Local Differential Privacy: Refined Mechanism Design and Utility Analysis

by

Ye Zheng

A proposal submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in Computing and Information Sciences

B. Thomas Golisano College of Computing and Information Sciences
Rochester Institute of Technology

[Month and year of Dissertation Acceptance was signed]

GCCIS Ph.D. PROGRAM IN COMPUTING AND INFORMAITON SCIENCES
ROCHESTER INSTITUTE OF TECHNOLOGY
ROCHESTER, NEW YORK

CERTIFICATE OF APPROVAL

Ph.D. DEGREE PROPOSAL

The Ph.D. Degree Proposal of Ye Zheng
has been examined and approved by the
proposal committee as satisfactory for the
proposal required for the
Ph.D. degree in Computing and Information Sciences

Ph.D. Program Director Date

Yidan Hu, Proposal Advisor Date

[External Chair's name], External Chair Date

[Committee member's name] Date

[Committee member's name] Date

Local Differential Privacy: Refined Mechanism Design and Utility Analysis

by

Ye Zheng

Submitted to the
B. Thomas Golisano College of Computing and Information Sciences Ph.D. Program in Computing and Information Sciences
in partial fulfillment of the requirements for the
Doctor of Philosophy Degree
at the Rochester Institute of Technology

Abstract

Local Differential Privacy (LDP) is a privacy model that enables users to perturb their data locally before sharing it with untrusted data collectors for analysis. This privacy model provides provable privacy guarantees for each individual user. Owing to these guarantees, it has been widely deployed by major technology companies, inducing Apple, Google, and Microsoft, to protect user privacy while still enabling the collection of data for analytics and machine learning tasks. However, a fundamental challenge of LDP is the trade-off between privacy and data utility: stronger privacy guarantees for users typically result in lower data utility for data collectors. Addressing this challenge, a central direction of theoretical LDP research is to design mechanisms that provide better data utility under the same privacy guarantee. This dissertation focuses on the central direction, aiming to advance both the design of LDP mechanisms and the analysis of data utility under LDP.

This dissertation makes four main contributions: (i) It introduces correlated perturbation of multiple attributes to improve data utility under LDP, generalizing existing independent-perturbation mechanisms; (ii) it establishes the optimality of piecewise-based mechanisms, a state-of-the-art category of LDP mechanisms for collecting bounded numerical data; (iii) building on piecewise-based mechanisms, it proposes two mechanisms for collecting individual trajectory data, which achieve higher efficiency and data utility by operating in continuous space instead of previously studied discrete space; (iv) it provides a quantification framework for theoretically analyzing data utility of classifiers under LDP-perturbed inputs, making a first step towards connecting LDP mechanisms with robustness.

Acknowledgments

This is the acknowledgements text

To the three cold, sober years I lived in Rochester,
whose quiet shaped me and whose memory I shall always carry.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Local differential privacy (LDP) mechanisms protect individual users' data privacy against untrusted data collectors by allowing each user to locally perturb their data before sharing it [6, 8]. Though the data collector receives only perturbed data, they can still learn valuable statistics while being unable to infer much about any individual user's true data, with privacy guarantees quantified by the privacy parameter ε. Due to these provable privacy guarantees, LDP mechanisms have been widely adopted by major technology companies, including Apple's operating systems [26], Google Chrome [9], and Microsoft Office [28] for collecting user statistics on-device. Furthermore, LDP is a key privacy-enhancing component in federated learning, a decentralized machine learning paradigm where users collaboratively train a global model without sharing sensitive data with a central server [1, 14].

Figure 1.1 illustrates the LDP system model, in which an untrusted data collector may attempt to infer users' true (sensitive) data. The LDP mechanism $\mathcal{M}$ acts as a guard at the trust boundary,

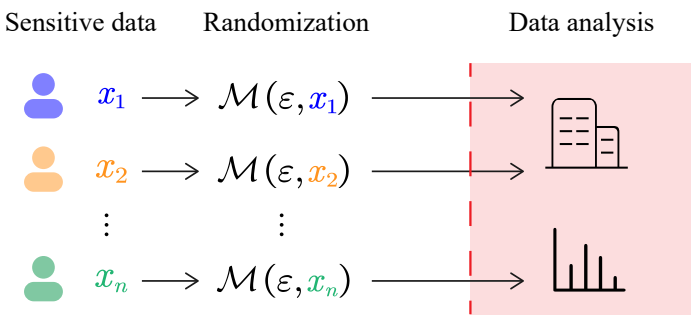

Figure 1.1: LDP system model. The dashed red line indicates the trust boundary. Each user locally perturbs their true data x_i using an LDP mechanism $\mathcal{M}$ before sending it to the untrusted data collector. The data collector has access only to the perturbed data $\tilde{x}_i = \mathcal{M}(\varepsilon, x_i)$, which it uses to perform statistical analysis.

Table 1.1: Comparison of typical privacy-enhancing techniques.

Threat Model	Technique	Privacy Guarantee	Complexity	Data Utility
Trusted Collector	k-anonymity [29]	Syntactic	Medium	High
	Central DP [7]	Semantic (ε-DP)	Simple	ε-dependent
Untrusted Collector[a]	HE[b] [27]	Semantic (IND-CPA)	Complex	High
	MPC[c] [36]	Semantic (Real/Ideal)	Complex	High
	LDP [6]	Semantic (ε-LDP)	Simple	ε-dependent

[a] For correct service functionality, an untrusted data collector is often modeled as honest-but-curious.

[b] Homomorphic Encryption (HE) typically requires that the service functionality be expressible using operations supported by the encryption scheme.

[c] Secure Multi-Party Computation (MPC) usually assumes that only a subset of parties may be malicious.

preserving privacy by injecting unreversible randomness (quantified by ε) into each user's true data x_i before it leaves their device.

Advantages over other privacy-enhancing techniques. Orthogonal to LDP, various privacy-enhancing techniques (PETs) have been proposed to protect user data privacy in data collection and analysis. They are designed under different threat models and offer different senses of privacy guarantee. Table 1.1 summarizes and compares typical PETs.

k-anonymity [29] and central DP [7] assume a trusted data collector with direct access to users' true data, and place the responsibility for safeguarding individual privacy against external inference attacks on the collector. Among them, (i) k-anonymity is a *syntactic* privacy model that protects the linkage between users' identities and their data by masking quasi-identifiers in the dataset, rather than directly protecting the sensitive data themselves. Designing an effective masking scheme is often difficult for high-dimensional data, and such syntactic models have been shown to be vulnerable to various attacks [13, 16, 22]. (ii) Central DP, in contrast, provides a *semantic* privacy guarantee by adding randomness to aggregated statistics before releasing them, making it difficult to infer any individual user's data from the published statistics. A subtle but important issue in central DP is the definition of "neighboring datasets" (e.g. record removal vs. record replacement), which can lead to different required noise levels under the same privacy parameter ε [23, 24, 25].

Homomorphic Encryption (HE) [27] and Secure Multi-Party Computation (MPC) [36] are cryptographic techniques that can also protect user data privacy when the data collector is untrusted. (i) HE enables certain function evaluations to be performed directly on encrypted data without decryption. Its privacy guarantee relies on indistinguishability under chosen-plaintext attack (IND-CPA), which ensures that an adversary cannot distinguish the encryptions of any two chosen

plaintexts. (ii) MPC allows multiple parties to jointly compute a function over their private inputs without revealing those inputs to one another. Its privacy guarantee is formalized via the real/ideal simulation paradigm, which requires that an adversary's (e.g. a malicious party's) view during the protocol execution can be simulated using only its own input and output, implying that no additional information is leaked to the adversary. Both HE and MPC typically require sophisticated protocol design and incur substantial computational and communication overhead.

Compared with these PETs, LDP offers several advantages. *It provides semantic privacy guarantees for individual users' data, quantified by the clean ε-LDP notion, agnostic to the data collector and has low complexity suitable for resource-constrained devices.*

Despite these advantages, LDP mechanisms are fundamentally constrained by a privacy-utility trade-off: achieving stronger privacy guarantees (i.e. smaller ε) requires injecting more randomness into users' data, which in turn degrades the utility of the collected data for downstream statistical analysis. Figure 1.2 illustrates this trade-off by showing the worst-case expected error of $\mathcal{M}(\varepsilon, x_i)$ for four LDP mechanisms on a numerical data domain $x_i \in [0, 1]$. As ε increases, the worst-case expected error decreases for all mechanisms, indicating improved data utility. However, for a fixed privacy parameter ε, the utility achieved by different mechanisms can vary significantly. Some mechanisms, such as OGPM, achieve lower expected error across most ε values than others, demonstrating a better privacy-utility trade-off. A central direction of theoretical LDP research is therefore to design mechanisms that optimize this trade-off, as they directly determine the building blocks of practical LDP systems deployed in real-world applications. Beyond this, it is also important to quantify the utility of complex data analysis tasks under LDP, such as classifier performance, which cannot be inferred solely from the expected error of the underlying LDP mechanisms. Understanding these utility implications is crucial for guiding the selection and configuration of LDP mechanisms in practice.

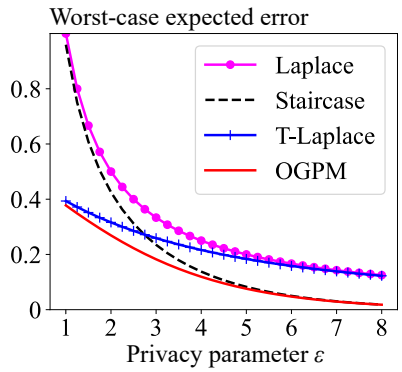

Figure 1.2: Privacy-utility curves.

To summarize, there are two fundamental challenges in advancing current LDP research:

- Designing LDP mechanisms with optimized privacy-utility trade-offs.
- Quantifying the utility of complex data analysis tasks under LDP.

Contributions of this dissertation. Aiming to address these fundamental challenges, this dissertation makes the following contributions:

- Chapter 3 introduces correlated perturbation into LDP mechanisms for multiple users' data, which generalizes existing LDP mechanisms that perturb each user's data independently and achieves improved privacy-utility trade-offs.
- Chapter 4 establishes the optimality of piecewise-based mechanisms, state-of-the-art category of LDP mechanisms for collecting bounded numerical data.
- Chapter 5 proposes two mechanisms for collecting individual trajectory data, which achieve higher efficiency and data utility by operating in continuous space instead of previously studied discrete space.
- Chapter 6 provides a quantification framework for theoretically analyzing data utility of classifiers under LDP-perturbed inputs, making a first step towards connecting LDP mechanisms with robustness.

The remainder of this chapter briefly overviews the research questions and key ideas underlying each contribution. Implementations and evaluations are available at `https://github.com/ZhengYeah/`.

1.1 Correlated Perturbation for LDP

The most classical and widely used application of LDP is frequency estimation, where the data collector aims to estimate the number (or proportion) of users possessing a certain attribute or data value. Randomized Response (RR) [35] is the first known and most classical LDP protocol for frequency estimation on binary data (e.g. yes/no questions). In RR, each user perturbs their true binary data independently by flipping it with a probability determined by the privacy parameter ε. Due to its simplicity and effectiveness, RR has been widely adopted as a building block in many LDP mechanisms for diverse data types and analysis tasks [2, 4, 33, 34]. A common feature of these RR-based mechanisms, and of LDP mechanisms more broadly, is that each user's data is perturbed independently, resulting in a large amount of total randomness. This raises a natural research question: *Can the data utility of RR be improved by introducing correlations among the random perturbations performed by different users?*

Chapter 3 investigates correlated random perturbations for frequency estimation to improve data utility without weakening LDP guarantees. The key insight is that the total randomness injected into users' data can be reduced by partitioning data contributors into disjoint pairs and introducing carefully designed correlations into each pair's random perturbations, as illustrated

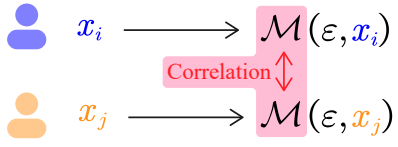

Figure 1.3: Correlated perturbation.

in Figure 1.3. Crucially, no additional information is revealed to
the data collector as long as the group membership remains hidden. A novel Joint Randomized Response (JRR) mechanism is proposed based on this idea. With appropriately chosen parameters, JRR achieves substantially higher data utility in the vast majority of cases, while providing the same level of LDP protection as classical RR.

Specifically, this chapter makes the following contributions:

- (Correlated perturbation) It makes the first attempt in the LDP literature to introduce correlated perturbations into LDP mechanisms, thereby improving the data utility of frequency estimation.
- (The JRR mechanism) It proposes a general JRR mechanism that provides the same level of LDP protection as classical RR, while substantially improving data utility in most cases, particularly when the number of data contributors is large.
- (Practical instantiations) It presents a practical instantiation of JRR that leverages either a non-colluding auxiliary server or an MPC protocol to conceal group membership from the data collector.

1.2 Optimal Piecewise-based Mechanism

Numerical data with bounded domains is a fundamental data type in personal devices and sensor networks. These bounded domains can be categorized into two types: linear ranges, such as sensor readings in $[0, 1)^*$, referred to as the *classical domain*; and cyclic ranges, such as angular measurements in $[0, 2\pi)$, referred to as the *circular domain*. The Laplace mechanism [7] is the most classical LDP mechanism for numerical data privacy: it adds random noise drawn from a Laplace distribution determined by ε to the sensitive data. However, the unbounded support of the Laplace noise makes it unsuitable for bounded domains.

State-of-the-art LDP mechanisms for numerical data on bounded domains are *piecewise-based mechanisms* [17, 20, 31]. These mechanisms randomize sensitive data to values drawn from carefully designed piecewise probability distributions. Existing instantiations use different pieces and probabilities, but are all designed for classical (linear) domains. Their applicability to other bounded domains, such as the circular domains of angular sensors that frequently arise in personal devices, remains unexplored.

*To ease interval operations (e.g. union and intersection), this dissertation uses left-closed right-open intervals. They are equivalent to closed intervals in implementation and practical applications.

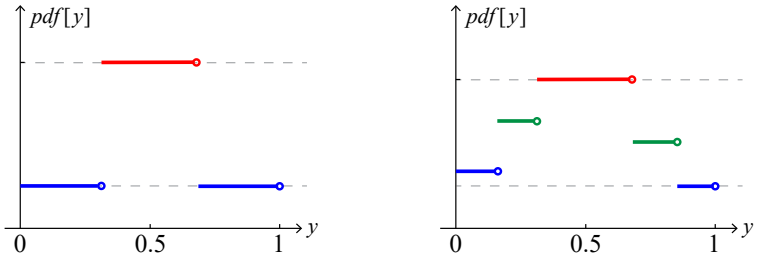

Figure 1.4: Illustration of piecewise distributions. Existing piecewise-based mechanisms [17, 20, 31] are special cases of 3-piece distributions (left). In contrast, a general piecewise distribution can have an arbitrary number of pieces and locations (right), which may improve data utility as existing evidence [10, 30] suggests.

The optimality of piecewise-based mechanisms remains a challenging open problem. Existing instantiations can be viewed as heuristic forms of the 3-piece mechanism (TPM). As a special case with exactly 3 pieces and pre-defined functional forms, TPM is too restrictive to fully characterize the optimality of piecewise-based mechanisms. Evidence from the staircase Laplace mechanism [10] for unbounded numerical data shows that the asymptotically optimal mechanism has a staircase (multi-piece) structure. For categorical data, the Staircase Randomized Response mechanism (SRR) [30] improves data utility in location collection compared to classical RR. These and other results suggest that increasing the diversity of probabilities over the data domain, i.e., using more pieces, can improve data utility. Motivated by this, a fundamental question for piecewise-based mechanisms is: *What is the optimal instantiation of a piecewise-based mechanism?* In designing such mechanisms, the number of pieces, as well as their probabilities and sizes, can be arbitrary. Finding the optimal instantiation within this large design space is challenging, as it requires jointly optimizing the number of pieces and the associated probabilities and sizes.

Chapter 4 studies the optimality of piecewise-based mechanisms in their most general form. It extends TPM to a generalized piecewise-based mechanism (GPM) with m pieces, where each piece has no predefined functional form. Within this GPM framework, it formulates an optimization problem that minimizes the distance between the sensitive and randomized data. By combining numerical solutions of this optimization problem with analytical proofs, it derives a closed-form characterization of the optimal GPM for classical domains. For circular domains, where the distance metric is periodic (e.g. the distance between 0 and 2π is zero), it explicitly incorporates this property into the mechanism design and reduce the search for the optimal mechanism to related problems on classical domains.

Specifically, this chapter makes the following contributions:

- (Solving framework) It is the first work to study the closed-form optimal piecewise-based mechanism in its most general form. A framework is proposed that combines analytical proofs

with off-the-shelf optimization solvers to derive the closed-form optimal mechanism, providing a practical foundation for achieving optimal data utility under LDP for bounded numerical data.

- (Closed-form instantiations) It derives closed-form optimal mechanisms for both the classical and circular domains. These mechanisms can be directly used as building blocks in applications such as sensor networks and federated learning.

1.3 Trajectory Collection in Continuous Space under LDP

Trajectory data from users—sequences of locations that describe movement over time—are a fundamental resource for activity analysis and location-based services, such as activity classification and routine detection. Existing LDP methods for trajectory collection are primarily designed for discrete spaces. They rely on discrete-domain mechanisms, such as the Exponential mechanism [21], to perturb trajectory data. Because discrete LDP mechanisms are explicitly defined with respect to the size of the location space, these methods either partition the continuous space into grids [30] or assume that the location space is a finite set of labeled locations (points of interest) [5, 38].

Limitations of discrete-domain mechanisms. Relying on discrete spaces has three limitations: (i) Their privacy guarantees are inherently tied to the chosen discrete set. For example, a discrete location space with 10 points provides weaker protection than one with 100 points, since any inference attack has at least a 10% chance of success in the former, regardless of the privacy parameter ε. (ii) Their efficacy and efficiency are often constrained by the size of the discrete set. As the set size increases, the probability of selecting the true location decreases, while the widely used Exponential mechanism incurs linear sampling complexity in the size of the location set for each perturbed sample. (iii) Discrete methods are not directly applicable to inherently continuous location spaces, such as flying and sailing trajectories or sensor trajectories from wearable devices.

Chapter 5 addresses these limitations by *shifting the focus from discrete to continuous spaces for trajectory collection under LDP*. A continuous space models locations as real-valued coordinates, such as GPS positions in $[-180, 180] \times [-90, 90]$, and therefore contains infinitely many candidate locations. Collecting trajectory data directly in continuous spaces is natural for many applications and offers three key advantages over first discretizing the space and then applying discrete methods: (i) The privacy guarantee holds over the entire continuous space, independent of the chosen discretization. (ii) The sampling mechanism operates directly on the continuous domain, so its efficacy and efficiency are decoupled from the "number" of locations. (iii) Perturbed locations in the continuous space can be post-processed (e.g. by rounding) to any discrete space contained within it, making the approach

applicable to both continuous and discrete settings.

This chapter proposes two new LDP methods for trajectory collection in continuous spaces. The key idea is to decompose the 2D continuous space into 1D subspaces and design mechanisms for each subspace, building on piecewise-based mechanisms [17, 31, 39] for 1D bounded numerical domains.[†] Depending on the chosen decomposition of the continuous space, we obtain two methods, TraCS-D and TraCS-C.

Specifically, this chapter makes the following contributions:

- This is the first work to provide trajectory collection methods for continuous spaces under pure LDP. It introduces TraCS-D and TraCS-C, which employ novel location perturbation mechanisms that exploit both directional and coordinate information of trajectories in continuous spaces.
- TraCS is also applicable to discrete spaces. Compared with existing methods for discrete spaces, TraCS achieves significantly lower computational complexity when generating perturbed locations.

1.4 Quantification of Classifier Utility under LDP

Classifiers map input data to class labels and underpin a wide range of industrial applications, including predictive modeling, data analysis, and image recognition [12, 15, 19]. When deployed as services, classifiers require users to submit input data that often contain sensitive information, such as medical records or financial attributes, thereby raising serious privacy concerns. While users seek to benefit from classification services, they may be unwilling to disclose their sensitive data. A common mitigation strategy is client-side data perturbation, such as adding noise to numerical data or applying blurring and other obfuscation techniques to images before sending them to the classifier. Data perturbation remains a lightweight and intuitive solution for privacy-preserving classification [3, 37, 40]. Among these approaches, LDP mechanisms provide users with provable privacy guarantees. However, such perturbations inevitably degrade the utility of the classifier, leading to a fundamental research question: *How can we quantify the utility of classifiers when their inputs are perturbed by LDP mechanisms?*

Utility of classifiers under LDP. For simple queries such as summation, utility can be quantified

[†]Another category of LDP mechanisms for bounded numerical domains is truncated mechanisms, e.g. the truncated Laplace mechanism [11, 18]. While they can be incorporated into TraCS, they are more complex and less effective than piecewise-based mechanisms.

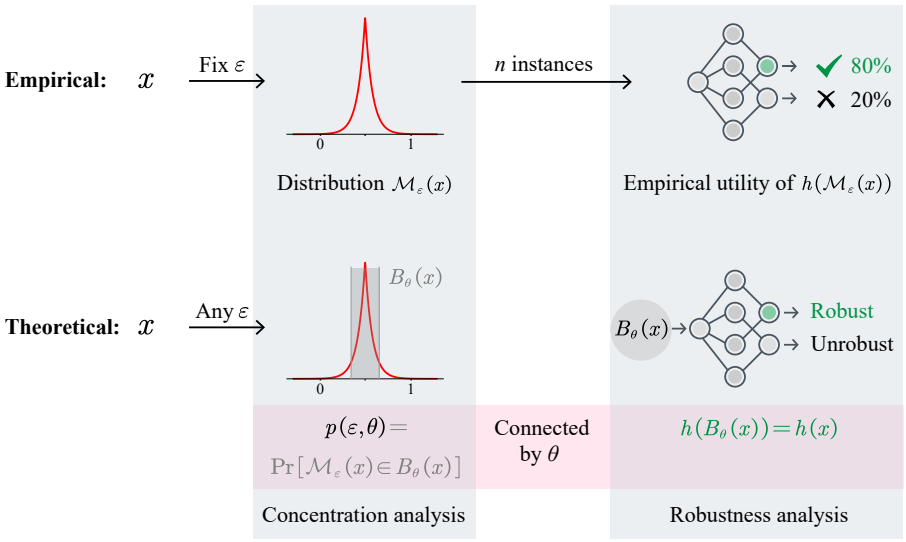

Figure 1.5: From empirical to theoretical utility analysis. Chapter 6 analytically quantifies classifier utility under LDP mechanisms by linking the *concentration analysis* of LDP mechanisms with the *robustness analysis* of classifiers.

analytically via the mean squared error (MSE) of LDP mechanisms [31, 32]. For classification tasks, however, utility is typically measured by classification accuracy, which cannot be written analytically in terms of MSE. A straightforward way to evaluate classifier utility under an LDP mechanism with a given privacy parameter ε is to repeatedly perturb the data and measure the proportion of correctly classified instances, as illustrated in Figure 1.5. Although practical, this empirical approach has significant limitations: it applies only to specific choices of ε and particular perturbed datasets; changing ε requires time-consuming re-evaluation, and different random perturbations lead to different results. Moreover, it does not reveal how utility depends on the privacy parameter, making it unsuitable for systematic comparison of LDP mechanisms. In contrast, an analytical utility quantification framework—analogous to MSE for summation queries—would provide principled guidance for the design and deployment of classifiers under LDP, but such a framework is currently lacking.

Quantifying the relationship between privacy and classifier utility presents significant analytical challenges. These challenges stem from two aspects: (i) LDP mechanisms inherently introduce perturbations across the entire data domain, potentially causing no utility for classifiers. (ii) Classifiers are often complex or even black-box functions, making it difficult to analyze their behavior under perturbations.

Chapter 6 develops a framework to theoretically quantify classifier utility under LDP mechanisms. This framework addresses the above challenges via two key insights: (i) LDP mechanisms generate perturbed data that, with high probability, concentrates within a bounded region around the original

data, making extreme perturbations rare. (ii) Within this concentration region, utility analysis of a classifier can be reformulated as a robustness analysis problem. Here, robustness characterizes how reliably a classifier preserves its predictions under input perturbations. Established robustness analysis techniques can determine the maximum permissible perturbation that leaves the classifier's output unchanged, thereby preserving utility. By combining the concentration analysis of LDP mechanisms with the robustness analysis of classifiers, a theoretical characterization of data utility can be obtained as a function of the privacy parameter ε, formalized as: *Given classifier h, LDP mechanism $\mathcal{M}_\varepsilon$, and raw data x, with probability at least $p(\varepsilon, \theta)$, h preserves its correct classification result under $\mathcal{M}_\varepsilon(x)$.*

Applications. The proposed utility quantification framework has direct applications in privacy-preserving classification systems: (i) It enables a comparative analysis of different LDP mechanisms for a given classifier by evaluating their probability guarantees $p(\varepsilon, \theta)$. An LDP mechanism that provides a higher $p(\varepsilon, \theta)$ ensures better utility at the same privacy level. (ii) The framework facilitates the selection of an appropriate privacy parameter ε to meet specific utility requirements. For example, given a utility threshold p^*, the framework identifies the ε that satisfies $p(\varepsilon, \theta) \geq p^*$, achieving a precise privacy-utility balance when using the classifier.

Specifically, this chapter makes the following contributions:

- (Quantification framework) It introduces the first analytical framework for quantifying classifier utility under LDP mechanisms by linking the concentration analysis of LDP mechanisms with the robustness analysis of classifiers, enabling principled utility evaluation.
- (Refinement techniques) It develops two refinement techniques that enhance utility quantification. The first extends robustness from a scalar "robustness radius" to an axis-aligned "robustness hyperrectangle", enabling tighter robustness analysis. The second adapts the PAC privacy notion by introducing a new privacy indicator and an extended Gaussian mechanism that is applicable to any ε.

1.5 Dissertation Organization

The remainder of this dissertation presents the necessary background on LDP in Chapter 2, followed by four chapters that each address one of the research questions in detail. Each chapter begins with a technical overview of the problem and key ideas, and then presents the proposed methods, discussions, and experimental evaluations. Related work for each research question is reviewed at the end of the corresponding chapter.

Bibliography

[1] Martín Abadi, Andy Chu, Ian J. Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, Vienna, Austria, October 24-28, 2016*, pages 308–318. ACM, 2016.

[2] Raef Bassily and Adam D. Smith. Local, private, efficient protocols for succinct histograms. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 127–135. ACM, 2015.

[3] Keke Chen and Ling Liu. Privacy preserving data classification with rotation perturbation. In *Proceedings of the 5th IEEE International Conference on Data Mining ICDM 2005, 27-30 November 2005, Houston, Texas, USA*, pages 589–592. IEEE Computer Society, 2005.

[4] Graham Cormode, Tejas Kulkarni, and Divesh Srivastava. Answering range queries under local differential privacy. *Proc. VLDB Endow.*, 12(10):1126–1138, 2019.

[5] Teddy Cunningham, Graham Cormode, Hakan Ferhatosmanoglu, and Divesh Srivastava. Real-world trajectory sharing with local differential privacy. *Proc. VLDB Endow.*, 14(11):2283–2295, 2021.

[6] John C. Duchi, Michael I. Jordan, and Martin J. Wainwright. Local privacy and statistical minimax rates. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, Berkeley, CA, USA, October, 26-29, 2013*, pages 429–438. IEEE Computer Society, 2013.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam D. Smith. Calibrating noise to sensitivity in private data analysis. In Shai Halevi and Tal Rabin, editors, *Theory of Cryptography, Third Theory of Cryptography Conference, TCC 2006, New York, NY, USA, March 4-7, 2006, Proceedings*, volume 3876 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2006.

[8] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.

[9] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: randomized aggregatable privacy-preserving ordinal response. In Gail-Joon Ahn, Moti Yung, and Ninghui Li, editors, *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*, pages 1054–1067. ACM, 2014.

[10] Quan Geng and Pramod Viswanath. The optimal mechanism in differential privacy. In *2014 IEEE International Symposium on Information Theory, Honolulu, HI, USA, June 29 - July 4, 2014*, pages 2371–2375. IEEE, 2014.

[11] Naoise Holohan, Spiros Antonatos, Stefano Braghin, and Pól Mac Aonghusa. The bounded laplace mechanism in differential privacy. *J. Priv. Confidentiality*, 10(1), 2020.

[12] Anil K. Jain, M. Narasimha Murty, and Patrick J. Flynn. Data clustering: A review. *ACM Comput. Surv.*, 31(3):264–323, 1999.

[13] Daniel Kifer and Ashwin Machanavajjhala. No free lunch in data privacy. In Timos K. Sellis, Renée J. Miller, Anastasios Kementsietsidis, and Yannis Velegrakis, editors, *Proceedings of the ACM SIGMOD International Conference on Management of Data, SIGMOD 2011, Athens, Greece, June 12-16, 2011*, pages 193–204. ACM, 2011.

[14] Muah Kim, Onur Günlü, and Rafael F. Schaefer. Federated learning with local differential privacy: Trade-offs between privacy, utility, and communication. In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*, pages 2650–2654. IEEE, 2021.

[15] Kamran Kowsari, Kiana Jafari Meimandi, Mojtaba Heidarysafa, Sanjana Mendu, Laura E. Barnes, and Donald E. Brown. Text classification algorithms: A survey. *Inf.*, 10(4):150, 2019.

[16] Ninghui Li, Tiancheng Li, and Suresh Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In Rada Chirkova, Asuman Dogac, M. Tamer Özsu, and Timos K. Sellis, editors, *Proceedings of the 23rd International Conference on Data Engineering, ICDE 2007, The Marmara Hotel, Istanbul, Turkey, April 15-20, 2007*, pages 106–115. IEEE Computer Society, 2007.

[17] Zitao Li, Tianhao Wang, Milan Lopuhaä-Zwakenberg, Ninghui Li, and Boris Skoric. Estimating numerical distributions under local differential privacy. In David Maier, Rachel Pottinger, AnHai Doan, Wang-Chiew Tan, Abdussalam Alawini, and Hung Q. Ngo, editors, *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*, pages 621–635. ACM, 2020.

[18] Fang Liu. Statistical properties of sanitized results from differentially private laplace mechanism with univariate bounding constraints. *Trans. Data Priv.*, 12(3):169–195, 2019.

[19] D. Lu and Q. Weng. A survey of image classification methods and techniques for improving classification performance. *International Journal of Remote Sensing*, 28(5):823–870, 2007.

[20] Fei Ma, Renbo Zhu, and Ping Wang. PTT: piecewise transformation technique for analyzing numerical data under local differential privacy. *IEEE Trans. Mob. Comput.*, 23(10):9518–9531, 2024.

[21] Frank McSherry and Kunal Talwar. Mechanism design via differential privacy. In *48th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2007, Providence, RI, USA, October 20-23, 2007, Proceedings*, pages 94–103. IEEE Computer Society, 2007.

[22] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *2008 IEEE Symposium on Security and Privacy (SP 2008), 18-21 May 2008, Oakland, California, USA*, pages 111–125. IEEE Computer Society, 2008.

[23] Joseph P Near, David Darais, Naomi Lefkovitz, and Gary S Howarth. Guidelines for evaluating differential privacy guarantees.

[24] Natalia Ponomareva, Hussein Hazimeh, Alex Kurakin, Zheng Xu, Carson Denison, H. Brendan McMahan, Sergei Vassilvitskii, Steve Chien, and Abhradeep Guha Thakurta. How to dp-fy ML: A practical guide to machine learning with differential privacy. *J. Artif. Intell. Res.*, 77:1113–1201, 2023.

[25] Gauri Pradhan, Joonas Jälkö, Santiago Zanella-Bèguelin, and Antti Honkela. Beyond membership: Limitations of add/remove adjacency in differential privacy.

[26] Apple Machine Learning Research. Learning with privacy at scale, 2017.

[27] R L Rivest, L Adleman, and M L Dertouzos. On data banks and privacy homomorphisms. *Foundations of Secure Computation, Academia Press*, pages 169–179, 1978.

[28] Sharmistha Chatterjee Sapient, Senior Manager Data Sciences at Publicis. A guide to differential privacy at scale, December 2020.

[29] Latanya Sweeney. k-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl. Based Syst.*, 10(5):557–570, 2002.

[30] Han Wang, Hanbin Hong, Li Xiong, Zhan Qin, and Yuan Hong. L-SRR: local differential privacy for location-based services with staircase randomized response. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *Proceedings of the 2022 ACM SIGSAC Conference on*

Computer and Communications Security, CCS 2022, Los Angeles, CA, USA, November 7-11, 2022, pages 2809–2823. ACM, 2022.

[31] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. Collecting and analyzing multidimensional data with local differential privacy. In *35th IEEE International Conference on Data Engineering, ICDE 2019, Macao, China, April 8-11, 2019*, pages 638–649. IEEE, 2019.

[32] Tianhao Wang, Jeremiah Blocki, Ninghui Li, and Somesh Jha. Locally differentially private protocols for frequency estimation. In Engin Kirda and Thomas Ristenpart, editors, *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, pages 729–745. USENIX Association, 2017.

[33] Tianhao Wang, Ninghui Li, and Somesh Jha. Locally differentially private frequent itemset mining. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 127–143. IEEE Computer Society, 2018.

[34] Tianhao Wang, Ninghui Li, and Somesh Jha. Locally differentially private heavy hitter identification. *IEEE Trans. Dependable Secur. Comput.*, 18(2):982–993, 2021.

[35] Stanley L. Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[36] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *23rd Annual Symposium on Foundations of Computer Science, Chicago, Illinois, USA, 3-5 November 1982*, pages 160–164. IEEE Computer Society, 1982.

[37] Nan Zhang, Shengquan Wang, and Wei Zhao. A new scheme on privacy-preserving data classification. In Robert Grossman, Roberto J. Bayardo, and Kristin P. Bennett, editors, *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Chicago, Illinois, USA, August 21-24, 2005*, pages 374–383. ACM, 2005.

[38] Yuemin Zhang, Qingqing Ye, Rui Chen, Haibo Hu, and Qilong Han. Trajectory data collection with local differential privacy. *Proc. VLDB Endow.*, 16(10):2591–2604, 2023.

[39] Ye Zheng, Sumita Mishra, and Yidan Hu. Optimal piecewise-based mechanism for collecting bounded numerical data under local differential privacy. *Proc. Priv. Enhancing Technol.*, 2025(4):146–165, 2025.

[40] Ezgi Zorarpaci and Selma Ayse Özel. Privacy preserving classification over differentially private data. *WIREs Data Mining Knowl. Discov.*, 11(3), 2021.

Appendices