Local Privacy-preserving Mechanisms and Applications in Machine Learning

Likun Qin, Tianshuo Qiu

Department of Electrical and Computer Engineering

Shandong University, Jinan, China

arXiv:2401.13692v1 [cs.CR] 8 Jan 2024

Abstract—**The emergence and evolution of Local Differential Privacy (LDP) and its various adaptations play a pivotal role in tackling privacy issues related to the vast amounts of data generated by intelligent devices, which are crucial for data-informed decision-making in the realm of crowdsensing. Utilizing these extensive datasets can provide critical insights but also introduces substantial privacy concerns for the individuals involved. LDP, noted for its decentralized framework, excels in providing strong privacy protection for individual users during the stages of data collection and processing. The core principle of LDP lies in its technique of altering each user's data locally at the client end before it is sent to the server, thus preventing privacy violations at both stages. There are many LDP variances in the privacy research community aimed to improve the utility-privacy tradeoff. On the other hand, one of the major applications of the privacy-preserving mechanisms is machine learning. In this paper, we firstly delves into a comprehensive analysis of LDP and its variances, focusing on their various models, the diverse range of its adaptations, and the underlying structure of privacy mechanisms; then we discuss the state-of-art privacy mechanisms applications in machine learning.**

I. INTRODUCTION

Collecting and analyzing data introduces significant privacy concerns because it often includes sensitive user information. With the advent of sophisticated data fusion and analysis methods, user data becomes even more susceptible to breaches and exposure in this era of big data. For instance, by studying appliance usage, adversaries can deduce daily routines or behaviors of individuals, like when they are home or their specific activities such as watching TV or cooking. It's crucial to prioritize the protection of personal data when gathering information from diverse devices. Currently, the European Union (EU) has released the GDPR [1], which oversees EU data protection laws for its citizens and outlines the specifics related to the handling of personal data. Similarly, the U.S. National Institute of Standards and Technology (NIST) is in the process of crafting privacy frameworks. These frameworks aim to more effectively recognize, evaluate, and address privacy risks, enabling individuals to embrace innovative technologies with increased trust and confidence [2], [3].

From a privacy-protection standpoint, differential privacy (DP) has been introduced over a decade ago [4], [5]. Recognized as a robust framework for safeguarding privacy, it's often termed as global DP or centralized DP. DP's strength lies in its mathematical rigor; it operates independent of an adversary's background knowledge and assures potent privacy protection for users. It has found applications across various domains [6].

However, DP assumes the presence of a trustworthy server, which can be a challenge since many online platforms or crowdsourcing systems might have untrustworthy servers keen on user data statistics [7], [8].

Emerging from the concept of DP, local differential privacy (LDP) was introduced [9]. LDP stands as a decentralized version of DP, offering individualized privacy assurances and making no assumptions about third-party server trustworthiness. LDP has become a focal point in privacy research due to its theoretical significance and practical implications [10]. Numerous corporations, including Apple's iOS [11], Google Chrome, and the Windows operating system, have integrated LDP-driven algorithms into their systems. Owing to its robust capabilities, LDP has become a preferred choice to address individual privacy concerns during various statistical and analytical operations. This includes tasks like frequency and mean value estimation [12], the identification of heavy hitters [13], k-way marginal release, empirical risk minimization (ERM), federated learning, and deep learning.

While LDP is powerful, it's not without its challenges, notably in striking an optimal balance between utility and privacy [14]. To address this, there are two primary approaches. Firstly, by devising improved mechanisms - leading to the introduction of numerous LDP-based protocols and sophisticated mechanisms in academic circles. Secondly, by revisiting the definition of LDP itself, with researchers suggesting more flexible privacy concepts to better cater to the utility-privacy balance required for real-world applications. Given the growing significance of LDP, a thorough survey of the topic is both timely and essential. While there exists some literature reviewing LDP, the focus has often been narrow. They either focus on specific applications or certain types of mechanisms.

In this paper, we delve deep into the world of LDP and its various offshoots, meticulously studying their recent advancements and associated mechanisms. We embark on a thorough exploration of the foundational principles that drive LDP and the evolutionary trajectories of its multiple variants. We aim to identify the cutting-edge developments, shedding light on the innovations that have shaped these privacy tools and the challenges they aim to address in our contemporary digital landscape. Furthermore, we analyze the specific mechanisms that support and enhance the capabilities of LDP, understanding their technical intricacies and the real-world applications they cater to. Through this comprehensive study, we aspire to provide readers with a panoramic view of

the current state of LDP research, setting the stage for future inquiries and innovations in this critical domain.

Recognizing the complexity of privacy issues in the digital age, we delve into how LDP mechanisms have been tailored and adapted to meet diverse needs and environments. We explore the balance between privacy preservation and the practical utility of data, a central challenge in the deployment of LDP.

Our survey extends to the innovative methodologies that have been developed within the LDP framework. These include refined algorithms and protocols that enhance privacy protection while maintaining data usefulness. We investigate how these advancements have been implemented in various real-world scenarios, from online platforms to smart device ecosystems, and their effectiveness in protecting user privacy.

Moreover, the paper examines the evolution of LDP in response to emerging technologies and changing data landscapes. We discuss how the concept of LDP has been broadened or altered to accommodate new types of data collection and analysis methods, including the growing field of federated learning and the increasing reliance on deep learning algorithms. In addressing these aspects, we also highlight the legal and ethical considerations surrounding LDP. The paper considers how LDP aligns with global data protection regulations like the GDPR and how it is influencing policy-making and privacy standards worldwide.

Finally, we present the state-of-the-art applications of the privacy models in machine learning. Specifically, we focus on supervised and unsupervised machine learning, Empirical Risk Minimization, Reinforcement Learning, Deep learning, and Federated Learning.

II. LOCAL DIFFERENTIAL PRIVACY, PROPERTIES AND MECHANISMS

This section explores Local Differential Privacy (LDP) and its associated mechanisms. We begin with an understanding of LDP's definition.

Definition 1: ε-Local Differential Privacy (ε-LDP) - A randomized mechanism M adheres to ε-LDP if, for any two input values v and v' within M's domain, and for any output y in Y, the following holds:

$$P[M(v) = y] \leq e^\varepsilon \times P[M(v') = y], \tag{1}$$

where $P[\cdot]$ signifies probability, and ε represents the privacy budget. Smaller ε values imply stronger privacy.

LDP's key characteristics are:

Composition - If two mechanisms, M_1 and M_2, offer ε_1-LDP and ε_2-LDP respectively, their combined use results in $(\varepsilon_1 + \varepsilon_2)$-LDP.

$$M(v) = (M_1(v), M_2(v)) \Rightarrow M \text{ provides } (\varepsilon_1 + \varepsilon_2)\text{-LDP} \tag{2}$$

Post-processing Immunity - Functions applied to an ε-LDP mechanism's output maintain the ε-LDP property.

If $M(v)$ offers ε-LDP, then $f(M(v))$ also ensures ε-LDP.
$$\tag{3}$$

Resistance to Auxiliary Information - LDP remains effective even if adversaries possess additional information.

Utility-Privacy Balancing - A lower ε generally means increased privacy but can reduce the data's usefulness. Independence from Background Knowledge - LDP's privacy assurances are unaffected by an adversary's prior knowledge. Subsequently, we examine LDP-compliant mechanisms: Randomized Response Mechanism - This technique, effective for binary data, involves a user responding truthfully or randomly with equal probability. Its probability mass function (pmf) ensures ε-LDP with $\varepsilon = \ln(2)$. Laplace Mechanism - This approach adds Laplace-distribution noise to data, with the noise level determined by the sensitivity of the function and ε. Gaussian Mechanism - Similar to the Laplace Mechanism, but uses Gaussian distribution noise. The noise amount is based on the desired ε and function sensitivity. Exponential Mechanism - Selects outputs based on a scoring function, with selection probability proportional to the exponential of their score. Perturbed Histogram Mechanism - Rather than altering each item, this method modifies the entire histogram of data items, adding Laplace-distribution noise. The effectiveness of each mechanism is closely tied to the query's sensitivity, denoted as Δf. In LDP, high sensitivity can necessitate substantial noise, potentially reducing data utility.

Increasing input support size complicates maintaining desired privacy levels, as high sensitivity noise can obscure actual data, leading to potential misinterpretation or meaningless results. This highlights a significant tradeoff between data utility and privacy. Stronger privacy often comes at the cost of accuracy and utility, posing challenges for precision-requiring applications. While LDP offers robust theoretical privacy, its practical application demands a delicate balance between utility and privacy, spurring ongoing research for more effective mechanisms or modified privacy models.

III. LDP VARIANTS AND MECHANISMS

In this section, we introduce LDP variants that aim to provide better utility-privacy tradeoff in different applications.

A. Variants and Mechanisms of LDP

1) (ε, δ)-LDP: Drawing parallels with how (ε, δ)-DP [15] extends ε-DP, (ε, δ)-LDP (sometimes termed as approximate LDP) serves as a more flexible counterpart to ε-LDP (or pure LDP).

Definition 1 (Approximate Local Differential Privacy)**.** *A randomized process M complies with (ε, δ)-LDP if, for all input pairs v and v' within M's domain and any probable output $y \in Y$, the following holds:*

$$P[M(v) = y] \leq e^\varepsilon \cdot P[M(v') = y] + \delta.$$

Here, δ is customarily a small value.

In essence, (ε, δ)-LDP implies that M achieves ε-LDP with a likelihood not less than $1 - \delta$. If $\delta = 0$, (ε, δ)-LDP converges to ε-LDP.

2) BLENDER Model: BLENDER [16], a fusion of global DP and LDP, optimizes data utility while retaining privacy. It classifies users based on their trust in the aggregator into two categories: the opt-in group and clients. BLENDER enhances utility by balancing data from both. Its privacy measure mirrors that of (ε, δ)-DP [17].

3) Geo-indistinguishability: Originally tailored for location privacy with global DP, Geo-indistinguishability [18] uses the data's geographical distance. Alvim et al. [19] argued for metric-based LDP's advantages in specific contexts.

Definition 2 (Geo-indistinguishability)**.** *A randomized function M adheres to Geo-indistinguishability if, for any input pairs v and v' and any output $y \in Y$, the subsequent relation is met:*

$$P[M(v) = y] \leq e^{\varepsilon \cdot d(v,v')} \cdot P[M(v') = y],$$

where $d(.,.)$ designates a distance metric.

This model adjusts privacy depending on data distance, augmenting utility for datasets like location or smart meter consumption that are sensitive to distance.

4) Local Information Privacy: Local Information Privacy (LIP) was originally proposed in [20] as a prior-aware version of LDP, and then, in [21], Jiang et al relax the prior-aware assumption to partial prior-aware (Bounded Prior in their version). The definition of LIP is shown as follows:

Definition 3. (ϵ, δ)-*Local Information Privacy [22] A mechanism M satisfies (ϵ, δ)-LIP, if $\forall x \in \mathcal{X}, y \in Range(\mathcal{M})$:*

$$\begin{aligned} P(Y = y) &\geq e^{-\epsilon} P(Y = y | X = x) - \delta, \\ P(Y = y) &\leq e^{\epsilon} P(Y = y | X = x) + \delta. \end{aligned} \quad (4)$$

The operational meaning of LIP is, the output Y provides limited additional information about any possible input X, and the amount of the additional information is measured by the privacy budget ϵ and failure probability δ.

In [23], multiple LIP mechanisms were proposed and testified, showing that even though ϵ-LIP is stronger than 2ϵ-LDP in terms of privacy protection. The mechanisms achieve more than 2 times of utility gain.

5) CLDP: Recognizing LDP's diminished utility with fewer users, Gursoy et al. [24] introduced the metric-based model of condensed local differential privacy (CLDP).

Definition 4 (α-CLDP)**.** *For all input pairs v and v' in M's domain and any potential output $y \in Y$, a randomized function M satisfies α-CLDP if:*

$$P[M(v) = y] \leq e^{\alpha \cdot d(v,v')} \cdot P[M(v') = y],$$

where $\alpha > 0$.

In CLDP, a decline in α compensates for a growth in distance d. Gursoy et al. employed an Exponential Mechanism variant to devise protocols, particularly benefitting scenarios with limited users.

6) PLDP: PLDP [25] offers user-specific privacy levels. Here, users can modify their privacy settings, denoted by ε.

Definition 5 (ε-PLDP)**.** *For a user U, and all input pairs v and v' in M's domain and any potential output $y \in Y$, a randomized function M meets ε_U-PLDP if:*

$$P[M(v) = y] \leq e^{\varepsilon_U} \cdot P[M(v') = y].$$

Approaches like the personalized count estimation protocol and advanced combination strategy cater to users with varying privacy inclinations.

7) Utility-optimized LDP (ULDP): Traditional LDP assumes all data points have uniform sensitivity, often causing excessive noise addition. Recognizing that not all personal data have equivalent sensitivity, the Utility-optimized LDP (ULDP) model was proposed. In this model, let $KS \subseteq K$ be the sensitive dataset and $KN = K \setminus KS$ be the non-sensitive dataset. Let $Y_P \subseteq Y$ be the protected output set and $Y_I = Y \setminus Y_P$ be the invertible output set. The formal definition of ULDP is:

Definition 6. *Given $KS \subseteq K$, $Y_P \subseteq Y$, a mechanism M adheres to (KS, Y_P, ϵ)-ULDP if:*

- *For every $y \in Y_I$, there is a $v \in KN$ with $P[M(v) = y] > 0$ and $P[M(v') = y] = 0$ for any $v' \neq v$.*
- *For all $v, v' \in K$ and $y \in Y_P$, $P[M(v) = y] \leq e^{\epsilon} \cdot P[M(v') = y]$.*

In simpler terms, (KS, Y_P, ϵ)-ULDP ensures that sensitive inputs are mapped only to the protected output set.

8) Input-Discriminative LDP (ID-LDP): While ULDP classifies data as either sensitive or non-sensitive, the ID-LDP model offers a more nuanced approach by acknowledging varying sensitivity levels among data. It is defined as:

Definition 7. *Given a set of privacy budgets $E = \{\epsilon_v\}_{v \in K}$, a mechanism M adheres to E-ID-LDP if for all input pairs v and v', and any output $y \in Y$:*

$$P[M(v) = y] \leq e^{r(\epsilon_v, \epsilon_{v'})} \cdot P[M(v') = y]$$

where $r(\cdot, \cdot)$ is a function of two privacy budgets.

The study in [26] primarily utilizes the minimum function between ϵ_v and $\epsilon_{v'}$ and introduces the MinID-LDP as a specialized case.

9) Parameter Blending Privacy (PBP): PBP was proposed as a more flexible LDP variant [27]. In PBP, let Θ represent the domain of privacy parameters. Given a privacy budget $\theta \in \Theta$, let $P(\theta)$ denote the frequency with which θ is selected. PBP is defined as:

Definition 8. *A mechanism M adheres to r-PBP if, for all $\theta \in \Theta, v, v' \in K, y \in Y$, there exists a $\theta' \in \Theta$ such that:*

$$P(\theta)P[M(v; \theta) = y] \leq e^{r(\theta)} \cdot P(\theta')P[M(v'; \theta') = y]$$

10) Perfect Privacy: Another variant yet widely investigated privacy notion is perfect privacy. Perfect privacy can be understood as a special case of DP(LDP) when $\epsilon = 0$, or simply, the input is independent of the output.

Definition 9. *A mechanism M achieves perfect privacy for the input if, for all $x \in \mathcal{X}$:*

$$P(X|Y) = P(X).$$

From the information perspective, perfect privacy can also be understood as the mutual information between the input and the output is zero. Since perfect privacy means the input and the output is independent, it provides no utility in general. However, recently there are several adaptions of perfect privacy mechanisms that protecting a latent variable of the input. Hence, when the latent variable is independent of the output, the input can still be correlated to it and therefore provides certain utility [28], [29].

IV. Local Privacy Mechanisms for Machine Learning

Machine learning, a critical method in data analysis, finds application across various fields. However, it's susceptible to several types of attacks during the training process. These include membership inference attacks [30], model inversion attacks [31], and memorizing model attacks [32]. For instance, adversaries could exploit these vulnerabilities to access sensitive user data, as demonstrated in cases where memorized information is extracted during training [30]. A particularly concerning example by Fredrikson et al. [31] highlighted how facial recognition systems can be compromised under model inversion attacks, thereby exposing the frailties of a trained machine learning model.

Extensive research has been conducted on machine learning algorithms incorporating global Differential Privacy (DP) through private training [33]–[35]. The advent of Local Differential Privacy (LDP) further led to investigations into machine learning algorithms under the LDP framework, aiming to enhance privacy protection in a distributed manner. The subsequent subsections provide a comprehensive overview of existing machine learning algorithms that utilize LDP. These algorithms are categorized and summarized from various perspectives, including supervised learning, unsupervised learning, empirical risk minimization, deep learning, reinforcement learning, and federated learning.

A. Unsupervised Learning

Clustering, a fundamental problem in data analysis, has been explored under centralized Differential Privacy (DP) [36], [37]. In the realm of Local Differential Privacy (LDP), Nissim and Stemmer [38] studied 1-clustering by identifying a minimum enclosing ball. Additionally, Sun et al. [39] delved into non-interactive clustering under LDP, extending the Bit Vector mechanism from [40]. They enhanced the encoding process and introduced the $k\text{Cluster}$ algorithm for clustering in an anonymous space. Further, Li et al. [41] proposed a mechanism for local-clustering-based collaborative filtering, employing the kNN algorithm to categorize items while ensuring item-level privacy.

In local clustering, individuals randomize their data before sharing it with an untrusted data curator, aiming for stronger privacy protection. While local clustering may not be as accurate as its central counterpart, it offers enhanced user privacy and is more adaptable to varied privacy specifications, like personalized privacy parameters [42]. Xia et al. [43] applied LDP to K-means clustering, perturbing user data directly. They introduced a budget allocation scheme to mitigate noise and enhance accuracy. Despite these advances, research on clustering under LDP is still in its nascent stages.

B. Supervised Learning

Supervised learning algorithms are designed to train models for predicting data classes using labeled datasets. In this domain, Yilmaz et al. [44] proposed a method for training a Naïve Bayes classifier under Local Differential Privacy (LDP). This classifier aims to determine the most probable label for a new instance by maintaining the relationship between feature values and class labels during data perturbation. Yilmaz et al. achieved this by first transforming each user's data and label into a new value, followed by LDP perturbation. Similarly, Xue et al. [45] also focused on training a Naïve Bayes classifier with LDP, utilizing joint distributions to compute conditional distributions. Additionally, Berrett and Butucea [46] explored binary classification problems within the LDP framework.

High dimensionality presents a significant challenge in training classifiers with LDP, often leading to increased time costs and reduced accuracy. Traditional solutions like Principal Component Analysis (PCA) [47] are commonly used for dimensionality reduction, but effective methods compatible with LDP in machine learning require further exploration. Another common approach in LDP-based model learning involves partitioning users, as seen in the work of Xue et al. [45], who divided users into groups for different calculations. However, this simple partitioning can compromise estimation accuracy. Therefore, the field of supervised learning with LDP is still evolving and requires more in-depth research.

C. Empirical Risk Minimization

Empirical risk minimization (ERM) in machine learning refers to the process of calculating an optimal model from a set of parameters by minimizing the expected loss [48]. The loss function $\mathcal{L}(\theta; x, y)$, parameterized by x and y, maps the parameter vector θ to a real number. ERM is applicable to various learning tasks like logistic regression, linear regression, and support vector machine (SVM) by selecting appropriate loss functions. Existing literature discusses both interactive and non-interactive models under Local Differential Privacy (LDP) for natural learning problems, with the interactive model offering better accuracy but at the cost of higher network delay and weaker privacy guarantees. In contrast, the non-interactive model is more robust and practical in most settings.

Smith et al. [49] pioneered the study of interaction in LDP for natural learning problems, highlighting the network delays caused by sequential information exchange in convex optimization. They explored the necessity of interactivity for optimizing convex functions and introduced new algorithms that are either non-interactive or require minimal interaction.

Building on this, Zheng et al. [50] developed more efficient algorithms using Chebyshev expansion in a non-interactive LDP setting, achieving quasi-polynomial sample complexity.

However, the sample complexity in high dimensions remains a challenge, as noted by Wang et al. [32], who proposed LDP algorithms with error bounds based on Gaussian width. They improved upon Smith et al.'s work, but the sample complexity was still exponential with dimensionality. Their subsequent work [51] enhanced this to quasi-polynomial complexity. For generalized linear models (GLM), Wang et al. [52] demonstrated that LDP algorithms can attain fully polynomial sample complexity when the GLM's feature vector is sub-Gaussian with a bounded ℓ_1-norm. Additionally, Wang and Xu [53] addressed the principal component analysis (PCA) problem under non-interactive LDP and established the lower bounds of the minimax risk in both low and high dimensional settings.

Classical machine learning models under LDP are typically built using ERM, solved via stochastic gradient descent (SGD). This approach is applied to linear regression, logistic regression, and SVM classification. With LDP, the SGD algorithm updates the parameter vector θ iteratively, ensuring privacy by perturbing each gradient into a noisy version before aggregation. In the context of data center networks (DCN), Fan et al. [54] examined LDP-based support vector regression (SVR) classification for cloud-supported data centers, employing the Laplace mechanism for LDP. Yin et al. [55] approached LDP-based logistic regression classification through steps of noise addition, feature selection, and logistic regression. LDP has also been applied to online convex optimization by van2019user, ensuring privacy in adaptive online learning without disclosing parameters. Furthermore, Jun and Orabona [56] investigated the parameter-free SGD problem under LDP, proposing BANCO, which achieves the convergence rate of tuned SGD without multiple runs, thereby reducing privacy loss and conserving the privacy budget.

D. Reinforcement Learning

Reinforcement learning, a key component in artificial intelligence (AI), facilitates interactive model learning for an agent but is susceptible to attacks that can cause significant privacy breaches [57].

Addressing these concerns, Gajane et al. [58] were among the first to explore the multi-armed bandits (MAB) problem in the context of Local Differential Privacy (LDP). They developed an LDP-compliant bandit algorithm targeted at arms with Bernoulli rewards. Building on this, Basu et al. [59] introduced a set of foundational privacy definitions for MAB algorithms, incorporating both the graphical and LDP models. Their work provided insights into both distribution-dependent and distribution-free regret lower bounds.

In the realm of distributed reinforcement learning, Ono and Takahashi [60] proposed the Private Gradient Collection (PGC) framework. PGC enables the private learning of a model using noisy gradients, where each local agent submits perturbed gradients in compliance with LDP to a central aggregator, which then updates the global parameters. Furthering

this research, Ren et al. [61] focused on regret minimization in MAB problems under LDP, establishing a tight regret lower bound. They introduced two algorithms that achieve LDP, one based on Laplace perturbation and the other on Bernoulli response.

Reinforcement learning is a pivotal aspect of AI [62], and LDP emerges as a promising technique to protect sensitive information within this field. However, the application of LDP in reinforcement learning is still in its early stages, indicating a significant potential for future developments.

E. Deep Learning

Deep learning has significantly impacted fields like natural language processing and image classification. Yet, it is vulnerable to adversaries who can inject malicious algorithms during training to approximate and extract sensitive user information [63].

To mitigate these risks, Arachchige et al. [64] developed an LDP-compliant mechanism called LATENT for deep learning models, such as convolutional neural networks (CNNs). LATENT, like other LDP frameworks, incorporates a randomization layer (LDP layer) to protect against untrusted learning servers. A primary challenge in applying LDP to deep learning is the extremely high sensitivity, which in LATENT is expressed as lr, where l is the binary string length and r is the number of neural network layers. To overcome this, Arachchige et al. enhanced the One-time Unary Encoding (OUE) [65], which has a sensitivity of 2, by introducing Modified OUE (MOUE). MOUE offers greater control over the randomization of 1s and increases the likelihood of retaining 0 bits in their original state. They also proposed the Utility Enhancing Randomization (UER) mechanism, which further improves the utility of randomized binary strings.

Additionally, Zhao [66] explored distributed deep learning under Differential Privacy (DP) using the teacher-student paradigm, allowing for personalized privacy parameter choices for each distributed data entity under LDP. Xu et al. [67] implemented LDP in a deep inference-based edge computing framework, enabling the private construction of complex deep learning models. Despite these advancements, research in deep learning with LDP is still in its infancy. Future research is necessary to enhance privacy protection, address high dimensionality challenges, and improve accuracy.

F. Federated Learning

Federated learning (FL), a pivotal technology for advancing modern artificial intelligence (AI), offers a collaborative learning framework for multiple data owners or parties [68], [69]. While FL adeptly balances utility and privacy in machine learning [70], transmitting or exchanging model parameters can still lead to privacy breaches. Hence, Local Differential Privacy (LDP) has been widely implemented in FL systems to provide robust privacy assurances, as evidenced in applications like smart electric power systems [71] and wireless channels [72]. Early studies in FL predominantly used global Differential Privacy (DP) [73] for safeguarding sensitive information. However, given FL's distributed nature,

LDP is more suited to FL systems. Truex et al. [74] integrated LDP into FL for the joint training of deep neural networks, offering a method that efficiently handles complex models and defends against inference attacks while providing personalized LDP. Additionally, Wang et al. [75] introduced FedLDA, an LDP-based latent Dirichlet allocation (LDA) model within an FL framework. FedLDA uses a novel random response with prior, enhancing privacy without dependency on the dictionary size and improving accuracy through adaptive, non-uniform sampling.

Bhowmick et al. [76] proposed a relaxed optimal LDP mechanism for private FL to enhance model fitting and prediction. Li et al. [77] developed an efficient LDP algorithm for meta-learning, applicable to personalized FL. In federated Stochastic Gradient Descent (SGD), LDP has been employed to protect gradient privacy. However, as the dimensionality d increases, the privacy budget depletes rapidly, and noise escalation leads to diminished model accuracy for large d. To address this, Liu et al. [78] introduced FedSel, which selects only the top-k most important dimensions, stabilizing the learning process. Sun et al. [79] suggested improving accuracy by reducing noise variance through splitting and shuffling, thus mitigating privacy degradation.

Naseri et al. [80] recently presented an analytical framework to empirically evaluate the effectiveness of LDP and CDP in FL. Their findings indicate that both LDP and global DP can counter backdoor attacks but are less effective against property inference attacks.

V. CONCLUSION

In the field of data privacy, Local Differential Privacy (LDP) is an essential tool for protecting user data. Our research covers a wide range of LDP mechanisms and variations, each designed for specific challenges and data types, from BLENDER for categorical data to specialized versions like CLDP for small datasets. We also explore how LDP is applied in different machine learning areas, such as supervised learning, deep learning, and federated learning, highlighting its adaptability. The choice of LDP variant is critical, depending on the data and privacy requirements. Overall, LDP's versatility makes it a key player in ensuring data privacy across various contexts.

REFERENCES

[1] European Parliament and Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council. [Online]. Available: https://data.europa.eu/eli/reg/2016/679/oj

[2] R. MEYER. (2014) Facebook's mood manipulation experiment might have been illegal.

[3] "No free lunch in data privacy," in *Proceedings of SIGMOD 2011 and PODS 2011*, ser. Proceedings of the ACM SIGMOD International Conference on Management of Data. Association for Computing Machinery, 2011, pp. 193–204.

[4] C. Dwork, F. McSherry, and K. Nissim, "Calibrating noise to sensitivity in private data analysis," in *Theory of Cryptography: Third Theory of Cryptography Conference*, 2006, pp. 265–284. [Online]. Available: https://doi.org/10.1007/11681878_14

[5] C. Dwork, "Differential privacy," in *Automata, Languages and Programming: 33rd International Colloquium, ICALP 2006, Part II*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds., 2006, pp. 1–12. [Online]. Available: https://doi.org/10.1007/11787006_1

[6] Y. Xiao and L. Xiong, "Protecting locations with differential privacy under temporal correlations," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15, New York, NY, USA, 2015, p. 1298–1309.

[7] W. Primoff and S. Kess, "The equifax data breach: What cpas and firms need to know now," *The CPA Journal*, vol. 87, no. 12, pp. 14–17, 2017.

[8] J. Lu, "Assessing the cost, legal fallout of capital one data breach," *Legal Fallout Of Capital One Data Breach (August 15, 2019)*, 2019.

[9] C. Dwork, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation: 5th International Conference, TAMC*, M. Agrawal, D. Du, and Z. Duan, Eds., 2008, pp. 1–19. [Online]. Available: https://doi.org/10.1007/978-3-540-79228-4_1

[10] Úlfar Erlingsson, V. Pihur, and A. Korolova, "Rappor: Randomized aggregatable privacy-preserving ordinal response," in *Proceedings of the 21st ACM CCS*, 2014.

[11] A. Greenberg, "Apple's 'differential privacy' is about collecting your data—but not your data," 2016. [Online]. Available: https://www.wired.com/2016/06/apples-differential-privacy-collecting-data/

[12] T. Wang, J. Blocki, N. Li, and S. Jha, "Locally differentially private protocols for frequency estimation," in *26th USENIX Security 17*. USENIX Association, 2017, pp. 729–745. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/wan

[13] Z. Qin, Y. Yang, and T. Yu, "Heavy hitter estimation over set-valued data with local differential privacy," in *Proceedings of the 2016 ACM SIGSAC*, ser. CCS '16, 2016, pp. 192–203. [Online]. Available: http://doi.acm.org/10.1145/2976749.2978409

[14] C. Huang, P. Kairouz, X. Chen, L. Sankar, and R. Rajagopal, "Context-aware generative adversarial privacy," *CoRR*, vol. abs/1710.09549, 2017. [Online]. Available: http://arxiv.org/abs/1710.09549

[15] R. Bassily, "Linear queries estimation with local differential privacy," *CoRR*, vol. abs/1810.02810, 2018. [Online]. Available: http://arxiv.org/abs/1810.02810

[16] B. Avent, A. Korolova, D. Zeber, T. Hovden, and B. Livshits, "BLENDER: Enabling local search with a hybrid differential privacy model," in *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 747–764. [Online]. Available: https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/ave

[17] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, ser. Lecture Notes in Computer Science, S. Vaudenay, Ed., vol. 4004. Springer, 2006, pp. 486–503. [Online]. Available: https://doi.org/10.1007/11761679_29

[18] M. E. Andrés, N. E. Bordenabe, and K. Chatzikokolakis, "Geo-indistinguishability: Differential privacy for location-based systems," *CoRR*, vol. abs/1212.1984, 2012. [Online]. Available: http://arxiv.org/abs/1212.1984

[19] M. S. Alvim, K. Chatzikokolakis, C. Palamidessi, and A. Pazii, "Metric-based local differential privacy for statistical applications," *CoRR*, vol. abs/1805.01456, 2018. [Online]. Available: http://arxiv.org/abs/1805.01456

[20] B. Jiang, M. Li, and R. Tandon, "Context-Aware data aggregation with localized information privacy," in *2018 IEEE Conference on Communications and Network Security (CNS)*, May 2018.

[21] ——, "Local information privacy with bounded prior," in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*, May 2019, pp. 1–7.

[22] ——, "Local information privacy and its application to privacy-preserving data aggregation," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 3, pp. 1918–1935, 2022.

[23] B. Jiang, M. Seif, R. Tandon, and M. Li, "Context-aware local information privacy," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 3694–3708, 2021.

[24] M. E. Gursoy, A. Tamersoy, S. Truex, W. Wei, and L. Liu, "Secure and utility-aware data collection with condensed local differential privacy," *CoRR*, vol. abs/1905.06361, 2019. [Online]. Available: http://arxiv.org/abs/1905.06361

[25] Y. NIE, W. Yang, L. Huang, X. Xie, Z. Zhao, and S. Wang, "A utility-optimized framework for personalized private histogram estimation," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 4, pp. 655–669, 2019.

[26] T. Murakami and Y. Kawamoto, "Utility-optimized local differential privacy mechanisms for distribution estimation," in *Proceedings of the*

28th USENIX Conference on Security Symposium, ser. SEC'19. USA: USENIX Association, 2019, p. 1877–1894.

[27] S. Takagi, Y. Cao, and M. Yoshikawa, "Poster: Data collection via local differential privacy with secret parameters," in *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security*, ser. ASIA CCS '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 910–912. [Online]. Available: https://doi.org/10.1145/3320269.3405441

[28] A. Machanavajjhala and J. Gehrke, "On the efficiency of checking perfect privacy," in *Proceedings of the Twenty-Fifth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, ser. PODS '06. New York, NY, USA: Association for Computing Machinery, 2006, p. 163–172. [Online]. Available: https://doi.org/10.1145/1142351.1142375

[29] B. Jiang, M. Seif, R. Tandon, and M. Li, "Answering count queries for genomic data with perfect privacy," *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 3862–3875, 2023.

[30] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *2017 IEEE Symposium on Security and Privacy (SP)*, 2017, pp. 3–18.

[31] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 1322–1333. [Online]. Available: https://doi.org/10.1145/2810103.2813677

[32] C. Song, T. Ristenpart, and V. Shmatikov, "Machine learning models that remember too much," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 587–601. [Online]. Available: https://doi.org/10.1145/3133956.3134077

[33] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep learning with differential privacy," in *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '16. New York, NY, USA: Association for Computing Machinery, 2016, p. 308–318. [Online]. Available: https://doi.org/10.1145/2976749.2978318

[34] J. Zhao, Y. Chen, and W. Zhang, "Differential privacy preservation in deep learning: Challenges, opportunities and solutions," *IEEE Access*, vol. 7, pp. 48 901–48 911, 2019.

[35] J. Lee and D. Kifer, "Concentrated differentially private gradient descent with adaptive per-iteration privacy budget," in *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, ser. KDD '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 1656–1665. [Online]. Available: https://doi.org/10.1145/3219819.3220076

[36] K. Nissim, U. Stemmer, and S. Vadhan, "Locating a small cluster privately," in *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, 2016, pp. 413–427.

[37] D. Su, J. Cao, N. Li, E. Bertino, and H. Jin, "Differentially private k-means clustering," in *Proceedings of the sixth ACM conference on data and application security and privacy*, 2016, pp. 26–37.

[38] K. Nissim and U. Stemmer, "Clustering algorithms for the centralized and local models," in *Algorithmic Learning Theory*. PMLR, 2018, pp. 619–653.

[39] L. Sun, J. Zhao, and X. Ye, "Distributed clustering in the anonymized space with local differential privacy," *arXiv preprint arXiv:1906.11441*, 2019.

[40] D. Karapiperis, A. Gkoulalas-Divanis, and V. S. Verykios, "Distance-aware encoding of numerical values for privacy-preserving record linkage," in *2017 IEEE 33rd International Conference on Data Engineering (ICDE)*. IEEE, 2017, pp. 135–138.

[41] Y. Li, S. Liu, J. Wang, and M. Liu, "A local-clustering-based personalized differential privacy framework for user-based collaborative filtering," in *Database Systems for Advanced Applications: 22nd International Conference, DASFAA 2017, Suzhou, China, March 27-30, 2017, Proceedings, Part I 22*. Springer, 2017, pp. 543–558.

[42] M. Akter and T. Hashem, "Computing aggregates over numeric data with personalized local differential privacy," in *Information Security and Privacy: 22nd Australasian Conference, ACISP 2017, Auckland, New Zealand, July 3–5, 2017, Proceedings, Part II 22*. Springer, 2017, pp. 249–260.

[43] C. Xia, J. Hua, W. Tong, and S. Zhong, "Distributed k-means clustering guaranteeing local differential privacy," *Computers & Security*, vol. 90, p. 101699, 2020.

[44] E. Yilmaz, M. Al-Rubaie, and J. M. Chang, "Locally differentially private naive bayes classification," *CoRR*, vol. abs/1905.01039, 2019. [Online]. Available: http://arxiv.org/abs/1905.01039

[45] Q. Xue, Y. Zhu, and J. Wang, "Joint distribution estimation and naïve bayes classification under local differential privacy," *IEEE Transactions on Emerging Topics in Computing*, vol. 9, no. 4, pp. 2053–2063, 2021.

[46] T. Berrett and C. Butucea, "Classification under local differential privacy," *arXiv preprint arXiv:1912.04629*, 2019.

[47] S. Y. Kung, *Kernel methods and machine learning*. Cambridge University Press, 2014.

[48] K. Chaudhuri, C. Monteleoni, and A. D. Sarwate, "Differentially private empirical risk minimization." *Journal of Machine Learning Research*, vol. 12, no. 3, 2011.

[49] A. Smith, A. Thakurta, and J. Upadhyay, "Is interaction necessary for distributed private learning?" in *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2017, pp. 58–77.

[50] K. Zheng, W. Mou, and L. Wang, "Collect at once, use effectively: Making non-interactive locally private learning possible," in *International Conference on Machine Learning*. PMLR, 2017, pp. 4130–4139.

[51] D. Wang, C. Chen, and J. Xu, "Differentially private empirical risk minimization with non-convex loss functions," in *International Conference on Machine Learning*. PMLR, 2019, pp. 6526–6535.

[52] D. Wang, H. Zhang, M. Gaboardi, and J. Xu, "Estimating smooth glm in non-interactive local differential privacy model with public unlabeled data," in *Algorithmic Learning Theory*. PMLR, 2021, pp. 1207–1213.

[53] D. Wang and J. Xu, "Principal component analysis in the local differential privacy model," *Theoretical computer science*, vol. 809, pp. 296–312, 2020.

[54] W. Fan, J. He, M. Guo, P. Li, Z. Han, and R. Wang, "Privacy preserving classification on local differential privacy in data centers," *Journal of Parallel and Distributed Computing*, vol. 135, pp. 70–82, 2020.

[55] C. Yin, B. Zhou, Z. Yin, and J. Wang, "Local privacy protection classification based on human-centric computing," *Human-centric Computing and Information Sciences*, vol. 9, pp. 1–14, 2019.

[56] K.-S. Jun and F. Orabona, "Parameter-free locally differentially private stochastic subgradient descent," *arXiv preprint arXiv:1911.09564*, 2019.

[57] X. Pan, W. Wang, X. Zhang, B. Li, J. Yi, and D. Song, "How you act tells a lot: Privacy-leaking attack on deep reinforcement learning." in *AAMAS*, vol. 19, no. 2019, 2019, pp. 368–376.

[58] P. Gajane, T. Urvoy, and E. Kaufmann, "Corrupt bandits for preserving local privacy," in *Algorithmic Learning Theory*. PMLR, 2018, pp. 387–412.

[59] D. Basu, C. Dimitrakakis, and A. Tossou, "Differential privacy for multi-armed bandits: What is it and what is its cost?" *arXiv preprint arXiv:1905.12298*, 2019.

[60] H. Ono and T. Takahashi, "Locally private distributed reinforcement learning," *arXiv preprint arXiv:2001.11718*, 2020.

[61] W. Ren, X. Zhou, J. Liu, and N. B. Shroff, "Multi-armed bandits with local differential privacy," *arXiv preprint arXiv:2007.03121*, 2020.

[62] Y. Li, "Deep reinforcement learning: An overview," *arXiv preprint arXiv:1701.07274*, 2017.

[63] S. A. Osia, A. S. Shamsabadi, S. Sajadmanesh, A. Taheri, K. Katevas, H. R. Rabiee, N. D. Lane, and H. Haddadi, "A hybrid deep learning architecture for privacy-preserving mobile analytics," *IEEE Internet of Things Journal*, vol. 7, no. 5, pp. 4505–4518, 2020.

[64] P. C. M. Arachchige, P. Bertok, I. Khalil, D. Liu, S. Camtepe, and M. Atiquzzaman, "Local differential privacy for deep learning," *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 5827–5842, 2019.

[65] T. Wang, N. Li, and S. Jha, "Locally differentially private frequent itemset mining," in *2018 IEEE Symposium on Security and Privacy (SP)*, vol. 00, pp. 578–594.

[66] J. Zhao, "Distributed deep learning under differential privacy with the teacher-student paradigm," in *Workshops at the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[67] C. Xu, J. Ren, L. She, Y. Zhang, Z. Qin, and K. Ren, "Edgesanitizer: Locally differentially private deep inference at the edge for mobile data analytics," *IEEE Internet of Things Journal*, vol. 6, no. 3, pp. 5140–5151, 2019.

[68] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.

[69] Q. Li, Z. Wen, and B. He, "Federated learning systems: Vision, hype and reality for data privacy and protection. arxiv 2019," *arXiv preprint arXiv:1907.09693*.

[70] H. Zheng, H. Hu, and Z. Han, "Preserving user privacy for machine learning: Local differential privacy or federated machine learning?" *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 5–14, 2020.

[71] Y. Zhao, J. Zhao, M. Yang, T. Wang, N. Wang, L. Lyu, D. Niyato, and K.-Y. Lam, "Local differential privacy-based federated learning for internet of things," *IEEE Internet of Things Journal*, vol. 8, no. 11, pp. 8836–8853, 2020.

[72] M. Seif, R. Tandon, and M. Li, "Wireless federated learning with local differential privacy," in *2020 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2020, pp. 2604–2609.

[73] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.

[74] M. Kim, O. Günlü, and R. F. Schaefer, "Federated learning with local differential privacy: Trade-offs between privacy, utility, and communication," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 2650–2654.

[75] D. Wang, M. Gaboardi, and J. Xu, "Empirical risk minimization in non-interactive local differential privacy revisited," *Advances in Neural Information Processing Systems*, vol. 31, 2018.

[76] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," *arXiv preprint arXiv:1812.00984*, 2018.

[77] J. Li, M. Khodak, S. Caldas, and A. Talwalkar, "Differentially private meta-learning," *arXiv preprint arXiv:1909.05830*, 2019.

[78] R. Liu, Y. Cao, M. Yoshikawa, and H. Chen, "Fedsel: Federated sgd under local differential privacy with top-k dimension selection," in *Database Systems for Advanced Applications: 25th International Conference, DASFAA 2020, Jeju, South Korea, September 24–27, 2020, Proceedings, Part I 25*. Springer, 2020, pp. 485–501.

[79] L. Sun, J. Qian, and X. Chen, "Ldp-fl: Practical private aggregation in federated learning with local differential privacy," *arXiv preprint arXiv:2007.15789*, 2020.

[80] M. Naseri, J. Hayes, and E. De Cristofaro, "Toward robustness and privacy in federated learning: Experimenting with local and central differential privacy," *arXiv preprint arXiv:2009.03561*, 2020.