# Novelty Assessment Report

**Paper**: INO-SGD: Addressing Utility Imbalance under Individualized Differential Privacy

**PDF URL**: https://openreview.net/pdf?id=HMapYMkcrl

**Venue**: ICLR 2026 Conference Submission

**Year**: 2026

**Report Generated**: 2025-12-29

## Abstract

Differential privacy (DP) is widely employed in machine learning to protect confidential or sensitive training data from being revealed. As data owners gain greater control over their data due to personal data ownership, they are more likely to set their own privacy requirements, necessitating individualized DP (IDP) to fulfil such requests. In particular, owners of data from more sensitive subsets, such as positive cases of stigmatized diseases, likely set stronger privacy requirements, as leakage of such data could incur more serious societal impact. However, existing IDP algorithms induce a critical utility imbalance problem: Data from owners with stronger privacy requirements may be severely underrepresented in the trained model, resulting in poorer performance on similar data from subsequent users during deployment. In this paper, we analyze this problem and propose the INO-SGD algorithm, which strategically down-weights data within each batch to improve performance on the more private data across all iterations. Notably, our algorithm is specially designed to satisfy IDP, while existing techniques addressing utility imbalance neither satisfy IDP nor can be easily adapted to do so. Lastly, we demonstrate the empirical feasibility of our approach.

**Disclaimer**

This report is **AI-GENERATED** using Large Language Models and WisPaper (a scholar search engine). It analyzes academic papers' tasks and contributions against retrieved prior work. While this system identifies **POTENTIAL** overlaps and novel directions, **ITS COVERAGE IS NOT EXHAUSTIVE AND JUDGMENTS ARE APPROXIMATE**. These results are intended to assist human reviewers and **SHOULD NOT** be relied upon as a definitive verdict on novelty.

Note that some papers exist in multiple, slightly different versions (e.g., with different titles or URLs). The system may retrieve several versions of the same underlying work. The current automated pipeline does not reliably align or distinguish these cases, so human reviewers will need to disambiguate them manually.

If you have any questions, please contact: mingzhang23@m.fudan.edu.cn

## Core Task Landscape

This paper addresses: **Addressing Utility Imbalance under Individualized Differential Privacy**

A total of **41 papers** were analyzed and organized into a taxonomy with **10 categories**.

### Taxonomy Overview

The research landscape has been organized into the following main categories:

- **Individualized and Personalized Privacy Mechanisms**
- **Fairness and Utility Imbalance under Differential Privacy**
- **Privacy-Utility Trade-offs and Empirical Interactions**
- **Domain-Specific and Application-Oriented DP Frameworks**

### Complete Taxonomy Tree

- Addressing Utility Imbalance under Individualized Differential Privacy Survey Taxonomy
- Individualized and Personalized Privacy Mechanisms
  - Personalized Local Differential Privacy for Data Collection (4 papers)
  - [12] Local differential privacy-based federated learning under personalized settings (Xia Wu, 2023) View paper
  - [14] {Utility-optimized} local differential privacy mechanisms for distribution estimation (Takao Murakami, 2019) View paper
  - [15] AdaPDP: Adaptive Personalized Differential Privacy (Ben Niu, 2021) View paper
  - [20] Local Differential Privacy Preservation via the Novel Encoding Method (Niu Zhang, 2023) View paper
  - Personalized Privacy in Federated Learning (7 papers)
  - [2] Adaptive utility optimization for personalized local differential privacy (Linhai Cheng, 2025) View paper
  - [4] Enhancing Convergence, Privacy and Fairness for Wireless Personalized Federated Learning: Quantization-Assisted Min-Max Fair Scheduling (Xiyu Zhao, 2025) View paper
  - [6] Personalized privacy-preserving federated learning: Optimized trade-off between utility and privacy (Jinhao Zhou, 2022) View paper
  - [17] Fairness-Aware Client Selection and Payment Determination for Differentially Private Federated Learning (Zhao Xin, 2025) View paper
  - [23] Optimizing Federated Learning with Local Differential Privacy: A Game-Theoretic Approach for Privacy-Utility Tradeoff (QingKui Zeng, 2090) View paper
  - [25] Crowdsourcing in federated learning: a scalable approach to collaborative model training (D.V.S. Rashmika, 2025) View paper
  - [33] PLFa-FL: Personalized Local Differential Privacy for Fair Federated Learning (Hongyun Cai, 2024) View paper
  - Individualized Privacy for Centralized Learning (3 papers)
  - [18] Protecting Regression Models With Personalized Local Differential Privacy (Xiaoguang Li, 2022) View paper
  - [36] Individualized PATE: Differentially Private Machine Learning with Individual Privacy Guarantees (Franziska Boenisch, 2022) View paper
  - [39] Heterogeneous Differential Privacy (Mohammad Alaggan, 2015) View paper
- Fairness and Utility Imbalance under Differential Privacy
  - Utility Imbalance and Subgroup Disparity Analysis ★ (5 papers)
  - [0] INO-SGD: Addressing Utility Imbalance under Individualized Differential Privacy (Anon et al., 2026) View paper
  - [5] Neither private nor fair: Impact of data imbalance on utility and fairness in differential privacy (Tom Farrand, 2020) View paper
  - [8] Privacy at a Price: Exploring its Dual Impact on AI Fairness (Yang Mengmeng, 2024) View paper
  - [11] On the Fairness of Privacy Protection: Measuring and Mitigating the Disparity of Group Privacy Risks for Differentially Private Machine Learning (Yang Zhi, 2025) View paper

  - [24] Don't Hash Me Like That: Exposing and Mitigating Hash-Induced Unfairness in Local Differential Privacy (Berkay Kemal Balioglu, 2025) View paper
  - Joint Fairness and Privacy Optimization in Federated Learning (6 papers)
  - [10] FedFDP: Fairness-Aware Federated Learning with Differential Privacy (Ling, 2024) View paper
  - [16] Utility fairness for the differentially private federated-learning-based wireless IoT networks (Sheeraz A. Alvi, 2022) View paper
  - [28] Fairness and Privacy Guarantees in Federated Contextual Bandits (Solanki, 2024) View paper
  - [31] Federated Learning Meets Fairness and Differential Privacy (P. Manisha, 2021) View paper
  - [35] Fair Differentially Private Federated Learning Framework (Varshney, 2023) View paper
  - Fairness-Aware Mechanisms for Centralized DP Models (9 papers)
  - [3] FairDP: Achieving Fairness Certification with Differential Privacy (Khang Tran, 2023) View paper
  - [9] Achieving Hilbert-Schmidt Independence Under RÃ©nyi Differential Privacy for Fair and Private Data Generation (Panagiotou, 2025) View paper
  - [19] Achieving Differential Privacy and Fairness in Logistic Regression (Depeng Xu, 2019) View paper
  - [21] Fairness Meets Privacy: Integrating Differential Privacy and Demographic Parity in Multi-class Classification (Lilian Say, 2025) View paper
  - [27] Privacy for Fairness: Information Obfuscation for Fair Representation Learning with Local Differential Privacy (Xie, 2024) View paper
  - [32] Differentially Private Fair Binary Classifications (Hrad Ghoukasian, 2024) View paper
  - [38] Achieving Differential Privacy and Fairness in Machine Learning (Depeng Xu, 2021) View paper
  - [40] Differential Privacy for Fair Deep Learning Models (Ahmed El Ouadrhiri, 2021) View paper
  - [41] Differentially Private and Fair Classification via Calibrated Functional Mechanism (Jiahao Ding, 2020) View paper
  - Post-Processing and Fairness in DP Outputs (2 papers)
  - [1] Privacy and Fairness Analysis in the Post-Processed Differential Privacy Framework (Ying Zhao, 2025) View paper
  - [34] On the Impact of Output Perturbation on Fairness in Binary Linear Classification (Perrot, 2024) View paper
- Privacy-Utility Trade-offs and Empirical Interactions (2 papers)
  - [30] Privacy-Utility-Bias Trade-offs for Privacy-Preserving Recommender Systems (Shiva Parsarad, 2025) View paper
  - [37] (Local) Differential Privacy has NO Disparate Impact on Fairness (HÃ©ber H. Arcolezi, 2023) View paper
- Domain-Specific and Application-Oriented DP Frameworks
  - Blockchain and Crowdsourcing with DP (3 papers)
  - [7] On Fair Ordering and Differential Privacy (Cohen, 2025) View paper
  - [22] Privacy-First Crowdsourcing: Blockchain and Local Differential Privacy in Crowdsourced Drone Services (Junaid Akram, 2024) View paper
  - [26] Research on a Blockchain Adaptive Differential Privacy Mechanism for Medical Data Protection (Wang Feier, 2025) View paper
  - Domain-Aware DP for Sensitive Applications (1 papers)
  - [29] FedMentor: Domain-Aware Differential Privacy for Heterogeneous Federated LLMs in Mental Health (Dipta, 2025) View paper

### Narrative

Core task: Addressing utility imbalance under individualized differential privacy. The field has evolved around four main branches that reflect distinct but interconnected concerns. The first branch, Individualized and Personalized Privacy Mechanisms, develops frameworks that allow heterogeneous privacy budgets across users, enabling personalized protection levels as seen in works like Utility-Optimized Local Privacy[14] and Personalized Privacy Federated[6]. The second branch, Fairness and Utility Imbalance under Differential Privacy, examines how privacy mechanisms can inadvertently create disparities in model utility across subgroups, with studies such as Neither Private Nor Fair[5] and FairDP[3] highlighting these tensions. The third branch focuses on Privacy-Utility Trade-offs and Empirical Interactions, exploring how different privacy regimes affect learning performance and convergence, while the fourth branch addresses Domain-Specific and Application-Oriented DP Frameworks, tailoring differential privacy to federated learning, crowdsourcing, and other specialized settings.

Recent work has intensified around the interplay between personalized privacy and fairness guarantees, with many studies seeking mechanisms that simultaneously respect individual privacy preferences and ensure equitable utility across demographic groups. INO-SGD[0] sits within the utility imbalance and subgroup disparity analysis cluster, addressing how individualized noise injection can exacerbate performance gaps between subpopulations. Its emphasis on balancing per-user privacy levels with group-level utility contrasts with approaches like Privacy Fairness Post-Processed[1] and Adaptive Utility Optimization[2], which apply post-hoc corrections or adaptive budget allocation to mitigate disparity. Nearby works such as Hash-Induced Unfairness[24] and Privacy at Price[8] further explore how algorithmic choices and economic incentives shape fairness outcomes under personalized privacy, underscoring ongoing debates about whether utility imbalance is an inherent cost of individualization or a design challenge amenable to principled solutions.

## Related Works in Same Category

The following **1 sibling papers** share the same taxonomy leaf node with the original paper:

### 1. Privacy at a Price: Exploring its Dual Impact on AI Fairness

**Authors**: Yang Mengmeng, Ding Ming, Mengmeng Yang, Qu, Youyang, et al. (15 authors total) | **Year/Venue**: 2024 • arXiv.org | **URL**: View paper

#### Abstract

The worldwide adoption of machine learning (ML) and deep learning models, particularly in critical sectors, such as healthcare and finance, presents substantial challenges in maintaining individual privacy and fairness. These two elements are vital to a trustworthy environment for learning systems. While numerous studies have concentrated on protecting individual privacy through differential privacy (DP) mechanisms, emerging research indicates that differential privacy in machine learning models...

#### Relationship Analysis

Both papers belong to the same taxonomy category examining how differential privacy mechanisms affect utility and fairness across demographic subgroups. While the original paper (INO-SGD) focuses on addressing utility imbalance under individualized differential privacy by proposing a novel algorithm that strategically weights gradients to improve performance on more private data, the candidate paper empirically investigates the non-monotonic relationship between privacy levels and fairness disparities, demonstrating that accuracy gaps across subgroups initially grow but then diminish at higher privacy levels. The key difference is that the original paper proposes a solution (INO-SGD algorithm) for utility imbalance under individualized privacy budgets, whereas the candidate paper provides an empirical analysis of how standard differential privacy affects fairness without proposing algorithmic interventions for individualized settings.

## Contributions Analysis

**Overall novelty summary.** The paper proposes INO-SGD, an algorithm that addresses utility imbalance arising when users set heterogeneous privacy requirements under individualized differential privacy (IDP). It resides in the 'Utility Imbalance and Subgroup Disparity Analysis' leaf, which contains five papers examining how DP exacerbates accuracy gaps across subgroups. This leaf sits within the broader 'Fairness and Utility Imbalance under Differential Privacy' branch, indicating a moderately populated research direction focused on understanding and mitigating disparate impacts of privacy mechanisms.

The taxonomy reveals that neighboring leaves address related but distinct concerns: 'Joint Fairness and Privacy Optimization in Federated Learning' (six papers) focuses on FL-specific fairness-privacy trade-offs, while 'Fairness-Aware Mechanisms for Centralized DP Models' (nine papers) develops training-time interventions for centralized settings. The 'Individualized and Personalized Privacy Mechanisms' branch (eleven papers across three leaves) explores heterogeneous privacy budgets but does not explicitly target utility imbalance. INO-SGD bridges these areas by proposing a centralized training algorithm that both satisfies IDP and mitigates the resulting utility gaps.

No literature search was conducted for this analysis, so no candidate papers were examined and no refutation statistics are available. The contribution-level analysis shows zero candidates examined for all three contributions: the INO-SGD algorithm, the IDP-induced utility imbalance analysis, and the INO-SGM generalization. Without empirical search results, we cannot assess whether prior work overlaps with these specific algorithmic or analytical contributions. The taxonomy context suggests the problem space is recognized, but the novelty of the proposed solution remains unverified by this limited analysis.

Given the absence of a literature search, this assessment relies solely on taxonomy structure and sibling paper positioning. The paper appears to occupy a recognized but not overcrowded niche at the intersection of individualized privacy and utility imbalance. A full novelty evaluation would require examining the sibling papers and related leaves to determine whether INO-SGD's strategic down-weighting approach or its IDP-specific design represents a substantive advance over existing disparity mitigation techniques.

This paper presents **3 main contributions**, each analyzed against relevant prior work:

### Contribution 1: INO-SGD algorithm for addressing IDP-induced utility imbalance

**Description**: The authors introduce the Individualized Noisy Ordered SGD (INO-SGD) algorithm that addresses utility imbalance arising from individualized differential privacy requirements. The algorithm strategically assigns importance scores to gradients based on loss ordering, down-weighting less important gradients while preserving IDP guarantees and improving model performance on data from owners with stronger privacy requirements.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

### Contribution 2: Analysis of IDP-induced utility imbalance problem

**Description**: The authors identify and theoretically analyze a critical utility imbalance problem in individualized differential privacy settings, showing that data from owners with stronger privacy requirements may be severely underrepresented in trained models. They demonstrate that this problem differs from standard data imbalance and cannot be solved by existing techniques.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

### Contribution 3: INO-SGM mechanism generalizing INO-SGD

**Description**: The authors develop a generalized individualized differential privacy mechanism called INO-SGM that extends the INO-SGD approach beyond stochastic gradient descent. This mechanism provides a broader framework for applying score-based ordering while maintaining IDP guarantees.

This contribution was assessed against **0 related papers** from the literature. Papers with potential prior art are analyzed in detail with textual evidence; others receive brief assessments.

## Appendix: Text Similarity Detection

No high-similarity text segments were detected across any compared papers.

## References

- [0] INO-SGD: Addressing Utility Imbalance under Individualized Differential Privacy View paper
- [1] Privacy and Fairness Analysis in the Post-Processed Differential Privacy Framework View paper
- [2] Adaptive utility optimization for personalized local differential privacy View paper
- [3] FairDP: Achieving Fairness Certification with Differential Privacy View paper
- [4] Enhancing Convergence, Privacy and Fairness for Wireless Personalized Federated Learning: Quantization-Assisted Min-Max Fair Scheduling View paper
- [5] Neither private nor fair: Impact of data imbalance on utility and fairness in differential privacy View paper
- [6] Personalized privacy-preserving federated learning: Optimized trade-off between utility and privacy View paper
- [7] On Fair Ordering and Differential Privacy View paper
- [8] Privacy at a Price: Exploring its Dual Impact on AI Fairness View paper
- [9] Achieving Hilbert-Schmidt Independence Under RÃ©nyi Differential Privacy for Fair and Private Data Generation View paper
- [10] FedFDP: Fairness-Aware Federated Learning with Differential Privacy View paper
- [11] On the Fairness of Privacy Protection: Measuring and Mitigating the Disparity of Group Privacy Risks for Differentially Private Machine Learning View paper
- [12] Local differential privacy-based federated learning under personalized settings View paper
- [13] FedFDP: Federated Learning with Fairness and Differential Privacy View paper
- [14] {Utility-optimized} local differential privacy mechanisms for distribution estimation View paper
- [15] AdaPDP: Adaptive Personalized Differential Privacy View paper
- [16] Utility fairness for the differentially private federated-learning-based wireless IoT networks View paper
- [17] Fairness-Aware Client Selection and Payment Determination for Differentially Private Federated Learning View paper
- [18] Protecting Regression Models With Personalized Local Differential Privacy View paper
- [19] Achieving Differential Privacy and Fairness in Logistic Regression View paper
- [20] Local Differential Privacy Preservation via the Novel Encoding Method View paper
- [21] Fairness Meets Privacy: Integrating Differential Privacy and Demographic Parity in Multi-class Classification View paper
- [22] Privacy-First Crowdsourcing: Blockchain and Local Differential Privacy in Crowdsourced Drone Services View paper

- [23] Optimizing Federated Learning with Local Differential Privacy: A Game-Theoretic Approach for Privacy-Utility Tradeoff View paper
- [24] Don't Hash Me Like That: Exposing and Mitigating Hash-Induced Unfairness in Local Differential Privacy View paper
- [25] Crowdsourcing in federated learning: a scalable approach to collaborative model training View paper
- [26] Research on a Blockchain Adaptive Differential Privacy Mechanism for Medical Data Protection View paper
- [27] Privacy for Fairness: Information Obfuscation for Fair Representation Learning with Local Differential Privacy View paper
- [28] Fairness and Privacy Guarantees in Federated Contextual Bandits View paper
- [29] FedMentor: Domain-Aware Differential Privacy for Heterogeneous Federated LLMs in Mental Health View paper
- [30] Privacy-Utility-Bias Trade-offs for Privacy-Preserving Recommender Systems View paper
- [31] Federated Learning Meets Fairness and Differential Privacy View paper
- [32] Differentially Private Fair Binary Classifications View paper
- [33] PLFa-FL: Personalized Local Differential Privacy for Fair Federated Learning View paper
- [34] On the Impact of Output Perturbation on Fairness in Binary Linear Classification View paper
- [35] Fair Differentially Private Federated Learning Framework View paper
- [36] Individualized PATE: Differentially Private Machine Learning with Individual Privacy Guarantees View paper
- [37] (Local) Differential Privacy has NO Disparate Impact on Fairness View paper
- [38] Achieving Differential Privacy and Fairness in Machine Learning View paper
- [39] Heterogeneous Differential Privacy View paper
- [40] Differential Privacy for Fair Deep Learning Models View paper
- [41] Differentially Private and Fair Classification via Calibrated Functional Mechanism View paper
- [42] Removing disparate impact on model accuracy in differentially private stochastic gradient descent View paper
- [43] Weighted Loss Methods for Robust Federated Learning under Data Heterogeneity View paper
- [44] Technical Report for the Forgotten-by-Design Project: Targeted Obfuscation for Machine Learning View paper
- [45] Removing disparate impact of differentially private stochastic gradient descent on model accuracy View paper
- [46] Bayesian Pseudo Posterior Mechanism for Differentially Private Machine Learning View paper
- [47] Adaptive Token-Weighted Differential Privacy for LLMs: Not All Tokens Require Equal Protection View paper
- [48] Personalized Differential Privacy for Ridge Regression Under Output Perturbation View paper
- [49] On Optimal Hyperparameters for Differentially Private Deep Transfer Learning View paper
- [50] Wavelet-Domain Privacy SGD (WDP-SGD): FrequencySelective Privacy-Preserving Medical AI. View paper
- [51] Privacy-preserving Case-based Explanations for Federated Learning Models View paper
- [52] Personalized federated learning with gaussian processes View paper
- [53] Personalized DP-SGD using sampling mechanisms View paper
- [54] Privacy-preserving matrix factorization for recommendation systems using Gaussian mechanism and functional mechanism View paper
- [55] Application of the joint clustering algorithm based on Gaussian kernels and differential privacy in lung cancer identification. View paper
- [56] Gradient-based Riemannian Gaussian Differential Privacy View paper
- [57] Per-Attribute Privacy in Large Language Models Using Matrix-Variate Gaussian Mechanism View paper
- [58] TPMDP: Threshold Personalized Multi-party Differential Privacy via Optimal Gaussian Mechanism View paper
- [59] Fed-DPSDG-WGAN: Differentially Private Synthetic Data Generation for Loan Default Prediction via Federated Wasserstein GAN View paper
- [60] DP-MTFL: Differentially Private Multi-Tier Federated Learning for IoT applications View paper
- [61] Personalized trajectory privacy data publishing scheme based on differential privacy View paper
- [62] Conservative or liberal? Personalized differential privacy View paper
- [63] Communication-Efficient and Utility-Enhanced Local Differential Privacy-Based Personalized Federated Compressed Learning View paper
- [64] Generating synthetic personal health data using conditional generative adversarial networks combining with differential privacy View paper
- [65] Personalized Local Differential Privacy for Multi-dimensional Range Queries over Mobile User Data View paper
- [66] Differential Privacy in HyperNetworks for Personalized Federated Learning View paper
- [67] Utility-aware Exponential Mechanism for Personalized Differential Privacy View paper
- [68] Differential privacy and fairness in decisions and learning tasks: A survey View paper
- [69] Convergence-Privacy-Fairness Trade-Off in Personalized Federated Learning View paper