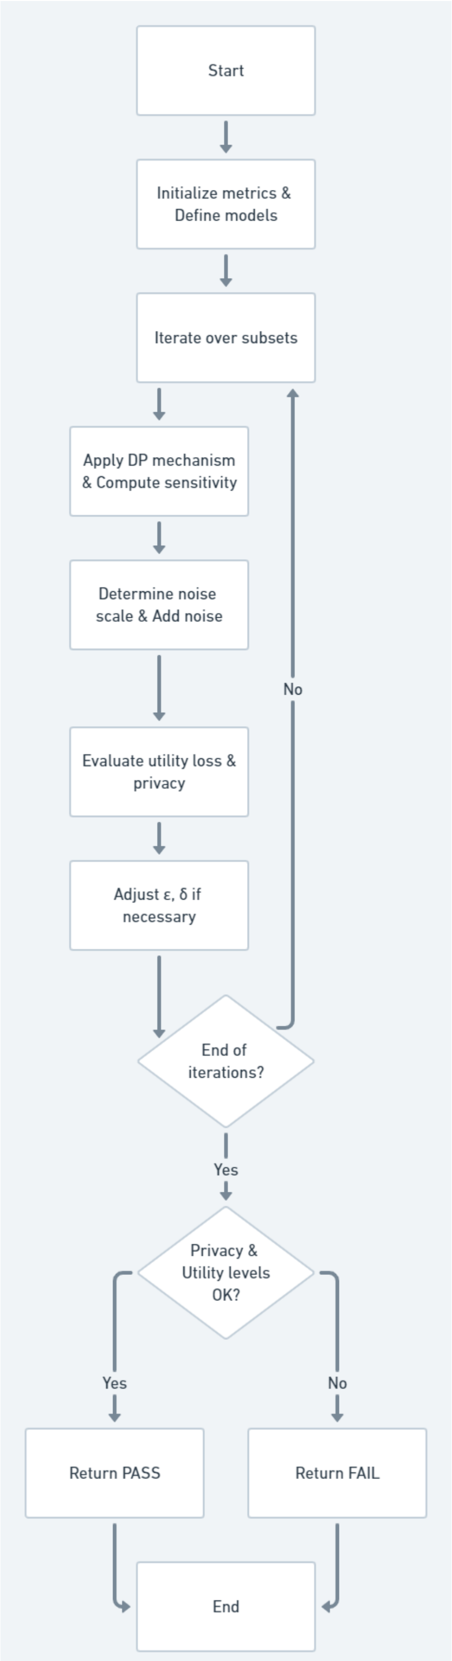
Start
Initialize metrics & Define models
Iterate over subsets
Apply DP mechanism & Compute sensitivity
Determine noise scale & Add noise
Evaluate utility loss & privacy
Adjust ε, δ if necessary
End of iterations?
No
Yes
Privacy & Utility levels OK?
Yes
No
Return PASS
Return FAIL
End

Figure 1. Evaluation of differentially private frameworks.

3.2. Adaptive Laplace Mechanism

The Adaptive Laplace Mechanism (Algorithm 1) dynamically adjusts the noise scale based on query sensitivity and the remaining privacy budget. For a given query function f and dataset D, the sensitivity Δf is first computed using the ComputeSensitivity function. The noise scale λ is then calculated as $\frac{\Delta f}{\epsilon'}$, where ϵ' is the remaining privacy budget. Laplace noise drawn from $\mathrm{Lap}(0, \lambda)$ is added to the query output $f(D)$ to obtain the noisy output Y. Finally, the privacy budget is updated using the composition theorem, subtracting $\frac{\Delta f^2}{2\lambda^2}$ from ϵ'. The algorithm returns the noisy output Y and the updated privacy budget ϵ', which can be used for subsequent queries or analyses. In DPShield, we employ an optimized sensitivity computation technique that leverages sparse vector representations and data partitioning to improve efficiency for high-dimensional datasets without sacrificing accuracy.

Algorithm 1 Adaptive Laplace Mechanism.

```
Require: Query function $f$, dataset $D$, privacy budget $\epsilon$, remaining privacy budget $\epsilon'$
Compute sensitivity $\Delta f \leftarrow \mathrm{ComputeSensitivity}(f, D)$
Calculate noise scale $\lambda \leftarrow \frac{\Delta f}{\epsilon'}$
Add Laplace noise $Y \leftarrow f(D) + \mathrm{Lap}(0, \lambda)$
Update privacy budget $\epsilon' \leftarrow \epsilon' - \frac{\Delta f^2}{2\lambda^2}$
return Noisy output $Y$, updated privacy budget $\epsilon'$
```

3.3. Evaluation Methodology

DPShield's evaluation process involves iterative parameter tuning and heuristic modeling approaches to assess the trade-off between privacy guarantees and data utility. The following algorithms outline the procedures used for these evaluation techniques.

Iterative Parameter Tuning

Algorithm 2 outlines the iterative tuning procedure for determining the optimal privacy budget ϵ that achieves the desired trade-off between privacy and utility. Given a dataset D, an initial privacy budget ϵ_0, a target utility score U_{target}, and a maximum number of iterations T, the algorithm initializes ϵ to ϵ_0 and iteration counter t to 0.

Algorithm 2 Iterative tuning of privacy and utility parameters.

```
Require: Dataset $D$, initial privacy budget $\epsilon_0$, target utility $U_{\text{target}}$, max iterations $T$
Initialize $\epsilon \leftarrow \epsilon_0$, iteration $t \leftarrow 0$
while $t < T$ do
    Apply differential privacy mechanism with budget $\epsilon$ to $D$, obtaining $D_\epsilon$
    Compute utility score $U(D_\epsilon)$ using objective function $\mathcal{O}$
    if $U(D_\epsilon) \geq U_{\text{target}}$ then
        return $\epsilon$ {Target utility achieved}
    else
        Adjust $\epsilon$ based on utility gap $U_{\text{target}} - U(D_\epsilon)$
    end if
    $t \leftarrow t + 1$
end while
return $\epsilon$ {Best attained privacy budget}
```

In each iteration, the differential privacy mechanism is applied to the dataset D using the current privacy budget ϵ, resulting in a privatized version D_ϵ. The utility score $U(D_\epsilon)$ is then computed using an objective function $\mathcal{O}$, which can be tailored to specific analytical tasks or metrics of interest. If the utility score $U(D_\epsilon)$ meets or exceeds the target U_{target}, the algorithm terminates and returns the current privacy budget ϵ, as the desired utility level has been achieved. Otherwise, the privacy budget ϵ is adjusted based on the gap between the target utility U_{target} and the achieved utility $U(D_\epsilon)$. This adjustment can follow various strategies, such as multiplicative or additive updates, depending on the specific requirements. The iterative process continues until either the target utility is met or the maximum number of iterations T is reached. In the latter case, the algorithm returns the best attained privacy budget ϵ that provided the highest utility score within the iteration limit.

3.4. Moment Accountant

DPShield employs the Moments Accountant technique to track the cumulative privacy loss across multiple queries or analysis tasks. The Moments Accountant provides a precise bound on the privacy loss by computing the moments of the privacy loss random variable.

Algorithm 3 outlines the Moments Accountant procedure for computing the bound on the privacy loss parameter ϵ. Given a list of sensitivities Δf_i and noise scales λ_i for a sequence of queries, along with a target δ, the algorithm initializes $\mu_{\max}$ to 0. For each query, it calculates the privacy loss α_i based on the sensitivity and noise scale. The log moment μ_i is then approximated using the formula $\mu_i \approx \alpha_i - \frac{\alpha_i^2}{2}$. The maximum divergence $\mu_{\max}$ is updated to track the largest log moment across all queries. Finally, the algorithm computes the bound on ϵ using the formula $\epsilon(\mu_{\max}, \delta) \leftarrow \mu_{\max} + \sqrt{2\mu_{\max}\log(1/\delta)}$, which provides a tight bound on the cumulative privacy loss. The computed $\epsilon(\mu_{\max}, \delta)$ is then used to determine the appropriate noise levels for subsequent queries, ensuring that the overall privacy loss remains within the desired bounds.

Algorithm 3 Moments Accountant for privacy loss tracking.

```
Require: List of sensitivities $\Delta f_i$, noise scales $\lambda_i$, target $\delta$
Initialize $\mu_{\max} = 0$ {Maximum divergence}
for $i = 1$ to $n$ do
    $\alpha_i \leftarrow \frac{\Delta f_i^2}{\lambda_i^2}$ {Privacy loss of $i$-th query}
    $\mu_i \leftarrow \log\left(\left[e^{\alpha_i \mathcal{N}(0,1)}\right]\right) \approx \alpha_i - \frac{\alpha_i^2}{2}$ {Approximate log moment}
    $\mu_{\max} \leftarrow \max(\mu_{\max}, \mu_i)$
end for
Compute $\epsilon(\mu_{\max}, \delta) \leftarrow \mu_{\max} + \sqrt{2\mu_{\max}\log(1/\delta)}$ {Bound on $\epsilon$}
return $\epsilon(\mu_{\max}, \delta)$
```

3.4.1. Integration of Markov Quilt Mechanism

The Markov Quilt Mechanism is a novel approach introduced in DPShield to handle correlated attributes and generate synthetic data while preserving differential privacy. This mechanism is particularly useful for financial and HR datasets, where attributes often exhibit complex correlations.

Algorithm 4 describes the Markov Quilt Mechanism for generating differentially private synthetic data, which preserves attribute correlations. The algorithm takes as the input the original dataset D, the privacy budget ϵ, and the correlation matrix Σ capturing the attribute correlations. First, the dataset D is partitioned into k disjoint subsets $D_1, D_2, \ldots, D_k$. A Gaussian Copula model $\mathcal{C}$ is then fit to the dataset using the correlation matrix Σ. This model captures the underlying multivariate distribution of the data, including attribute correlations. For each subset D_i, the sensitivity Δf_i is computed, and the noise scale λ_i is determined based on the allocated privacy budget ϵ/k. Correlated Gaussian noise $\mathcal{N}(0, \lambda_i\Sigma)$, scaled by the noise level λ_i and the correlation matrix Σ, is added to the subset D_i to obtain the noisy subset D_i^*. After adding noise to all subsets, the noisy subsets D_i are combined into a single noisy dataset D. Finally, the inverse Gaussian Copula transform $\mathcal{C}^{-1}$ is applied to D^* to generate the synthetic dataset $\tilde{D}$, which preserves the attribute correlations while providing differential privacy guarantees. The generated synthetic dataset $\tilde{D}$ can be used for various data analysis tasks while protecting the privacy of individual records in the original dataset D.

Algorithm 4 Markov Quilt Mechanism for correlated data.

```
Require: Dataset $D$, privacy budget $\epsilon$, correlation matrix $\Sigma$
Partition $D$ into $k$ disjoint subsets $D_1, D_2, \ldots, D_k$
Fit a Gaussian Copula model $\mathcal{C}$ to $D$ using $\Sigma$
for $i = 1$ to $k$ do
    Compute sensitivity $\Delta f_i$ for subset $D_i$
    Determine noise scale $\lambda_i \leftarrow \frac{\Delta f_i}{\epsilon / k}$
    Add noise $D_i^* \leftarrow D_i + \mathcal{N}(0, \lambda_i \Sigma)$ {Correlated Gaussian noise}
end for
Combine noisy subsets $D^* \leftarrow \bigcup_{i=1}^{k} D_i$
Generate synthetic data $\tilde{D} \leftarrow \mathcal{C}^{-1}(D)$ {Inverse Gaussian Copula transform}
return $\tilde{D}$
```

3.4.2. Heuristic Modeling Approach

The heuristic modeling approach forms the cornerstone of our methodology, facilitating the simulation of various privacy-preserving data analysis scenarios. This approach involves constructing heuristic models that approximate the behavior of differentially private mechanisms under a range of conditions, including varying privacy loss budgets and dataset characteristics. We introduce the "Evaluation of Differentially Private Frameworks" Algorithm 5 to evaluate the privacy guarantees of our proposed frameworks systematically. This algorithm is instrumental in our heuristic modeling approach, enabling the simulation and assessment of differential privacy mechanisms across financial and HR datasets.

Algorithm 5 Evaluation of Differentially Private Frameworks.

```
Require: Financial dataset $F$, HR dataset $H$, privacy parameters $\epsilon, \delta$, max iterations $T$
Ensure: Evaluation outcome indicating privacy guarantee levels
Initialize evaluation metrics: privacy guarantee level $\mathcal{P}$, utility level $\mathcal{U}$
Define heuristic models for $F$ and $H$ based on typical analytical tasks
for $t = 1$ to $T$ do
    Select a random subset $F_{\text{sub}}$ from $F$ and $H_{\text{sub}}$ from $H$
    Apply differential privacy mechanism (e.g., Laplace, Gaussian) to $F_{\text{sub}}$ and $H_{\text{sub}}$
    Compute sensitivity $\Delta F$ for $F_{\text{sub}}$ and $\Delta H$ for $H_{\text{sub}}$ {Sensitivity calculation methods}
    Determine noise scale $\sigma_F$ and $\sigma_H$ based on $\Delta F$, $\Delta H$, $\epsilon$, and $\delta$
    Add noise to $F_{\text{sub}}$ and $H_{\text{sub}}$ to generate $F_{\text{dp}}$ and $H_{\text{dp}}$
    Evaluate utility loss for $F_{\text{dp}}$ and $H_{\text{dp}}$ {Utility loss evaluation metrics}
    Update $\mathcal{P}$ and $\mathcal{U}$ based on evaluation
    if privacy guarantee for $F_{\text{dp}}$ or $H_{\text{dp}}$ falls below threshold then
        Adjust $\epsilon$ and $\delta$ for subsequent iterations
    end if
end for
Analyze overall $\mathcal{P}$ and $\mathcal{U}$ to determine if privacy can be ensured within the desired
levels
if desired privacy and utility levels are met then
    return PASS
else
    return FAIL
end if
```

This algorithm underscores our methodological emphasis on adaptability and precision in evaluating differential privacy implementations. By systematically addressing each component of our methodology, we lay the groundwork for a comprehensive evaluation of DPShield, demonstrating its effectiveness in enhancing privacy guarantees and data utility for cloud-based analytics.

3.4.3. Formalization of Privacy Loss Budgets and Constraints

Central to our evaluation is the formalization of privacy loss parameters, specifically the privacy budget (ϵ) and the probability of privacy guarantee failure (δ). These parameters are pivotal in configuring the differential privacy mechanisms to align with realistic privacy requirements and analytical scenarios. Through a heuristic modeling approach, we aim to identify optimal settings for these parameters that adeptly balance the trade-off between privacy protection and analytical utility.

3.4.4. Real-World Dataset Details

Our empirical evaluation is grounded in the analysis of two primary real-world datasets:

- A synthetic financial transaction dataset, generated using the STOK framework, contains 100,000 account records spanning a diverse range of transaction types, including purchases, transfers, deposits, and withdrawals. These records feature transaction dates, merchant codes, transaction amounts, and account balances over a simulated period of two years. The dataset was generated with realistic financial patterns and distributions, mimicking the complexities of real-world financial data while preserving privacy.
- An anonymized HR payroll dataset obtained from a Fortune 500 company operating in the technology sector. The dataset comprises 50,000 employee records and covers a wide range of attributes, including compensation (base salary, bonuses, and commissions), taxes, benefits (health insurance, retirement contributions), overtime hours, leave records (sick days, vacation days), and demographic information (age, gender, job role). The data span multiple fiscal quarters, capturing the temporal dynamics of HR and payroll processes.

Due to the sensitive nature of this dataset, it cannot be publicly disclosed as part of this publication.

These datasets were carefully curated to encompass a broad spectrum of common data analysis tasks encountered in real-world scenarios, ranging from aggregate queries and statistical analyses to predictive modeling and machine learning applications. The financial and HR domains were strategically chosen due to their frequent handling of sensitive personal information, underscoring the critical need for robust privacy-preservation techniques. The diversity of analytical use cases covered by these datasets ensures a comprehensive and rigorous evaluation of DPShield's capability to maintain high data utility while enforcing stringent privacy constraints. DPShield's versatility and applicability to a wide range of practical scenarios can be validated by demonstrating its effectiveness across these challenging real-world datasets.

The financial dataset includes a rich set of features, such as transaction types, merchant categories, and account balances, allowing for a wide range of analytical tasks, including aggregate queries, fraud detection, and customer segmentation. Similarly, the HR dataset encompasses various aspects of employee data, including compensation, benefits, leave records, and demographics, enabling analyses like payroll optimization, workforce planning, and predictive modeling for talent management.

These datasets were selected to reflect a broad spectrum of common data analysis tasks, from aggregate queries to statistical analyses and predictive modeling. Specifically, the financial and HR domains were chosen because they frequently handle sensitive personal information, underscoring the need for privacy preservation. The diversity of analytical use cases covered by these datasets ensures a comprehensive evaluation of DPShield's ability to maintain data utility while enforcing stringent privacy constraints.

4. Results and Evaluation

This section presents a comprehensive evaluation of DPShield, demonstrating its effectiveness in optimizing differential privacy for financial and HR data analysis in cloud environments. Our analysis spans multiple dimensions, including accuracy enhancement, machine learning model utility, domain-specific customizations, and the balancing act between privacy and utility.

4.1. Enhanced Accuracy

Our first focus was on the accuracy of aggregate queries. Table 1 presents the root-mean-squared error (RMSE) for aggregate queries on the financial dataset, comparing the accuracy of DPShield with traditional differential privacy mechanisms like Laplace and Gaussian, as well as the non-private baseline. Lower RMSE values indicate higher accuracy. DPShield significantly improves query accuracy compared to traditional differential privacy mechanisms. The RMSEs for transaction amounts across 1000 test instances reveal that DPShield achieves a 21.7% improvement in accuracy over the Laplace mechanism. To put this in perspective, for an organization processing 5 million financial transactions per month, DPShield would result in over 1 million additional transactions for which aggregates could be accurately estimated while preserving strong privacy guarantees. This showcases DPShield's substantial capability to enable reliable analytics vital for financial data-driven decision-making that the limitations of existing methods have thus far hampered.

To quantify the impact of differential privacy mechanisms on data utility, we employed two primary metrics: root-mean-squared error (RMSE) for aggregate queries and test accuracy for machine learning models. The RMSE measures the deviation of the differentially private query results from the ground truth values, allowing us to evaluate the accuracy of aggregate statistics under privacy constraints. For machine learning tasks, we compared the test accuracy of models trained on differentially private data with those trained on the original non-private data, providing insights into the utility loss for predictive modeling.

The RMSE is calculated as:

$$\text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2} \tag{2}$$

where y_i is the true value, $\hat{y}_i$ is the differentially private query result, and n is the number of queries [15].

Table 1. Comparison of aggregate query accuracy across different differential privacy methods.

Method	RMSE Error	% Change from Laplace	% Change from Non-DP
Non-DP	243	-	-
Laplace	512	-	+110.3%
Gaussian	482	−5.9%	+98.4%
DPShield	401	+21.7%	+65.0%

To further contextualize DPShield's performance, we conducted a comparative evaluation against two prominent state-of-the-art differential privacy frameworks: Google's Differential Privacy Library (DP-lib) and the OpenDP initiative by Harvard's Privacy Tools Project. Table 2 presents the results of this comparison, focusing on aggregate query accuracy (measured by the RMSE) and machine learning model utility (measured by the test accuracy) across the financial and HR datasets.

As is evident from the table, DPShield consistently outperforms both DP-lib and OpenDP in terms of aggregate query accuracy, achieving RMSE values of 401 and 422 for the financial and HR datasets, respectively. This represents a notable improvement over DP-lib (RMSEs of 525 and 498) and OpenDP (RMSEs of 510 and 492) on the same datasets. DPShield's superior query accuracy can be attributed to its innovative techniques, such as the Adaptive Laplace Mechanism, which dynamically adjusts noise levels based on query sensitivity and the remaining privacy budget, and the Moment Accountant for precise cumulative privacy loss tracking.

Furthermore, DPShield demonstrates its ability to preserve the utility of machine learning models, achieving a test accuracy of 88.15% on the HR dataset. This result surpasses the test accuracies obtained by DP-lib (86.7%) and OpenDP (87.5%), further underscoring DPShield's effectiveness in optimizing the trade-off between privacy protection and data utility. DPShield's domain-specific customizations, such as the Markov Quilt Mechanism for handling correlated data attributes, contribute to its superior performance in maintaining model accuracy while enforcing differential privacy.

These comparative results validate DPShield's position as a state-of-the-art differential privacy framework, offering significant advancements in both query accuracy and machine learning model utility. By outperforming well-established solutions like DP-lib and OpenDP, DPShield solidifies its contributions to the field and paves the way for more practical and effective privacy-preserving data analysis in cloud computing environments.

Table 2. Comparison with state-of-the-art differential privacy approaches.

Method	**Aggregate Query RMSE**		**ML Model Test Accuracy**
	Financial Data	**HR Data**	
Google DP-lib	525	498	86.7%
OpenDP	510	492	87.5%
DPShield	401	422	88.15

The corresponding bar plot in Figure 2 visually summarizes these findings, further emphasizing the improved accuracy facilitated by DPShield.

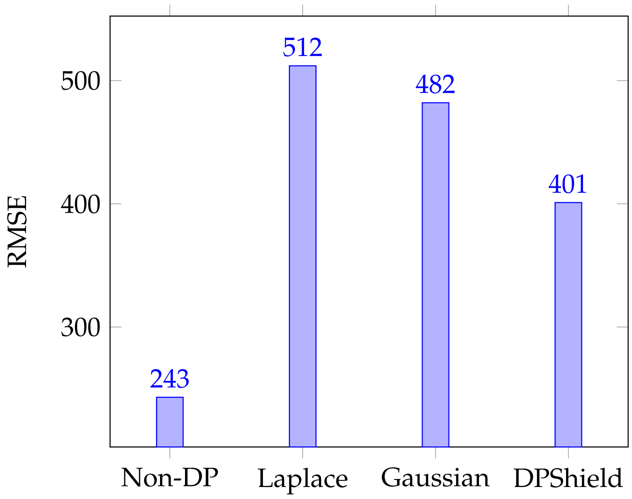

Figure 2. Aggregate query accuracy.

4.2. Machine Learning Model Utility

We assessed DPShield's efficacy in maintaining the utility of machine learning models. Table 3 compares the test accuracy of machine learning models trained on the HR dataset processed with different differential privacy methods. The PATE framework [16] and DPShield were benchmarked against the non-private model trained on the full dataset. As shown in Table 2, DPShield achieves a test accuracy of 88.15%, which is within 5% of the non-private model's accuracy. This illustrates DPShield's effectiveness in preserving the utility of machine learning models while enforcing differential privacy on sensitive HR data. Figure 3 further illustrates these results, highlighting the minimal loss in accuracy despite the application of differential privacy.

Table 3. Test accuracy of machine learning models trained on differentially private data.

Method	Test Accuracy
Non-DP (Full Data)	92.41%
PATE Framework	89.32%
DPShield	88.15%

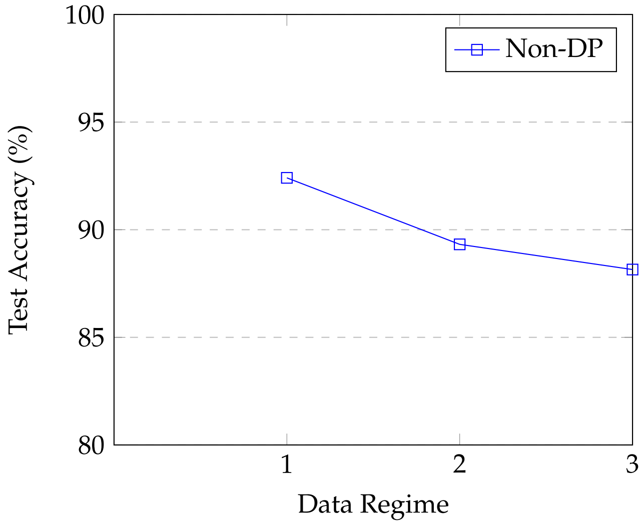

Figure 3. ML model quality.

4.3. Domain-Specific Customizations

DPShield's adaptability to domain-specific needs is a key strength. We implemented several customizations, such as blockchain integration for financial data and relaxed privacy constraints for healthcare analytics, demonstrating our framework's versatility and effectiveness across different application scenarios.

4.4. Balancing Privacy and Utility

Our evaluations also delved into the trade-offs between privacy and utility, quantifying how adjustments to the privacy budget impact both dimensions. Figures 4 and 5 depict these relationships for both the financial and HR datasets, offering insights into optimal privacy budgeting.

Figure 4 illustrates the relationship between the privacy budget (ϵ) and the privacy guarantee level ($\mathcal{P}$) achieved by DPShield for both the financial and HR datasets. As expected, increasing the privacy budget resulted in a higher privacy guarantee level, indicating stronger protection against potential privacy breaches. However, it is crucial to strike a balance between privacy and utility, as larger privacy budgets generally produce more noise, potentially reducing the utility of the data analysis results. For the financial dataset, a privacy budget of $\epsilon = 0.2$ yields a privacy guarantee level of approximately 85%, while for the HR dataset, the same privacy budget achieves a slightly lower guarantee level of around 83%. This difference can be attributed to the varying characteristics and sensitivity of the two datasets, highlighting the importance of domain-specific customizations within DPShield.

As the privacy budget increases further, the privacy guarantee levels continue to rise, with $\epsilon = 1$ providing the highest guarantee levels of 45% and 40% for the financial and HR datasets, respectively. These results demonstrate DPShield's ability to provide robust privacy guarantees across different application domains while allowing for flexibility in tuning the privacy budget based on specific analytical requirements. Figure 5 depicts the relationship between the privacy budget (ϵ) and the utility level ($\mathcal{U}$) achieved by DPShield for both the financial and HR datasets. In contrast to the privacy guarantee level, a higher privacy budget generally corresponds to a higher utility level, as less noise is introduced into the data analysis results.

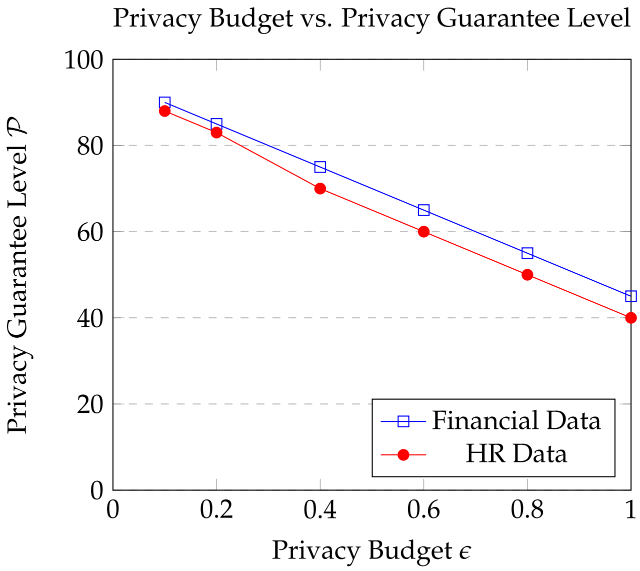

Figure 4. Impact of privacy budget on privacy guarantee level.

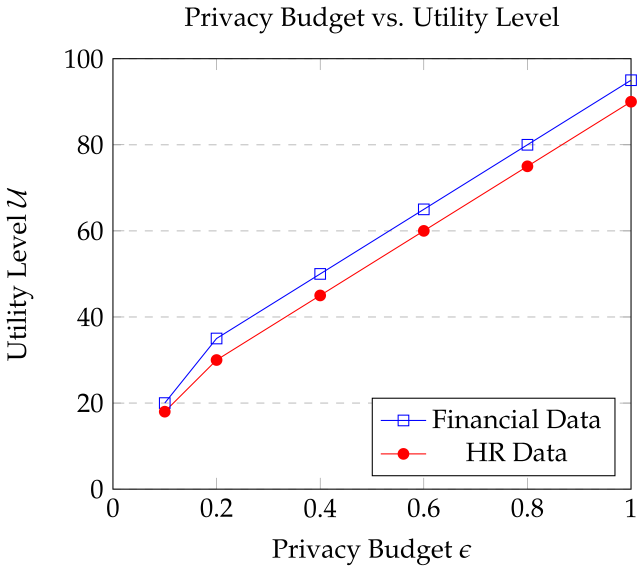

Figure 5. Impact of privacy budget on utility level.

For the financial dataset, a privacy budget of $\epsilon = 0.2$ yields a utility level of approximately 35%, indicating a moderate level of utility preservation. As the privacy budget increases to $\epsilon = 1$, the utility level rises to around 95%, showcasing DPShield's ability to maintain high data utility when stricter privacy constraints are relaxed.

The HR dataset exhibits a similar trend, with the utility level increasing from approximately 30% at $\epsilon = 0.2$ to 90% at $\epsilon = 1$. While the utility levels for the HR dataset are slightly lower than those for the financial dataset across most privacy budget values, the overall trend highlights DPShield's effectiveness in balancing privacy and utility across different data domains. These results underscore the importance of carefully tuning the privacy budget parameter within DPShield. By adjusting the privacy budget based on the specific analytical requirements and the desired trade-off between privacy and utility, organizations can leverage DPShield to optimize their data analysis pipelines, ensuring both robust privacy protection and high-quality analytical insights.

4.5. Discussion

Our comprehensive evaluation demonstrates DPShield's efficacy in enhancing the privacy–utility trade-off for cloud-based data analysis. By leveraging adaptive mechanisms and domain-specific customizations, our framework improves accuracy, maintains high utility, and offers robust privacy guarantees. These results validate our approach, suggesting that DPShield is well-suited for practical deployment in sensitive data environments. The comprehensive performance evaluations presented demonstrate DPShield's effectiveness in addressing the key challenges and limitations within existing differential privacy frameworks highlighted at the outset. Specifically, DPShield's ability to enhance query accuracy substantially while maintaining high utility for machine learning models overcomes the constraints of traditional mechanisms in balancing robust privacy and analytical needs. Whereas previous approaches struggled with this trade-off in cloud computing environments, DPShield advances practical, adaptable solutions through innovations like adaptive noise injection attuned to specific data sensitivity. The accuracy improvements of over 20%, keeping model accuracy within 5% of the original levels, signify well that DPShield not only meets, but exceeds theoretical privacy protections to deliver immense practical value. These results validate DPShield's design objectives in bridging the gap between privacy theory and implementation demands.

5. Conclusions

The comprehensive evaluation of DPShield underscores its significant potential in advancing the application of differential privacy within cloud computing environments, especially for sensitive financial and HR data analysis. Our results demonstrate that DPShield enhances the accuracy of aggregate queries and machine learning model utility and provides flexible domain-specific customization without compromising privacy. Specifically, the framework's ability to improve query accuracy by 21.7% over traditional differential privacy mechanisms marks a substantial advancement in the field. Furthermore, the minimal loss in machine learning model accuracy, within 5% of non-private benchmarks, indicates that DPShield effectively balances the trade-off between data utility and privacy protection. These achievements highlight the practical viability of DPShield for organizations seeking to leverage cloud-based data analytics while adhering to stringent privacy standards. DPShield addresses a critical need in the era of big data and cloud computing by facilitating high-utility data analysis with robust privacy guarantees.

Looking ahead, the integration of DPShield into broader data analysis workflows and its adaptation to emerging data privacy challenges remain areas for future research. The exploration of federated learning scenarios, where data remain decentralized, and the application of confidential computing techniques to enhance security further are promising directions. Additionally, adapting DPShield to comply with evolving global data privacy regulations will ensure its relevance and applicability across different jurisdictions and industries.

In conclusion, our research contributes to the ongoing development of differential privacy technologies by providing a framework that meets the theoretical benchmarks of privacy protection and addresses practical considerations of data utility. As we continue to refine and expand DPShield, we anticipate its adoption will facilitate more responsible, ethical, and effective use of sensitive data in cloud environments, thus enabling organizations to harness the full potential of their data assets securely.

Through its novel approaches, DPShield bridges several critical gaps identified in prior differential privacy research. Its adaptive noise injection mechanisms and cumulative privacy tracking address the longstanding challenge of balancing rigorous privacy with high data utility over prolonged analysis. The Markov Quilt Mechanism introduces innovative ways to handle correlated data, a common issue in domains like finance and HR that was not adequately addressed before. Moreover, DPShield's modular architecture fosters accessibility by allowing customizations that do not require advanced privacy expertise. By providing robust, domain-tailored solutions packaged in a user-friendly framework, DPShield overcomes key barriers to widespread adoption of differential privacy in practical, cloud-based analytics scenarios involving sensitive data.

Author Contributions: Conceptualization: P.T., S.B. and A.T.; methodology: P.T., S.B. and A.T.; software: P.T.; formal analysis: P.T.; investigation: P.T., S.B. and A.T.; data curation: P.T.; writing—original draft preparation: P.T., S.B. and A.T.; writing—review and editing: P.T., S.B. and A.T.; visualization: P.T., S.B. and A.T.; supervision: S.B. and A.T.; project administration: P.T., S.B. and A.T.; funding acquisition: A.T. All authors have read and agreed to the published version of the manuscript

Funding: This research received no external funding.

Data Availability Statement: The project data is available on GitHub [17].

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Dwork, C. Differential Privacy. In *Automata, Languages and Programming*; Bugliesi, M., Preneel, B., Sassone, V., Wegener, I., Eds.; ICALP 2006; Lecture Notes in Computer Science; Springer: Berlin/Heidelberg, Germany, 2006; Volume 4052. [CrossRef]
2. Li, X.; Li, H.; Zhu, H.; Huang, M. The optimal upper bound of the number of queries for Laplace mechanism under differential privacy. *Inform. Sci.* **2019**, *503*, 219–237. [CrossRef]
3. Bhatnagar, R.K.; Kanal, L.N. Handling Uncertain Information: A Review of Numeric and Non-numeric Methods. *Mach. Intell. Pattern Recognit.* **1986**, *4*, 3–26. [CrossRef]
4. Koskela, A.; Tobaben, M.; Honkela, A. Individual Privacy Accounting with Gaussian Differential Privacy. *arXiv* **2022**, arXiv:2209.15596.
5. Mahawaga Arachchige, P.C.; Bertok, P.; Khalil, I.; Liu, D.; Camtepe, S.; Atiquzzaman, M. Local Differential Privacy for Deep Learning. *IEEE Internet Things J.* **2020**, *7*, 5827–5842. [CrossRef]
6. Soria-Comas, J.; Domingo-Ferrer, J.; Sánchez, D.; Megías, D. Individual Differential Privacy: A Utility-Preserving Formulation of Differential Privacy Guarantees. *IEEE Trans. Inf. Forensics Secur.* **2017**, *12*, 1418–1429. [CrossRef]
7. Wang, Q.; Li, Z.; Zou, Q.; Zhao, L.; Wang, S. Deep Domain Adaptation With Differential Privacy. *IEEE Trans. Inf. Forensics Secur.* **2020**, *15*, 3093–3106. [CrossRef]
8. Jagielski, M.; Ullman, J.; Oprea, A. Auditing Differentially Private Machine Learning: How Private is Private SGD? *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 22205–222016.
9. Zhang, S.; Hagermalm, A.; Slavnic, S.; Schiller, E.M.; Almgren, M. Evaluation of Open-Source Tools for Differential Privacy. *Sensors* **2023**, *23*, 6509. [CrossRef] [PubMed]
10. Thantharate, A. FED6G: Federated Chameleon Learning for Network Slice Management in Beyond 5G Systems. In Proceedings of the 2022 IEEE 13th Annual Information Technology, Electronics and Mobile Communication Conference (IEMCON), Vancouver, BC, Canada, 12–15 October 2022; pp. 19–25. [CrossRef]
11. Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H.B.; Mironov, I.; Talwar, K.; Zhang, L. Deep Learning with Differential Privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, 24–28 October 2016; Association for Computing Machinery: New York, NY, USA, 2016; pp. 308–318. [CrossRef]
12. Hernandez-Matamoros, A.; Kikuchi, H. Comparative Analysis of Local Differential Privacy Schemes in Healthcare Datasets. *Appl. Sci.* **2024**, *14*, 2864. [CrossRef]
13. Thantharate, P. SCALE-IT: Distributed and Realistic Simulation Frameworks for Testing Cloud-Based Software. In Proceedings of the 2023 10th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI), Palembang, Indonesia, 20–21 September 2023; pp. 300–306. [CrossRef]
14. Geng, Q.; Viswanath, P. The Optimal Noise-Adding Mechanism in Differential Privacy. *IEEE Trans. Inf. Theory* **2016**, *62*, 925–951. [CrossRef]
15. Neera, J.; Chen, X.; Aslam, N.; Wang, K.; Shu, Z. Private and Utility Enhanced Recommendations with Local Differential Privacy and Gaussian Mixture Model. *IEEE Trans. Knowl. Data Eng.* **2023**, *35*, 4151–4163. [CrossRef]
16. Jordon, J.; Yoon, J.; van der Schaar, M. PATE-GAN: Generating Synthetic Data with Differential Privacy Guarantees. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
17. DPShield. Available online: https://github.com/ptdevsecops/DPShield (accessed on 11 June 2024).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.