

A Context-Aware Hybrid Local Differential Privacy–Federated Learning Framework for Digital Footprint Protection

Onyia Ogochukwu Sophia^{1*}; Akawuku Mirian Ogheneyovwino²;
Chekwube Georgina Nwankwo³

¹Department of Computer Science, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria

²Department of Adverting, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria

³Department of Computer Science, Chukwuemeka Odumegwu Ojukwu University, Uli, Anambra State, Nigeria

Corresponding Authors: Onyia Ogochukwu Sophia^{1*}

Publication Date: 2026/10/06

Abstract: Digital services increasingly depend on information generated from users' online activities, yet the same traces that make these services useful can expose individuals to re-identification and inference. This study develops and evaluates a context-aware Hybrid Local Differential Privacy–Federated Learning (LDP-FL) approach for protecting digital footprints while retaining useful machine-learning performance. The work distinguishes digital footprints into active, passive and hybrid categories and assigns privacy protection according to feature sensitivity rather than applying a single noise level to every attribute. The proposed pipeline perturbs sensitive features locally with a Laplace mechanism, clips local gradients, and uses Federated Averaging to train a global model without centralising raw client data. A synthetic dataset of 5,000 records and 20 normalised features was partitioned across ten simulated clients, with evaluations conducted over 20 communication rounds on a held-out set of 1,000 records. The hybrid configuration used $\epsilon_{\text{high}} = 0.5$ for high-sensitivity attributes and $\epsilon_{\text{low}} = 10.0$ for lower-sensitivity attributes. At the end of training, Hybrid LDP-FL achieved 79.50% accuracy and a weighted F1-score of 79.16%, compared with 64.20% and 62.52% for uniform Static LDP-FL. Its empirical nearest-neighbour linkage re-identification rate was 21.30%, compared with 0.70% for Static LDP-FL and 100% for the unprotected baseline. A separate privacy-budget sweep identified $\epsilon_{\text{high}} = 0.5$ as the largest tested high-sensitivity budget that remained below the study's 25% re-identification threshold. The results indicate that sensitivity-aware privacy allocation can retain substantially more model utility than uniform perturbation while materially reducing linkage risk.

Keywords: Digital Footprint; Local Differential Privacy; Federated Learning; Privacy-Preserving Machine Learning; Re-Identification; Privacy–Utility Trade-Off; Context-Aware Privacy.

How to Cite: Onyia Ogochukwu Sophia; Akawuku Mirian Ogheneyovwino; Chekwube Georgina Nwankwo (2026) A Context-Aware Hybrid Local Differential Privacy–Federated Learning Framework for Digital Footprint Protection. *International Journal of Innovative Science and Research Technology*, 11(9), 3015-3020. <https://doi.org/10.38124/ijisrt/26sep797>

I. INTRODUCTION

Digital technologies have made personal information unusually persistent. Browsing histories, search activity, social interactions, location traces, device characteristics and application telemetry are generated as people use ordinary online services. Some of these records are deliberately supplied by users, while others are produced quietly in the background. The resulting digital footprint can therefore extend well beyond what a person consciously chooses to disclose (Aswani *et al.*, 2020; Solove, 2021).

The privacy problem becomes harder to define when separate traces are combined. Active footprints may include posts, comments, searches and other information intentionally submitted by a user. Passive footprints include technical and behavioural traces such as device information, cookies, location signals and telemetry. A third layer emerges when systems combine these traces and infer attributes that were never directly disclosed. In this study, these inferred attributes are treated as hybrid digital footprints. Their significance lies in the fact that a collection of individually

modest signals can become much more revealing when analysed together (Zhang *et al.*, 2023; Shen *et al.*, 2022).

Conventional anonymisation is not sufficient for this setting. Literature has repeatedly shown that apparently anonymous records can remain linkable when an attacker has auxiliary information (Sweeney, 2002). Browser and behavioural traces can also be sufficiently distinctive to support re-identification, making privacy protection a problem of inference as well as direct disclosure (Oliveira *et al.*, 2023; Carey *et al.*, 2023). Cryptographic approaches, including homomorphic encryption, offer strong confidentiality properties but may introduce computational costs that are difficult to justify on ordinary edge devices (Gentry, 2009; Xu *et al.*, 2024).

Differential Privacy offers a different route. Rather than attempting to remove every identifying feature, it limits the influence that any one individual's information can have on an observable output by introducing controlled randomness (Dwork & Roth, 2014). Local Differential Privacy moves that protection to the point of data generation, so the raw value is perturbed before it is sent to an external service. Federated Learning addresses a related architectural concern by keeping raw data on participating devices and exchanging model updates instead of the original records (McMahan *et al.*, 2017).

Neither technique is a complete answer on its own. Federated Learning reduces the need to centralise data, but gradients and model updates can still leak information through reconstruction or membership inference attacks (Zhu *et al.*, 2019; Shokri *et al.*, 2017). Uniform Local Differential Privacy, on the other hand, can impose the same degree of distortion on every feature even when the privacy consequences of those features are very different. This can protect low-value information more aggressively than necessary while still failing to reflect the sensitivity of inferred or hybrid attributes.

The study therefore develops a context-aware Hybrid LDP-FL algorithm. The central idea is straightforward: privacy protection should be stronger where the information is more revealing and lighter where additional noise would mainly destroy useful learning signal. Sensitive features are assigned to a smaller privacy budget, while lower-sensitivity features receive a larger budget. The perturbed data are then used for local training, local gradients are clipped, and the resulting updates are aggregated through Federated Averaging. The study evaluates the approach against a standard FL baseline and a uniform Static LDP-FL configuration.

II. RELATED WORKS AND RESEARCH GAP

Research on digital-footprint privacy has developed along several related tracks. Studies of browsing and device fingerprinting have demonstrated that ordinary online behaviour can be surprisingly distinctive. Oliveira *et al.* (2023) examined the uniqueness of browsing traces, while

Papadogiannakis *et al.* (2023) examined browser and device fingerprinting in advertising environments. These studies are useful because they establish the problem empirically, but they are primarily diagnostic: they measure exposure rather than providing an end-to-end protection mechanism.

A second body of work focuses on privacy-preserving learning. Federated Learning keeps data at the edge and has become an important approach for distributed machine learning (McMahan *et al.*, 2017; Kairouz *et al.*, 2021). Yet the privacy benefit is not absolute. Gradient inversion and membership inference research has shown that model updates can contain enough information to expose characteristics of the underlying training data (Zhu *et al.*, 2019; Shokri *et al.*, 2017). Differential Privacy provides a mathematical mechanism for limiting such disclosure, and work combining Differential Privacy with Federated Learning has demonstrated the potential of the combined approach (Wei *et al.*, 2020).

More recent work has begun to personalise privacy settings. Wu *et al.* (2023), for example, considered personalised local differential privacy in federated learning, while Xu *et al.* (2024) examined local privacy mechanisms for edge and IoT environments. These studies point toward a more flexible privacy architecture, but the present study addresses a narrower and more concrete question: whether a simple sensitivity-aware allocation of privacy budgets can recover useful model performance while keeping an empirical linkage attack below a defined risk threshold.

➤ Selected Evidence from the Literature

Wei *et al.* (2020) combine Differential Privacy with Federated Learning; Sahu *et al.* (2023) demonstrate data locality in federated IoT learning; Oliveira *et al.* (2023) and Carey *et al.* (2023) show the persistence of re-identification risk in browsing traces and embeddings; Xu *et al.* (2024) examine local privacy at the edge; Wu *et al.* (2023) explore personalised LDP-FL; and NIST (2024) highlights gradient inversion and membership-inference risks in FL. Across these strands, the remaining issue is how to allocate privacy protection according to the sensitivity of different footprint features.

Taken together, the literature suggests a practical gap between identifying privacy risks and building a transparent protection pipeline that can be tested from data generation through model evaluation. The present work addresses that gap by implementing context-aware LDP at the client, retaining the decentralised FL architecture, adding gradient clipping, and measuring privacy using an explicit nearest-neighbour linkage attack. The contribution is therefore not a new privacy primitive; it is an experimentally evaluated way of combining familiar mechanisms around the sensitivity structure of digital footprints.

III. MATERIALS AND METHODS

The study adopted a Design Science Research approach because the research problem was addressed through the construction and evaluation of a technical artefact. The

artefact is the Hybrid LDP-FL algorithm, implemented as a modular Python pipeline. The evaluation compares three conditions under the same federated-learning architecture: Baseline FL without input perturbation, Static LDP-FL with one uniform privacy budget, and the proposed Hybrid LDP-FL with sensitivity-aware budgets.

➤ Experimental Data and Client Setting

The experiments used a synthetic digital-footprint dataset containing 5,000 records, 20 numerical features and a binary target representing privacy risk. All feature vectors were normalised to the $[0,1]$ interval. The simulated federated environment contained ten clients, with 80% of clients selected in each communication round. The training process ran for 20 global rounds, with three local epochs per client, a batch size of 32 and a learning rate of 0.01. Twenty percent of the records were reserved for testing, producing a held-out evaluation set of 1,000 records.

The feature design follows the thesis classification of digital footprints. High-sensitivity variables include location traces, inferred health interest, inferred political affiliation, financial behaviour, biometric patterns, inferred relationship status, night activity patterns and communication frequency.

Lower-sensitivity variables include search-query counts, application usage duration, page visits, video watch time, likes and comments. The purpose of this classification is not to claim that a feature has a universally fixed sensitivity; rather, it establishes the experimental context in which the proposed routing rule can be tested.

• Experimental Configuration

The system used 5,000 synthetic records, 20 normalised features, two classes, ten simulated clients, 0.80 client participation per round, 20 global rounds, three local epochs, batch size 32, learning rate 0.01, an 80/20 train-test split, random seed 42, and gradient clipping at an L2 norm of 1.0.

➤ Hybrid LDP-FL Architecture

The proposed architecture begins with the client device. Incoming footprint features are first classified according to their experimental sensitivity category. A privacy vector is then constructed so that high-sensitivity features receive $\epsilon_{\text{high}} = 0.5$ and lower-sensitivity features receive $\epsilon_{\text{low}} = 10.0$. Laplace noise is sampled feature-by-feature using the scale $\Delta f/\epsilon$, with global sensitivity $\Delta f = 1.0$. The noisy feature values are subsequently clipped back to $[0,1]$.




Fig 1 High-Level Architecture of the Proposed Context-Aware Hybrid LDP-FL Pipeline.

After perturbation, each client trains a local model using stochastic gradient descent. The resulting gradients are clipped to an L2 norm of 1.0 before the model update is applied. The server does not receive raw footprint records; it receives client model parameters and combines them with sample-size-weighted Federated Averaging. The updated global model is then redistributed for the next round. This preserves the central architectural idea of FL while adding a local privacy layer before learning begins.

➤ Privacy Mechanism and Evaluation

For a feature value x , the local mechanism is represented as $x' = x + Y$, where Y follows a zero-mean Laplace distribution with scale $b = \Delta f/\epsilon$. The hybrid design replaces a single ϵ with a feature-specific vector $\epsilon_{\text{vec}} = (\epsilon_1, \epsilon_2, \dots, \epsilon_d)$. In the implementation, the vector is built by checking each feature against the high- and low-sensitivity lists. This is the main difference between Hybrid LDP-FL and Static LDP-FL.

The clipping operation $x'' = \max(0, \min(x', 1))$ keeps perturbed values within the valid feature domain. Local gradients are then clipped according to $g' = g \cdot \min(1, C/\|g\|_2)$, with $C = 1.0$. Finally, the server applies FedAvg to the client updates. The same aggregation procedure is used in all three experimental conditions, so the comparison is intended to isolate the effect of privacy treatment rather than a change in the aggregation algorithm.

Utility was assessed using global test accuracy and weighted F1-score. Privacy risk was measured empirically through a nearest-neighbour linkage attack. The attacker receives a perturbed representation, computes Euclidean distances to the original test records, and succeeds when the nearest original record is the true counterpart. The resulting Re-identification Rate (Re-ID Rate) is the proportion of test records successfully linked back to their originals. This is an empirical attack measure, not a substitute for the formal privacy guarantee provided by differential privacy.

➤ Privacy-Budget Calibration

The high-sensitivity budget was not selected from the main experiment alone. A separate calibration experiment varied ϵ_{high} over $\{0.10, 0.25, 0.50, 0.75, 1.00, 2.00, 5.00, 10.00\}$, while ϵ_{low} remained fixed at 10.0. The same

synthetic data, client shards, model initialisation and random seed were reused for each candidate. Each candidate was trained for 20 rounds and evaluated using accuracy, weighted F1-score and Re-ID Rate. The study defined 25% as the maximum acceptable empirical linkage-attack rate and selected the candidate with the highest accuracy subject to that constraint.

IV. RESULTS AND DISCUSSION

The three experimental conditions produced clearly different privacy–utility profiles after 20 communication rounds. Baseline FL retained the most learning signal, while uniform LDP produced the strongest suppression of the measured linkage attack at the cost of substantial utility loss. Hybrid LDP-FL occupied the middle of the two extremes, retaining considerably more predictive performance than uniform LDP while reducing re-identification relative to the unprotected baseline.

Table 1. Comparative results after 20 communication rounds. Baseline FL: accuracy 88.00%, weighted F1 87.99%, Re-ID 100.00%; Static LDP-FL: accuracy 64.20%, weighted F1 62.52%, Re-ID 0.70%; Hybrid LDP-FL: accuracy 79.50%, weighted F1 79.16%, Re-ID 21.30%.




Fig 2 Accuracy and Weighted F1-Score after 20 Federated Communication Rounds.

The baseline result provides a useful reference point rather than a target. Its 88.00% accuracy and 87.99% weighted F1-score show what the model can learn when the input features are not perturbed. At the same time, the 100% empirical Re-ID rate indicates that the representations in this experimental setting remain directly linkable by the simulated attacker. This reinforces the distinction between keeping raw records off the server and preventing information about those records from being recoverable from protected representations.

Uniform Static LDP-FL reduced the empirical Re-ID rate to 0.70%, but accuracy fell to 64.20% and weighed F1-

score to 62.52%. In other words, the strong protection achieved by applying $\epsilon = 0.5$ to every feature came with a large reduction in usable signal. The result is consistent with the privacy–utility trade-off described in differential privacy research: stronger perturbation generally makes individual records harder to distinguish, but it can also make the learning task harder.

Hybrid LDP-FL achieved 79.50% accuracy and a weighted F1-score of 79.16%, while its Re-ID rate was 21.30%. Relative to Static LDP-FL, this represents a 15.30 percentage-point gain in accuracy and a 16.64 percentage-point gain in weighted F1-score. Relative to the unprotected

baseline, the hybrid model gives up 8.50 accuracy points while reducing the measured Re-ID rate from 100% to 21.30%. The result is consistent with the design premise: the system does not attempt to treat every feature as equally risky, so less sensitive signals remain more available to the learner.

➤ Privacy-Budget Calibration Results

• Calibration Results

At ϵ_{high} values of 0.10, 0.25, 0.50, 0.75, 1.00, 2.00, 5.00 and 10.00, the corresponding accuracy values were 80.10%, 81.00%, 81.10%, 82.10%, 82.50%, 85.30%, 87.80% and 87.50%; Re-ID rates were 12.00%, 14.40%, 21.10%, 27.00%, 34.00%, 61.90%, 97.00% and 100.00%, respectively.




Fig 3 Accuracy and Empirical Re-Identification Rate Across the ϵ_{High} Calibration Sweep.

The calibration sweep provides a more informative view of the trade-off than a single comparison. Accuracy rises as ϵ_{high} increases from 0.10 to 5.0, while the linkage rate rises much faster. At $\epsilon_{\text{high}} = 0.75$, the Re-ID rate reaches 27.00%, crossing the study's 25% threshold. At $\epsilon_{\text{high}} = 0.50$, accuracy is 81.10% and Re-ID is 21.10%, keeping the measured attack success below the threshold. Although $\epsilon_{\text{high}} = 0.75$ offers a further 1.00 percentage point of accuracy, it does so at the cost of crossing the predefined risk boundary. The calibration rule therefore selects $\epsilon_{\text{high}} = 0.50$.

One detail is worth noting when interpreting these numbers. The main comparison reports a final Hybrid LDP-FL accuracy of 79.50% and Re-ID rate of 21.30%, whereas the calibration sweeps at $\epsilon_{\text{high}} = 0.50$ reports 81.10% and 21.10%. This difference arises because the calibration experiment and the main comparison are separate runs, even though they use the same overall experimental design. The article retains both reported results rather than silently treating them as identical observations.

V. IMPLICATIONS OF THE FINDINGS

The practical value of the proposed design lies in the allocation of privacy rather than in the use of noise alone. A

uniform mechanism has no way to distinguish between a feature that is mainly useful for prediction and a feature that can expose an individual's location, biometric behaviour or inferred health interest. The hybrid mechanism introduces that distinction explicitly. In the experimental setting, this was enough to recover much of the predictive performance lost under uniform $\epsilon = 0.5$ perturbation.

The findings also illustrate why privacy claims should be tied to an attack model. The 21.30% Re-ID rate is evidence that the tested nearest-neighbour attacker was substantially less successful against the hybrid representation than against the unprotected baseline. It does not mean that all possible attacks are prevented. The thesis itself identifies gradient reconstruction, membership inference and more advanced attack strategies as areas for further evaluation. Likewise, the 25% threshold used for calibration is an experimental decision made in this study, not a universal definition of acceptable privacy risk.

The architecture is deliberately lightweight. The implementation uses Python 3.11, PyTorch, Pandas, NumPy, Scikit-learn and Streamlit, and it was designed around standard multi-core consumer hardware rather than specialised infrastructure. That choice is consistent with the intended edge setting, although the present experiments do

not establish performance under real network latency, heterogeneous hardware or commercial-scale deployment. Those questions require a subsequent study with real-world data and a broader systems evaluation.

VI. CONCLUSION

This study developed and evaluated a context-aware Hybrid LDP-FL approach for digital-footprint privacy protection. The approach combines local perturbation, sensitivity-aware privacy budgets, gradient clipping and Federated Averaging. Its purpose was not to remove all information from the learning process, but to place stronger protection around features whose disclosure is more consequential while preserving useful signal elsewhere.

Across the 20-round experiment, Hybrid LDP-FL achieved 79.50% accuracy and a weighted F1-score of 79.16%, compared with 64.20% and 62.52% under uniform Static LDP-FL. Its empirical linkage re-identification rate was 21.30%, compared with 0.70% for Static LDP-FL and 100% for the unprotected baseline. The calibration experiment further identified $\epsilon_{\text{high}} = 0.50$, with $\epsilon_{\text{low}} = 10.0$, as the largest tested high-sensitivity budget that remained below the study's 25% empirical linkage-risk threshold.

These results support the narrower claim that context-aware privacy allocation can provide a more usable privacy-utility balance than applying the same strong perturbation to every feature in the tested synthetic setting. They should not be read as evidence that the algorithm is ready for unrestricted production deployment. The principal limitations are the use of synthetic data, simulated clients and a controlled experimental environment, together with the use of a single nearest-neighbour linkage attack. Future work should test dynamic sensitivity estimation, larger real-world datasets, stronger gradient-leakage attacks and secure multi-party computation as a complement to the local privacy mechanism.

REFERENCES

- [1]. Aswani, R., Kar, A. K., Ilavarasan, P. V., & Dwivedi, Y. K. (2020). Digital footprints: Opportunities for personalization and risks for privacy. *Information Systems Frontiers*, 22(2), 389–407.
- [2]. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, 108, 2938–2948.
- [3]. Carey, C. J., et al. (2023). Measuring re-identification risk in user embeddings. *Proceedings of the ACM Conference on Computer and Communications Security*, 201–215.
- [4]. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407.
- [5]. Gentry, C. (2009). A fully homomorphic encryption scheme. Stanford University PhD thesis.
- [6]. Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- [7]. Kairouz, P., McMahan, H. B., Avenet, B., et al. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
- [8]. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. y. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, 54, 1273–1282.
- [9]. NIST. (2024). Privacy attacks in federated learning: A taxonomy and risk assessment. NIST Special Publication 800-225.
- [10]. Oliveira, D., Oliveira, D. S., Santiago, P. M., & João Paulo. (2023). The uniqueness of web browsing traces. *Proceedings of the Web Conference*, 1450–1460.
- [11]. Papadogiannakis, E. (2023). Measuring web fingerprinting in the ad ecosystem. *Proceedings of the Internet Measurement Conference*.
- [12]. Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77.
- [13]. Sahu, A. K., Gupta, D. N., & Kumar, S. (2023). Federated learning for IoT device identification. *IEEE Internet of Things Journal*, 10(15), 13456–13465.
- [14]. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. *2017 IEEE Symposium on Security and Privacy*, 3–18.
- [15]. Solove, D. J. (2021). *Privacy: A very short introduction*. Oxford University Press.
- [16]. Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), 557–570.
- [17]. Wei, K., Li, J., Ding, M., et al. (2020). Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15, 3454–3469.
- [18]. Wu, Y., Deng, Y., & Yu, H. (2023). PLDP-FL: Federated learning with personalized local differential privacy. *Information Sciences*, 630, 1–18.
- [19]. Xu, Z., Wang, L., & Zhang, X. (2024). A review of privacy preserving techniques in edge computing. *arXiv preprint arXiv:2401.11305*.
- [20]. Zhang, J., Liu, W., & Li, Y. (2023). Ambient intelligence and hybrid digital footprints. *Pervasive and Mobile Computing*, 83, 101612.
- [21]. Zhu, L., Liu, Z., & Han, S. (2019). Deep leakage from gradients. *Advances in Neural Information Processing Systems*, 32.
- [22]. Zuboff, S. (2020). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.