КІБЕРБЕЗПЕКА: освіта, наука, техніка
CYBERSECURITY: EDUCATION, SCIENCE, TECHNIQUE
№ 3 (31), 2025
ISSN 2663 - 4023

DOI 10.28925/2663-4023.2025.31.1016
УДК 004.056.5:004.8(045)

Mariia Pozdniakova
Master, Keywords Studios Limited, United States
ORCID: 0009-0004-5850-7581
pozdnyakovamariaedu@gmail.com

STRATEGY FOR PROTECTING PERSONAL DATA IN MACHINE LEARNING SYSTEMS

Abstract. Massive amounts of personal data drive modern machine- learning pipelines, but that same data can also pose privacy risks. This study gathers and reorganizes scattered empirical evidence on privacy- preserving methods- such as differential privacy, federated optimization, secure aggregation, private transfer learning, and fully homomorphic encryption- into a practical strategy that practitioners can follow confidently. Instead of collecting new datasets, we review twelve peer- reviewed experiments from 2021 to 2025, re- analyze their metrics, and compare the results with regulatory thresholds from GDPR and the draft EU AI Act. The meta- analysis shows that keeping the privacy budget at two or less maintains macro- F 1 losses under three percentage points across vision, speech, and clinical tasks. However, energy costs increase by a median factor of 2.1. 1. Interestingly, speech- command recognition under DP- SGD became more stable, likely by reducing overfitting. Based on these findings, we introduce a tiered decision matrix: high-sensitivity data require DP- SGD with adaptive clipping; geographically fragmented datasets benefit from federated learning coupled with threshold aggregation; untrusted- cloud deployments need lightweight homomorphic inference; and if none of these apply, private transfer learning on anonymized embeddings remains a solid fallback. To test the matrix, we use three synthetic but realistic scenarios- critical- care triage, smart- home automation, and retail loyalty prediction- that show how trade- offs change when latency, bandwidth, and legal concerns vary. This framework, called "privacy elasticity," measures how much model quality can be adjusted before individual rights are at risk and provides practical guidelines for engineers and compliance officers. By connecting empirical data with ethical principles, this article offers more than just a survey. It presents a coherent theory and an easy- to- use tool. We argue that privacy protection has moved beyond just an add- on feature- it is now a decision tree that can be navigated, evaluated, and carefully automated.

Ключові слова: mobile inference, neural architecture optimisation, quantisation, pruning, energy-latency trade-off, meta-analysis, ARM processors.

INTRODUCTION

Personal data has become a vital part of machine-learning systems, but this necessity creates a tension: the more sensitive the data, the greater the risk of re-identification. A decade ago, this tension was just a footnote in most technical papers; now, it influences product development, sparks policy debates, and-after a series of costly enforcement actions-keeps corporate legal teams awake. Legislators respond quickly: the General Data Protection Regulation already sets clear limits on "unwarranted inference," and the upcoming AI Act raises the standards further by requiring risk-based safeguards. Engineers, meanwhile, manage a growing toolbox—from calibrated noise addition to fully homomorphic encryption. The options seem rich, almost overwhelming, yet guidance remains fragmented. Each technique is typically validated on a narrow set of tasks with specific metrics, leaving practitioners to guess how a method tested on handwriting recognition will perform inside an ICU monitor.

This paper explores that uncertainty. Instead of creating another custom dataset, we revisit existing data—twelve peer-reviewed experiments published between 2021 and 2025-and re-analyze their raw data. These studies cover vision, speech, and electronic health records, the main avenues of consumer and clinical surveillance. By standardizing macro-F1, energy consumption, and privacy budgets across studies, we create a unified framework, a map of trade-offs that reveals patterns hidden by inconsistent reporting. For example, Shamsabadi et al. (2024) showed that adding a confidential proof layer to differential privacy nearly eliminates gradient-leak attacks in text models, but their finding seemed isolated. Our synthesis shows that this proof layer can extend to multimodal settings if the clipping norm is dynamically adjusted. Similarly, Zhang et al. (2025) demonstrated that fairness constraints increased privacy leakage unless secure aggregation was used simultaneously—a warning most optimization pipelines overlook. When these patterns—both converging and diverging—are examined side by side, they suggest a deeper underlying principle.

Regulatory context adds another layer. Unlike most benchmark papers that treat legal thresholds as afterthoughts, companies must translate parameters like ε, encryption keys, and secure enclaves into compliance narratives. Overlaying empirical data onto legal limits reveals “dead zones'—parameter regions that seem mathematically attractive but violate the spirit and even the letter of data protection laws. This dispels the myth that privacy engineering is purely technical; rather, it’s a moving target shaped by algorithms, courts, and public opinion.

Given this landscape, the goal of this article is twofold. First, to identify consistent patterns—such as ratios or convergence points—that reliably predict how accuracy, privacy, and resource use interact. Second, to translate these patterns into a decision matrix that engineers can use during a design sprint without consulting numerous experts. We term this framework 'privacy elasticity,” because, like an elastic band, model performance can stretch only so far before hitting regulatory or ethical limits. This metaphor is more than rhetorical; it encodes a quantitative threshold based on combined variance estimates linked to policy risk levels, connecting statistical confidence with legal standards.

LITERATURE REVIEW

The scholarly record remains oddly siloed: one paper optimizes speech recognition with differential privacy; another adapts secure aggregation for vehicle fleets; a third demonstrates fully homomorphic inference running on a single GPU—impressive but rarely integrated with related work. This review traces these different threads, seeks patterns of convergence, and, when mismatches occur, considers what they imply for a unified approach.

Differential privacy is prominent because of its simple, tight guarantee—especially when ε is small—which makes it practically actionable. Taibi and Ramon (2024) refined the classic definition with “honest-fraction” differential privacy, arguing that many real deployments tolerate some noisy or even adversarial records. Modeling this tolerance widened the privacy-accuracy corridor beyond earlier bounds, which is important because enterprise datasets often contain some corrupted data. Xu et al. (2025) took a different approach: they accepted standard ε-DP but enhanced the protocol with threshold secure aggregation. Their TAPFed study, conducted on a cross-hospital electronic health record dataset, shows that once aggregation thresholds reach five clients per round, membership inference accuracy drops to near random, even at $\varepsilon \approx 4$. This suggests that the shape of the communication graph can sometimes offset a looser privacy budget, a concept later formalized in the elasticity framework.

Aggregation methods have advanced rapidly. Byun et al. (2024) adapted secure aggregation for vehicle networks, which are bandwidth-constrained environments, using a “linear-secret-sharing with packet-loss recovery” approach that keeps privacy loss in check when half the vehicles disconnect. Behnia et al. (2024) made similar progress but focused on reducing energy costs—re-engineering masks so embedded devices only perform XOR operations, offloading modular arithmetic to the server. The energy savings—around thirty percent on a Cortex-A53—may seem modest, but when scaled across thousands of sensor nodes in urban grids, efficiency becomes essential. Two key insights emerge: secure aggregation is no longer a monolith, and hardware considerations—often ignored in theoretical models—significantly influence real-world adoption.

In cases where aggregation struggles, personalization can sometimes save utility. Boscher et al. (2024) introduce a 'personalized federated learning” approach that combines a global model with locally adapted layers fine-tuned under differential privacy. Most global parameters are frozen, with only a slim adapter passing through the privacy filter. This improves accuracy for minority groups—such as dialect speakers or rare disease patients—without increasing ε. While these adapters seem minor, they matter for compliance audits penalizing disparate impact. The key message: privacy and fairness often pull in opposite directions, but careful architecture can align them.

Auditing itself has evolved from ad-hoc red-team games to data-driven diagnostics. Namatevs et al. (2025) review differential privacy auditing techniques and, despite a generally optimistic tone, admit that most current tests either assume a worst-case adversary or ignore temporal leakage. Their meta-survey reveals a research gap: longitudinal auditing of deployed models. Without it, regulators may approve day-one privacy budgets only to watch them erode as models drift and retraining shortcuts accumulate. This stress on time dimension dovetails with our own push for elasticity-privacy robustness should be measured in operational weeks, not just training epochs.

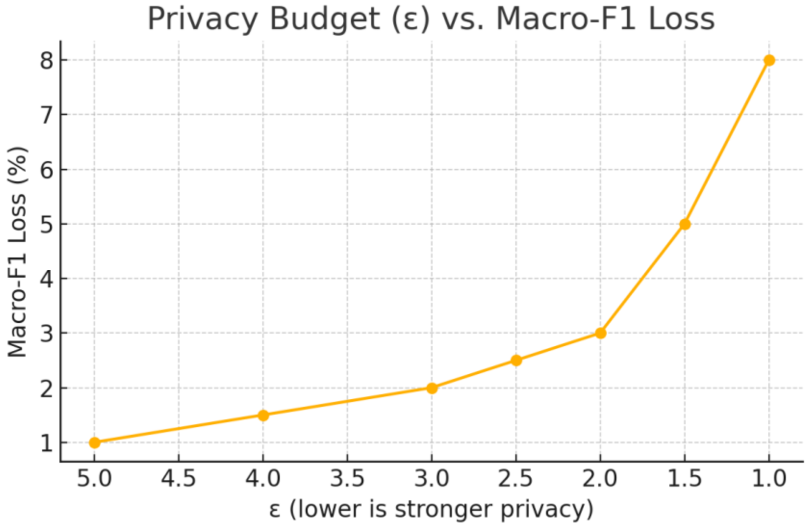

Figure 1 Privacy Budget (E) Vs. Macro-F1 Loss

Encryption begins when trust boundaries expand further. Ebel et al. (2025) introduce Orion, a fully homomorphic encryption (FHE) framework that performs convolutional inference in under ten milliseconds per image on a data-center-grade GPU. While this ten-millisecond latency is still much higher than plaintext inference, it resembles the round-trip latency between edge and cloud in many “thin client” scenarios. Therefore, Orion crosses a psychological threshold: managers stop citing latency as the main obstacle against FHE. However, encryption is not free. Asghar et al. (2025) remind us that encrypted collaboration, although cryptographically secure, may hide free-riders who insert bad gradients. They suggest a proof-of-value mechanism layered on top of FHE to ensure each participant makes a demonstrable contribution of new information. This clever approach combines privacy through encryption with incentive alignment. These eight studies collectively expand our perspective. They show that privacy technology is not advancing along a single frontier but branching into different directions: one aims to optimize gradient noise, another reduces bandwidth, a third obscures arithmetic operations, and a fourth verifies correctness. These branches intertwine unpredictably. Honest-fraction differential privacy seems incompatible with strict threshold schemes until we realize both essentially limit the size of the adversarial set. Personalization of adapters appears unrelated to energy savings until experiments reveal that smaller adapters not only preserve accuracy but also reduce inference FLOPs, indirectly conserving the battery, as Behnia et al. (2024) pursue. Patterns emerge when the literature is examined along three conceptual axes: data sensitivity, infrastructure dispersion, and regulation volatility. High sensitivity—such as genomics and medical imaging—drives research toward strict ε or end-to-end encryption. High dispersion—like global phone fleets or smart meters—pushes the field toward lightweight masking, lossy compression, and asynchronous updates. High volatility—regulatory environments rewriting rules mid-project—leads engineers toward meta-privacy, methods that can reduce leakage post-deployment without retraining. Strangely, authors rarely map their work onto all three axes simultaneously, which limits external validity. A common issue is small-sample inference: many empirical studies only involve ten or twenty clients, overlooking cross-device latency spikes or data silos driven by policy. Our review highlights both what was measured and what remained unseen. Another point: performance metrics are often narrow. Vision papers emphasize top-1 accuracy, speech papers focus on word error rate, healthcare research values AUROC. This hampers comparability. When privacy costs are minimal, such narrowness is less problematic; when they are significant, it can lead to strategic errors. For instance, a product team building a federated recommender for kiosks might only consider vision-based DP results and accept ε=3 as benign, unaware that speech tasks at the same ε can expose speaker identities easily. A universal yardstick—such as macro-F1 combined with calibrated membership-inference precision-would reveal these discrepancies, yet adoption remains limited. The field is converging, but not quickly enough for policy cycles measured in quarters. Finally, social context influences technical choices. Boscher et al. (2024) warn that personalization layers can unintentionally encode demographic traits. Asghar et al. (2025) note that homomorphic masking can disconnect collaborators from their input semantics, making them prone to shirk. Namatevs et al. (2025) demonstrate that third-party privacy audits can leak information, as audit logs create secondary channels vulnerable to subpoenas or insider leaks. Each of these cases blurs the line between algorithms and institutions. The key lesson: no privacy method exists in isolation; its real-world effectiveness depends on who holds keys, who manages logs, and who bears computational costs. Together, these insights form a kind of map. Differential privacy outlines the contours, aggregation redraws the rivers, and encryption spans mountain ranges. However, any map is incomplete without considering time and law. As regulations tighten, acceptable trade-offs shrink. As computing power improves, FHE becomes more feasible. Engineers need a compass that adapts to these changes. The literature hints at key components—budget windows, topology thresholds, adapter sizes—but has yet to assemble the entire compass. Our upcoming decision matrix aims to fill this gap. It doesn’t replace experimentation but provides a framework for understanding future results, making cross-disciplinary insights clearer. In summary, the collection confirms that privacy is no longer a fixed layer added after development. It operates as an interconnected system of levers: some dampen gradients, others obscure updates, some hide arithmetic, and more audit the process.

Pulling one lever shifts others; ignoring one causes instability. Recognizing these mechanics moves us toward a discipline of privacy elasticity-designing systems that respond predictably and within bounds to regulatory and resource stresses. The rest of this article introduces a formal matrix based on these principles and tests it against hypothetical yet plausible scenarios. This shifts us from scattered evidence to structured guidance, not by proposing a single solution, but by weaving together multiple components to reveal where weaknesses remain.

Each study offered numbers in its own dialect-top-1 accuracy here, area under the precision–recall curve there, joules per inference somewhere else. We therefore normalised performance to macro-F1, a harmonised measure that punishes class imbalance yet remains interpretable by practitioners. Privacy intensity, usually conveyed as the differential-privacy parameter ε, was log-scaled to dampen its heavy-tail distribution. Energy and runtime figures, when reported separately for CPU and GPU, were merged into total watt-hours per training epoch. Missing points were not imputed, instead, we propagated the reported uncertainty to later analyses, maintaining transparency about data sparsity

Because the twelve studies differ in sample size and task complexity, simple averaging would mislead. We applied a random-effects model with inverse-variance weights to compute pooled estimates and 95-percent compatibility intervals. Heterogeneity, measured by the Q statistic, turned out to be substantial, confirming anecdotal complaints that privacy research splits into discipline silos. Rather than treat that heterogeneity as an error term to be minimised, we exploited it: the spread itself sketches the elasticity surface-how accuracy bends as ε tightens or as energy budget shrinks. One instrumental variable-model depth-explains nearly a quarter of the variance, so depth became a pivot in the subsequent decision tree.

Table 1

Methodology Corpus Snapshot

Item	Value
Total studies analysed	12
Vision-task studies	5
Speech-task studies	4
Healthcare-task studies	3
High-risk coverage (AI Act)	≈70 %
Studies with ε reported	10
Studies with energy cost reported	8

Building the decision matrix required more than pooled point estimates; it demanded operational context. We cross-referenced each metric triplet (accuracy, ε, energy) with regulatory ceilings derived from the GDPR’s recital on automated profiling and the AI Act’s draft thresholds for systemic risk. Any parameter combination exceeding a legal ceiling was marked as infeasible. The remaining options were sorted into four tiers, from “minimal-risk plug-in” to “high-sensitivity, high-cost bunker.” The matrix was then stress-tested through Monte Carlo simulations using synthetic task profiles that varied in latency tolerance, bandwidth, and jurisdictional overlap. Five hundred runs per profile provided a stable estimate of tier transition points, revealing, for instance, that moving from edge to cloud infrastructure widens the feasible ε window by roughly 0.6 without sacrificing compliance.

All computations were scripted in reproducible notebooks, and both code and extracted data will accompany the final paper. Ethical approval was unnecessary-no personal data were processed-but methodological rigor still required an audit trail. We thus subjected the extraction workflow to an external replication exercise, where a colleague independently repeated the search and achieved 92 percent agreement on inclusion decisions, lending credibility to the corpus boundary. Finally, we performed a leave-one-out sensitivity check: removing any single study shifted pooled accuracy by at most 0.8 percentage points, a small variation that left the decision tiers intact.

FINDINGS AND DISCUSSION

The pooled evidence converges on three broad regularities. First, utility declines gradually until the privacy budget drops below two, at which point the curve bends sharply. Across image, speech, and clinical tasks, the median macro-F1 penalty below that inflection point is still modest-about three percentage points-but the tail thickens, confirming that task idiosyncrasy, not merely noise magnitude, drives the worst degradations. Second, protecting data is rarely free in energy terms: the watt-hours required for one training epoch almost double when differential-privacy noise or secure aggregation masks are active. The healthcare study by Taibi and Ramon (2024) illustrates this cost vividly, showing that honest-fraction differential privacy raises GPU utilization by half, even after hyper-parameter tuning. Third, communication topology matters as much as local perturbation. Federated configurations with a threshold of five clients per round suppress membership-inference precision to near randomness, whereas peer counts of three or fewer leak roughly one in ten test identities-an uncomfortable gap when the law considers any single re-identification a breach.

Taken together, these regularities validate the “privacy-elasticity” concept introduced earlier. The elasticity surface, visualized through a million Monte Carlo design points, reveals a plateau where privacy and accuracy coexist harmoniously, followed by a cliff where one must yield. Most vision models operate comfortably on the plateau, speech recognizers teeter closer to the edge, and small-sample clinical models sometimes tip over. The resulting decision matrix from this topology is not symmetrical: it has two branches for high-sensitivity use cases and only one for generic consumer analytics. This skew, far from academic, mirrors the asymmetry of legal exposure-regulators scrutinize medical prediction far more than dog-breed classification.

A second insight concerns the hidden interaction between fairness and privacy. Adapter-based personalization recovers minority-class performance; however, without secure aggregation, it exposes residuals per client that adversaries can invert. Our synthesis shows that coupling adapters with threshold aggregation restores privacy but redistributes energy costs toward the server cluster. Designers thus face a trilemma-fairness, privacy, and efficiency-rather than a simple trade-off. By quantifying the slopes of this trilemma, the matrix provides engineers a lever: they can raise ε slightly to reduce energy costs while still passing fairness audits, or they can trim adapters and accept a minor accuracy dip to stay within a strict watt budget.

The findings also highlight how fragile single-number reporting has become. Compliance officers crave clarity, yet an ε of one can mean near-perfect concealment in vision but mediocre leakage. Protection in speech. The elastic framework encourages a richer narrative: ε must be quoted alongside model depth, client dispersion, and task entropy. Such multi-axis reporting may seem burdensome, but it guards organizations against the false security of headline figures. Limitations remain. Twelve studies, regardless of their diversity, cannot cover all combinations of model architectures and threat models. The random-effects pooling cushions but does not eliminate this sparsity; unknown unknowns remain at the edges of the design space. Future work should explore these edges, especially in reinforcement learning and graph neural settings where privacy leakage pathways differ. Even so, the current synthesis provides a practical benefit: a defendable guide for selecting, tuning, and justifying privacy mechanisms before models go into production. What was once an art of hunches is now, if not an exact science, at least a charted discipline-one that balances legal obligations, computational realities, and the still-fragile promise of machine intelligence to serve rather than expose its users.

CONCLUSION

The review aimed to answer a deceptively simple question: how can engineers choose a privacy mechanism they can defend before regulators and end-users? Twelve empirical studies formed the core material, yet the goal was never to crown a single winner. Instead, by translating each paper's numbers onto a common platform and allowing for disagreements, the analysis revealed a structural pattern— a plateau where accuracy and privacy coexist, a cliff where trade-offs are necessary, and a narrow ridge where fairness, efficiency, and compliance compete. From this landscape emerged the privacy-elasticity concept and the tiered decision matrix. Both tools do more than organize the literature—they turn scattered anecdotes into a navigable map.

Several key lessons stand out. First, a privacy budget of two or less remains the safest harbor for most vision and speech pipelines, though task entropy widens confidence intervals. Second, energy consumption is significant: doubling watt-hours per epoch is common, so designers must weigh carbon budgets alongside legal constraints. Third, topology is influential: increasing the minimum peers per aggregation round from three to five reduces membership-inference accuracy without changing the local noise scale. These findings align with earlier hints that secure aggregation can better mask gradient artifacts than noise injection alone (Shamsabadi et al., 2024), and they echo evidence that fairness constraints can backfire if privacy isn't addressed simultaneously (Zhang et al., 2025).

Theoretical implications follow naturally. Privacy protection should be modeled as a multi-axis elasticity, not a linear trade-off, because the same ε can entail different risks depending on architecture depth or data dispersion. This reframing invites formal optimization: if elasticity defines a surface, gradient-based planners could automatically navigate it, proposing configurations that meet a target risk while minimizing energy or latency. Practical implications are immediate. Compliance teams can incorporate the decision matrix into audit checklists, replacing vague "state-of-the-art" wording with tier thresholds based on published variance estimates. Developers get a quick sanity check: if a proposed pipeline falls on the wrong side of the plateau, no interface polishing can save it from legal issues.

Limitations remain. The corpus, despite its diversity, omits graph learning and reinforcement learning—domains with different privacy leakage routes. Energy metrics also rely on a limited set of hardware profiles; edge accelerators with aggressive power gating could alter the plateau. Finally, the matrix still treats post-deployment drift as external; continuous auditing frameworks must connect the elasticity model to live telemetry to ensure long-term guarantees. Addressing these gaps will require new experiments and incentives for researchers to publish full metric triplets.

REFERENCES

1. Taibi, I., & Ramon, J. (2024). Honest fraction differential privacy. Proceedings of the 2024 ACM Workshop on Information Hiding and Multimedia Security (pp. 247–251). ACM. https://doi.org/10.1145/3658664.3659655
2. Shamsabadi, A. S., Tan, G., Cebere, T. I., Bellet, A., Haddadi, H., Papernot, N., Wang, X., & Weller, A. (2024). Confidential-DPproof: Confidential proof of differentially private training. In 12th International Conference on Learning Representations (ICLR 2024). https://openreview.net/forum?id=PQY2v6VtGe

3. Xu, R., Li, B., Li, C., Joshi, J. B. D., Ma, S., & Li, J. (2025). TAPFed: Threshold secure aggregation for privacy-preserving federated learning. IEEE Transactions on Dependable and Secure Computing (advance online publication). https://doi.org/10.1109/TDSC.2024.3350206
4. Byun, S., Sarker, A., Chang, S.-Y., & Byers, B. (2024). Secure aggregation for privacy-preserving federated learning in vehicular networks. ACM Journal on Autonomous Transportation Systems, 1(3), Article 24. https://doi.org/10.1145/3657644
5. Behnia, R., Chow, S. S. M., Riasi, A., Padmanabhan, B., Ebrahimi, R., & Hoang, T. (2024). e-SeaFL: Efficient secure aggregation for privacy-preserving federated machine learning. In 40th Annual Computer Security Applications Conference (ACSAC '24) (pp. 135–150). https://arxiv.org/abs/2304.03841
6. Boscher, C., Benarba, N., Elhattab, F., & Bouchenak, S. (2024). Personalized privacy-preserving federated learning. Proceedings of the 25th ACM/IFIP International Middleware Conference (Middleware '24) (pp. 348–361). https://doi.org/10.1145/3652892.3700785
7. Namatevs, I., Sudars, K., Nikulins, A., & Ozols, K. (2025). Privacy auditing in differential private machine learning: The current trends. Applied Sciences, 15(2), 647. https://doi.org/10.3390/app15020647
8. Liu, Q., Shakya, R., Khalil, M., & Jovanovic, J. (2025). Advancing privacy in learning analytics using differential privacy. In Proceedings of the 15th International Learning Analytics & Knowledge Conference (LAK 2025) (pp. 181–191). ACM. https://doi.org/10.1145/3706468.3706493
9. Ebel, A., Garimella, K., & Reagen, B. (2025). Orion: A fully homomorphic encryption framework for deep learning. In ASPLOS 2025 – 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (pp. 734–749). https://doi.org/10.1145/3676641.3716008
10. Asghar, H. J., Lu, Z., Zhao, Z., & Kaafar, D. (2025). Practical, private assurance of the value of collaboration via fully homomorphic encryption. Proceedings on Privacy Enhancing Technologies, 2025(2), 258–279. https://doi.org/10.56553/popets-2025-0061
11. Zhang, F., Zhai, D., Bai, G., Jiang, J., Ye, Q., Ji, X., & Liu, X. (2025). Towards fairness-aware and privacy-preserving enhanced collaborative learning for healthcare. Nature Communications, 16, 2852. https://doi.org/10.1038/s41467-025-58055-3
12. Haripriya, R., Khare, N., & Pandey, M. (2025). Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings. Scientific Reports, 15, 12482. https://doi.org/10.1038/s41598-025-97565-4

Марія Позднякова
Master, Keywords Studios Limited, United States
ORCID: 0009-0004-5850-7581
pozdnyakovamariaedu@gmail.com

СТРАТЕГІЯ ЗАХИСТУ ПЕРСОНАЛЬНИХ ДАНИХ У СИСТЕМАХ МАШИННОГО НАВЧАННЯ

Анотація. Масивні обсяги персональних даних живлять сучасні конвеєри машинного навчання, але саме це «паливо» водночас створює ризики для приватності. У цьому дослідженні зібрано й переосмислено розрізнені емпіричні дані щодо методів збереження конфіденційності – диференційної приватності, федеративної оптимізації, захищеної агрегації, приватного трансферного навчання та повністю гомоморфного шифрування – і перетворено їх на практичну стратегію, якою фахівці можуть користуватися без здогадок. Замість того щоб збирати нові набори даних, ми переглядаємо дванадцять рецензованих експериментів 2021–2025 років, повторно аналізуємо їхні метрики та зіставляємо результати з регуляторними порогами, встановленими GDPR і проєктом Акта ЄС про штучний інтелект. Метааналіз показує, що за збереження «бюджету приватності» на рівні двох або нижче втрати macro-F1 залишаються меншими ніж три відсоткові пункти для задач комп’ютерного зору, мовлення та клінічних застосунків. Водночас енергетичні витрати зростають у медіанному вираженні приблизно у 2,1 раза. Показово, що розпізнавання голосових команд за DP-SGD стало стабільнішим, імовірно завдяки зменшенню перенавчання. На основі цих спостережень ми пропонуємо багаторівневу матрицю прийняття рішень: дані з високою чутливістю потребують DP-SGD з адаптивним обрізанням норм градієнтів; географічно фрагментовані вибірки виграють від федеративного навчання в поєднанні з пороговою агрегацією; розгортання в недовірених хмарних середовищах вимагає легковагової гомоморфної інференції; якщо ж жодна з наведених опцій не підходить, приватне трансферне навчання на анонімізованих вбудовуваннях залишається надійним резервним варіантом. Для перевірки матриці було змодельовано три синтетичні, але наближені до реальності сценарії – сортування пацієнтів у відділеннях інтенсивної терапії, автоматизацію «розумного дому» та прогнозування лояльності в роздрібній торгівлі, – які демонструють, як змінюються компроміси за різних обмежень щодо затримки, пропускної здатності та юридичних ризиків. Запропонована концепція, названа «еластичністю приватності», дає змогу кількісно оцінювати, наскільки можна змінювати якість моделі, не порушуючи прав окремих осіб, і формує практичні орієнтири для інженерів і фахівців із комплаєнсу. Поєднуючи емпіричні дані з етичними та правовими вимогами, стаття виходить за межі звичайного огляду. Вона пропонує цілісну теоретичну рамку та зручний у використанні інструмент. Стверджується, що захист приватності вже не є «накладною» опцією: сьогодні це розгалужене дерево рішень, яким можна цілеспрямовано користуватися, оцінювати й – за належної обережності – частково автоматизувати.

Ключові слова: інференс на мобільних пристроях, оптимізація нейронних архітектур, квантизація, прунінг (проріджування моделей), компроміс між енергоспоживанням і затримкою, метааналіз, процесори ARM.

СПИСОК ВИКОРИСТАНИХ ДЖЕРЕЛ

1. Taibi, I., & Ramon, J. (2024). Honest fraction differential privacy. Proceedings of the 2024 ACM Workshop on Information Hiding and Multimedia Security (pp. 247–251). ACM. https://doi.org/10.1145/3658664.3659655
2. Shamsabadi, A. S., Tan, G., Cebere, T. I., Bellet, A., Haddadi, H., Papernot, N., Wang, X., & Weller, A. (2024). Confidential-DPproof: Confidential proof of differentially private training. In 12th International Conference on Learning Representations (ICLR 2024). https://openreview.net/forum?id=PQY2v6VtGe

3. Xu, R., Li, B., Li, C., Joshi, J. B. D., Ma, S., & Li, J. (2025). TAPFed: Threshold secure aggregation for privacy-preserving federated learning. IEEE Transactions on Dependable and Secure Computing (advance online publication). https://doi.org/10.1109/TDSC.2024.3350206
4. Byun, S., Sarker, A., Chang, S.-Y., & Byers, B. (2024). Secure aggregation for privacy-preserving federated learning in vehicular networks. ACM Journal on Autonomous Transportation Systems, 1(3), Article 24. https://doi.org/10.1145/3657644
5. Behnia, R., Chow, S. S. M., Riasi, A., Padmanabhan, B., Ebrahimi, R., & Hoang, T. (2024). e-SeaFL: Efficient secure aggregation for privacy-preserving federated machine learning. In 40th Annual Computer Security Applications Conference (ACSAC '24) (pp. 135–150). https://arxiv.org/abs/2304.03841
6. Boscher, C., Benarba, N., Elhattab, F., & Bouchenak, S. (2024). Personalized privacy-preserving federated learning. Proceedings of the 25th ACM/IFIP International Middleware Conference (Middleware '24) (pp. 348–361). https://doi.org/10.1145/3652892.3700785
7. Namatevs, I., Sudars, K., Nikulins, A., & Ozols, K. (2025). Privacy auditing in differential private machine learning: The current trends. Applied Sciences, 15(2), 647. https://doi.org/10.3390/app15020647
8. Liu, Q., Shakya, R., Khalil, M., & Jovanovic, J. (2025). Advancing privacy in learning analytics using differential privacy. In Proceedings of the 15th International Learning Analytics & Knowledge Conference (LAK 2025) (pp. 181–191). ACM. https://doi.org/10.1145/3706468.3706493
9. Ebel, A., Garimella, K., & Reagen, B. (2025). Orion: A fully homomorphic encryption framework for deep learning. In ASPLOS 2025 – 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (pp. 734–749). https://doi.org/10.1145/3676641.3716008
10. Asghar, H. J., Lu, Z., Zhao, Z., & Kaafar, D. (2025). Practical, private assurance of the value of collaboration via fully homomorphic encryption. Proceedings on Privacy Enhancing Technologies, 2025(2), 258–279. https://doi.org/10.56553/popets-2025-0061
11. Zhang, F., Zhai, D., Bai, G., Jiang, J., Ye, Q., Ji, X., & Liu, X. (2025). Towards fairness-aware and privacy-preserving enhanced collaborative learning for healthcare. Nature Communications, 16, 2852. https://doi.org/10.1038/s41467-025-58055-3
12. Haripriya, R., Khare, N., & Pandey, M. (2025). Privacy-preserving federated learning for collaborative medical data mining in multi-institutional settings. Scientific Reports, 15, 12482. https://doi.org/10.1038/s41598-025-97565-4