Balancing Privacy and Performance: Machine Learning Security in Distributed Edge Networks

Arian Li

Thomas S. Wootton High School, Rockville, Maryland, USA

arianwhli@gmail.com

Abstract. Edge computing has transformed how machine learning models are deployed across distributed networks. As computing power gets closer to where data is generated, protecting privacy is becoming more and more important. This review looks at how edge computing and privacy-preserving machine learning come together, especially the trade-offs between privacy, performance, and usability. Past studies have been organized into three main areas: algorithmic approaches like federated learning and differential privacy, cryptographic methods such as secure multi-party computation and homomorphic encryption, and hardware-based techniques that rely on trusted execution environments. This analysis reveals a persistent deployment gap between mathematically elegant privacy guarantees and the practical constraints of resource-limited edge devices, where battery life, latency, and computational overhead force uncomfortable compromises. Key challenges have been identified in multi-objective optimization and evaluate existing balancing mechanisms and propose a conceptual framework that interprets privacy-performance trade-offs through multi-dimensional optimization, clarifying when certain techniques work better depending on deployment scenarios. This review brings together what has been achieved so far and highlights where the field is heading - toward building scalable, regulation-compliant, and efficient privacy-preserving AI systems that can work smoothly in distributed edge environments.

Keywords: Edge computing, Privacy-preserving Machine Learning, Distributed Systems, Federated Learning, Performance Optimization.

1. Introduction

Artificial Intelligence (AI) and machine learning workloads have exploded in the past few years, completely changing how people think about computing. The old approach of relying entirely on centralized cloud systems is no longer sufficient, whether due to scalability issues or security concerns. For these reasons, distributed edge architectures are becoming a hot issue, especially for the applications that demand ultra-low latency and have to work within tight bandwidth constraints in real-time [1]. Edge computing has quietly become the driving force behind some of the smart technologies from natural language processing to predictive analytics and beyond. Yet unlike traditional cloud setups, edge computing makes privacy a must, not just a nice-to-have.

Edge computing brings along a unique set of challenges that go well beyond what centralized systems face. Devices at the edge often run on uneven hardware, deal with unstable network connections and operate under tight resource constraints [2]. In machine learning and AI, pulling off high-performance while maintaining strong privacy in restricted edge environments is not easy. Mature privacy-preserving methods like differential privacy, homomorphic encryption, and secure multi-party computation, though powerful, tend to introduce computational burdens that are often too heavy for these lightweight edge devices [3]. As researchers have noted, the tug-of-war between privacy and performance lies at the very heart of distributed machine learning, especially when sensitive data faces potential exposure through a variety of attacks.

The field of machine learning is becoming increasingly aware that privacy is no longer optional, it is a legal requirement under major regulations like the GDPR (General Data Protection Regulation [4]) and the CCPA (California Consumer Privacy Act [5]). What makes this issue even more pressing is that distributed edge systems often amplify vulnerabilities through insecure communication channels. As a result, edge models face elevated risks from attacks such as membership inference, gradient leakage, and other sophisticated threats. This review explores a range of existing approaches,

including federated learning, lightweight cryptography, and emerging frameworks, all aimed at reconciling privacy with performance in real-world applications. Exciting new developments, such as the idea of privacy elasticity, suggest treat trade-offs as dynamic balances rather than fixed compromises. It is believed that promising progress would depend on interdisciplinary work that may span technical, legal, and organizational perspectives to contribute sustainable real-world distributed edge networks systems.

2. Related Work

The core characteristics of edge computing include heterogeneous hardware, unstable network connections, strict latency requirements, and a huge number of nodes. Although various technologies have been developed for performance and privacy, due to the particularity of edge computing, many are not applicable to edge technologies. Viewed from these constraints, some methods and technologies that have the potential to be used in edge computing are introduced below.

2.1. Federated Learning and Differential Privacy

McMahan et al.'s seminal work demonstrated that models can be trained collaboratively without centralizing sensitive data; this represents a breakthrough for many privacy-critical applications [6]. The effectiveness and range of applications of federated learning have been greatly increased by later research that used adaptive aggregation to increase communication efficiency and fault tolerance. Since then, federated learning steadily rises to prominence though it still has certain blind spots. For instance, it frequently presumes honesty and cooperation while ignoring problems that arise in actual deployments, such as adversarial attacks and Byzantine failures.

In order to accommodate diverse edge environments, differential privacy has developed in tandem with edge computing. However, early implementations revealed that this approach uses a lot of resources because of its consistent privacy budget [7]. Many researchers have since improved this procedure, modifying resource budgets according to model requirements, data sensitivity, and device capacity. However, there may be a big disconnect between theory and practice because differential privacy research frequently places more emphasis on mathematical simplicity than on practical deployability.

2.2. Cryptographic and Hardware-Based Methods

Advances in approximate homomorphic encryption (HE) and batching have made limited secure neural network inference possible. Yet practical deployment is hampered by key management and operational complexity - issues often ignored in worst-case theoretical analyses. Similarly, secure multi-party computation (MPC) has been tailored for ML through preprocessing and more efficient protocols, but device heterogeneity and unstable networks make real-world adoption difficult. Most MPC studies also neglect organizational and human factors that determine whether systems are actually deployed.

Hardware-assisted methods approach the problem differently. Trusted execution environments (TEEs) provide hardware-backed guarantees with lower overhead than cryptography [8]. Hybrid solutions combining TEEs with DP or HE show promise, but TEEs bring their own risks: hardware vulnerabilities, large, trusted computing bases, and supply chain security issues.

Key technologies, their advancements, and lingering difficulties are compiled in Table 1 below. To increase the effectiveness of these technologies and methods, researchers are now looking to automation, evolutionary algorithms, and reinforcement learning.

Table 1: Key Technologies for Privacy-Preserving ML and Their challenges

Technique	Key Benefits	Deployment Limitations
Federated Learning (FL)	Enables collaborative model training; improves communication efficiency and fault tolerance; adaptive aggregation enhances convergence	Often assumes fully cooperative participants; vulnerable to adversarial behavior, Byzantine faults, and economic incentive misalignment
Differential Privacy (DP)	Enables formal privacy guarantees; adaptive budgets improve privacy-performance balance	Inefficient uniform budgets; solutions often impractical for resource-constrained devices; weak privacy protection
Homomorphic Encryption (HE)	Enables secure computation on encrypted data; improves efficiency	High overhead costs; challenges in management, integration, etc.
Secure Multi-Party Computation (MPC)	Enables secure distributed computing under limited bandwidth/latency; preprocessing improves efficiency	Difficult to translate simulation results to heterogeneous edge systems;
Trusted Execution Environments (TEEs)	Hardware-based security; low overhead costs and strong guarantees	Vulnerable to attacks; requires large, trusted computing bases
Optimization Frameworks	Enables adaptive privacy-performance tuning; evolutionary and RL-based optimization	Neglect contextual and socio-political requirements

Many of these technologies and methods have made significant progress outside of the lab and into the real world of edge computing, according to this survey. Progresses have been made to bridge the gap between real-world applications and advanced theory, yet no single technical solution can address every problem or fit in every scenario. Deeply human, privacy is influenced by culture, laws, and most importantly, trust. Every strategy has its own trade-off, like additional processing power, bandwidth, or storage, making it even harder to balance privacy and performance. Specifically in edge computing, models must contend with scarce power, data, and patience, and many other issues.

2.3. Theoretical Tensions and Strategic Interpretations

Every one of the edge computing technologies mentioned above has special traits and capacities. All of them must, however, maximize the trade-off between performance and privacy. This problem has not yet been resolved. There are many reasons for this, not least of which is that the trade-off between privacy and performance is not straightforward; in many cases, it is chaotic, nonlinear, and even counterintuitive. Furthermore, each privacy protection mechanism incurs additional overhead, computational complexity, increased communication, and increased storage requirements. As previously mentioned, edge computing systems and devices face inherent constraints, such as readily available hardware, uncertain network connectivity, strict latency requirements, and a large number of nodes. The burden of privacy protection can rapidly degrade the system response speed of many edge devices, which can be detrimental in many applications. Regarding privacy performance trade-off, privacy cannot be defined as a single, simple concept. Instead, it often encompasses multiple facets and dimensions. For example, data privacy concerns the protection of sensitive information when it is processed or transmitted. Differential privacy and encryption usually play a role in this. Then there is model privacy, a sneakier issue: trained parameters can actually leak information about the data they have seen. And finally, there's communication privacy, which deals with hiding traffic patterns and metadata, so participants are not exposed by their interactions. Each of these areas has its own headaches, and they interact in ways that are hard to predict.

On the performance side, things are not straightforward either. Edge computing forces researchers to rethink what "performance" even means. Accuracy and throughput still matter, sure, but now

latency becomes non-negotiable. In the case of autonomous driving, the response time of several milliseconds is critical to avoid accidents and save life [9]. Energy consumption is easily one of the biggest headaches in edge-computing. It is almost common sense, if a battery-powered device runs out of juice while managing a privacy-preserving task, everything grinds to a halt. It is frustrating to think that even the smartest algorithms cannot save a system that is simply out of power. And then there's communication efficiency, edge systems rely on constant data exchange to stay functional, but here is the catch: bandwidth is limited, and transmitting data isn't cheap. Every bit sent over the network matters because small details can make or break an edge system.

Performance and privacy have a nonlinear trade-off. Sometimes people can increase privacy without having a big effect on performance, but as soon as people push privacy a bit further, suddenly the resource demands spike. A privacy-performance curve must first be drawn, and a "sweet spot" identified on the curve due to this nonlinearity, which makes purely intuitive adjustments practically impossible. The curve can be driven by many different factors. It is also noted that this "sweet spot" can vary significantly across different applications and deployments. For example, in healthcare applications, stronger privacy protections must be prioritized compared to home devices or autonomous driving applications, even if this means sacrificing speed or other factors. Furthermore, many regulations set very strict requirements and boundaries for these scenarios, which must be adhered to regardless of performance impact. In the second example, in home devices or autonomous driving, people may prioritize convenience and user experience; autonomous driving is also extremely sensitive to latency.

Although finding the optimal point mentioned above is very difficult in practical applications, and no effective and mature method or strategy currently exists to balance efficiency and privacy, some newly developed technologies still offer ideas and potential contributions toward a final solution. For example, the adaptive privacy mechanism is an example. This method does not turn privacy protection on or off but fine-tunes it according to the sensitivity of the data or the operating state of the system. It is like a smart thermostat rather than a switch. Another example is the hybrid architecture, which has the advantage of combining different technologies together, which may produce unexpected synergistic effects [10]. In addition, there are layered protection strategies, which can provide system designers with greater flexibility in terms of performance and privacy protection, and can set different levels of privacy protection according to threat models and performance constraints. The main advantage of these methods is that they move beyond a "one size fits all" approach and can be adjusted to fit different situations.

Specialized hardware becomes another useful means to balance privacy and efficiency. In recent development, cryptographic accelerators can greatly reduce the cost of running secure computing protocols. They allow edge computing systems to use methods that were once too resource intensive. Likewise, neural processing units (NPUs) for secure machine learning and processors optimized for privacy also show strong potential, often outperforming traditional general-purpose hardware. By allocating separate resources to privacy-preserving functions and doing away with the need to compete with other processes for cycles, edge-specific processors push the envelope even further. These specialized chips mark a shift because they are designed specifically for privacy workloads. Privacy thus becomes a crucial factor in dedicated infrastructure design.

Edge systems are characterized by heterogeneous hardware, unstable network connections, strict latency requirements, and a large number of nodes. Any technology analysis must take the specific limitations into consideration. Despite the fact that none of the aforementioned approaches are flawless for edge computing, combining these individual methods, for example, adaptive algorithms, hybrid methods, and specialized hardware, the result can be a transformative shift. Based on the results of this survey, it is anticipated that systems would be developed where privacy and performance complement each other, rather than against each other. Numerous industries, including healthcare, finance and artificial intelligence, will benefit from these developments.

3. Methodology

This section examines how individuals can address the complex balance between privacy and performance in contemporary distributed edge networks. The discussion centers on three main areas: federated learning systems, lightweight privacy techniques, and emerging technologies that offer promising solutions to several persistent challenges.

3.1. Federated Learning and Balance Optimization

Early versions of federated learning were basically just parameter averaging, but modern FL has gotten much more sophisticated. Today, adaptive aggregation, compressed communication, asynchronous coordination, personalization, and stronger privacy guarantees are all aimed at keeping the system efficient without compromising too much. Aggregation is where the first real breakthrough happened. The old uniform weighting approach didn't cut it; devices are just too different in power, reliability, and data quality. Adaptive aggregation flips the script by weighing updates based on who's trustworthy, what data looks like, and how much each contribution affects privacy. It is not perfect, but it is a huge leap toward making FL work outside the lab.

Communication efficiency has been a key pain point, especially at the edge where bandwidth is precious. Gradient compression and quantization techniques have slashed communication costs, often by orders of magnitude, while keeping accuracy intact [11]. Combine that with asynchronous protocols, which no longer force everyone to wait for the slowest device, and FL finally feels ready for messy, real-world deployments where connectivity is uneven, and devices check in at random. Personalization has also changed the game. A single global model just doesn't fit all use cases. Meta-learning helps clients adapt quickly to their own local quirks, while clustering strategies build group-specific models when data patterns align. In practice, these approaches often outperform the one-size-fits-all strategy.

Still, big gaps remain. Most FL systems treat privacy needs as static, when they change over time. A temporal accounting framework - basically a private budget that can go into debt or credit - could help. Even more pressing, FL often assumes uniform privacy requirements, which almost never matches reality. Despite all the hype about decentralization, most FL deployments still lean heavily on centralized servers. This highlights that FL is only a solid foundation, and complementary, lightweight approaches are needed to truly make privacy-preserving AI work at the edge.

3.2. Lightweight Privacy-Preserving Techniques and Performance Optimization

Because edge devices have limited resources, privacy techniques that are tailored for low-power, memory-constrained settings are required. Many of the techniques used today require a lot of energy; encryption techniques, for instance, require a lot of computation. This has led researchers to look for low-power alternatives that balance privacy and performance while offering strong security [12]. These developments include reduced homomorphic encryption, intelligent selective privacy applications, optimized secure computation, approximate differential privacy, and other tactics tailored to edge environments.

On the basis of approximation and flexibility, numerous lightweight techniques have been created. In approximate differential privacy, the system can become effective enough for real-world deployment in limited edge environments if one is prepared to accept somewhat weaker privacy guarantees. Additionally, by employing batching techniques and approximations, lightweight homomorphic encryption lowers overhead costs, enabling encrypted computing on edge devices. In contrast to them, selective privacy applications become intelligent and adaptable: they can prioritize the protection of high-risk data by using sensitivity analysis, and dynamic protocols can modify the privacy level in response to emerging threats and resource constraints. Secure computing protocols can reduce costs through offline preprocessing and communication compression, thereby allowing multiple parties to work together without completely overwhelming network capacity [13]. Location-aware processing reduces data transmission between nodes, thereby reducing latency and privacy

risks. Hierarchical processing distributes computation across multiple layers, allowing us to apply privacy mechanisms in stages.

The survey also uncovered several promising lightweight approaches applicable to edge environments. One of them, the privacy gradient framework treats privacy as a dynamic property that adapts to data sensitivity, available resources, and real-time threats. In contrast, cooperative privacy enhancement allows powerful devices to assist their weaker peers in the system, forming a distributed collaborative privacy network. Preliminary results show that challenges remain, such as the failure of small-scale methods at large scale and the continued reliance on approximations without reliable privacy guarantees. The next wave of research needs to tackle these gaps in order to create systems that are both effective and demonstrably private in the real world.

Many approaches have been developed for edge computing and applied to different areas of AI and machine learning. Recent advances have overcome earlier limitations such as small memory, simple instruction sets, and weak attestation, and now TEE integrity can even be verified across distributed systems [14]. Specialized hardware is also speeding up privacy-preserving computation. ML-based methods like reinforcement learning and multi-objective optimization help systems learn from experience and find Pareto-optimal trade-offs [15]. New technical developments, such as neural processors, secure training with TEEs, and ML-driven optimization, make it possible to run powerful AI safely. With smart planning and predictive tweaks, autonomous AI respects privacy is starting to become a reality. Nevertheless, it should also be recognized that each method has strengths and weaknesses, and while combining them can be promising, it can also introduce new challenges. For instance, in high-privacy fields like healthcare, federated learning alone is not enough and must be paired with differential privacy, which can add noise and affects model performance.

4. Evaluation Metrics and Multi-Objective Assessment

It is especially difficult to evaluate privacy-preserving machine learning systems because of the special limitations that edge computing brings. In this regard, evaluation frameworks should take into account the intricate relationship between privacy, performance, and usability, rather than evaluating metrics separately as is often done in traditional approaches. To ensure both theoretical soundness and practical robustness, privacy assessment should be thorough, combining empirical tests like reconstruction attacks and membership inference with formal guarantees like differential privacy. To reflect the diverse demands of edge applications, performance evaluation, on the other hand, needs to take into consideration a number of factors, such as latency, throughput, resource usage, communication efficiency, and the like [16]. Table 2 below offers a comprehensive framework that combines privacy, performance, and usability in order to reflect this complexity.

Table 2: Holistic Evaluation Framework for Privacy-Preserving Edge ML Systems

Dimension	Sub-Dimensions / Metrics	Description / Notes
Privacy	Differential Privacy Guarantees	Formal ε-differential privacy bounds
	Empirical Privacy Leakage	Membership inference, reconstruction attacks
	Model Confidentiality	Leakage from model parameters or updates
	Communication Privacy	Metadata leakage, traffic analysis
	Regulatory Compliance & Contextual Sensitivity	Legal requirements, context-specific privacy needs
Performance	Computational Efficiency	Processing latency, and resource utilization
	Communication Efficiency	Bandwidth utilization, transmission latency, protocol overhead
	Energy Consumption	Power usage, battery life for edge devices
	Accuracy / Task Effectiveness	Model accuracy, inference quality
Usability	User Perception & Acceptance	Subjective privacy perception, ease-of-use
	System Usability & Transparency	System usability, adaptation to context, transparency
	Application Criticality	Safety-critical requirements, operational constraints

The aforementioned framework is intended to inspire researchers to develop truly multi-dimensional evaluation systems. At its core, a system needs a thorough approach that investigates the intricate connections and unavoidable trade-offs between privacy, performance, usability, and other elements like robustness [17]. Furthermore, in order to move beyond static, one-time evaluations, dynamic evaluation might take center stage. Static snapshots simply cannot capture the "privacy-performance trajectory" of a system, which can be traced over time to demonstrate how it changes or evolves in real-world scenarios. This temporal perspective may help assessments better capture the messy reality of distributed edge environments. By enabling testing in disorderly, unpredictable settings, scenario-based benchmarks avoid making evaluation outcomes unduly reliant on an ideal laboratory setup. For evaluation to be effective, context is essential. An approach that is effective in the privacy-sensitive healthcare sector, for instance, might not work in the reliability-critical industrial system of autonomous driving, where milliseconds matter. Instead of using a one-size-fits-all strategy, each domain needs a customized, well-balanced approach.

The human element is just as important. The trade-off between privacy and performance is about people's comfort, expectations, and trust. Evaluations ought to combine empirical data with user-centered studies on usability and privacy perception. It is necessary to evaluate whether people genuinely accept and trust these systems. Future research should adopt this comprehensive approach, observing the joint evolution of privacy, performance, and robustness across domains and over time.

5. Conclusions and Future Research

In distributed edge networks, striking a balance between privacy and performance is extremely challenging. Workloads can abruptly increase or decrease in the constraint edge environment where devices are frequently added or removed. In addition, user expectations and legal requirements may change at any time. The survey revealed that while privacy-preserving techniques for edge computing have made great strides, the core trade-offs between performance and privacy remain unresolved. Yet prior research indicates that rather than viewing these conflicts as obstacles, they should be seen as optimization problems that are difficult but ultimately manageable.

As an illustration, consider federated learning, which permits cooperative model training while maintaining privacy. The deployment is still difficult, though. Adaptive aggregation, communication costs, and system resilience remain unresolved. Even worse, many approaches that shine in lab settings break down in the wild because of heterogeneity, unpredictability, and device failures. This "deployment gap" is not surprising: academic incentives often reward novelty, not robustness. In many cases, the messy leap from proof-of-concept to production is chronically underestimated. Future studies must adopt more comprehensive, wide-ranging approaches rather than making isolated tuning adjustments. Several directions (see below) appear particularly promising based on the survey.

Automatic optimization: According to the survey, automatic optimization provides a workable solution that requires little human involvement to balance properly between privacy and performance. Systems can discover efficient trade-offs that would be challenging to manage manually by using machine learning and reinforcement learning to dynamically modify resource allocation, privacy budgets, and aggregation strategies in real time. It is similar to giving the system intuition, which relieves engineers of the need for constant manual adjustment. That flexibility is not only desirable, but necessary in large-scale edge environments. ML-driven optimization may not only reduce overhead as people may believe; it may also reveal astute trade-offs in real time that people might overlook.

Interdisciplinary collaboration: It takes more than just engineering to strike a desirable balance between performance and privacy. Collaboration between disciplines is necessary for real progress: psychologists to comprehend how people truly value privacy, economists to create incentives, lawyers to navigate changing laws like the CCPA and GDPR, and computer scientists to build. Without that synergy, even the most advanced technical solutions risk failing against the chaotic

reality. Even the most sophisticated technical solutions run the risk of collapsing against the chaotic reality in the absence of that synergy.

Hardware and Standards: The innovations in specialized hardware are quite promising. With TEEs, cryptographic accelerators and privacy-focused chips, edge systems could be made secure and powerful. However, the survey also found that cost and accessibility remain major obstacles. To prevent fragmentation, compare applications fairly and adjust to evolving regulations, new practical standards must also be established.

Long-Term Monitoring: Long-term monitoring is not just some technical checklist item - it is like keeping finger on the pulse of a system's real life. It is necessary to keep watching how privacy and performance balance out as user habits evolve, networks get faster (or slower), and laws evolve. The world does not hit pause, so neither can the system. Honestly, adaptive systems need constant feedback to keep learning and adjusting to the environments they are supposed to protect.

In concluding this survey on the balance between privacy and performance in edge systems, the true path forward in distributed edge computing lies in viewing performance and privacy as complementary forces, strengthened by interdisciplinarity, automation, hardware acceleration, and long-term validation. Alone, none of these can fully bridge the deployment gap, but together, they hold real promise. If researchers, practitioners, and policymakers work together instead of in silos, there is a real chance of building systems that not only succeed in controlled experiments but also perform reliably in the real world - achieving a balance between privacy and performance without constantly forcing trade-offs.

References

[1] F. R. Mughal, J. He, B. Das, et al., "Adaptive federated learning for resource-constrained IoT devices through edge intelligence and multi-edge clustering," Sci. Rep., vol. 14, art. 28746, 2024.

[2] J. Hong, Y. Hong, X. de Foy, M. Kovatsch, E. Schooler, and D. Kutscher, "Internet of Things (IoT) edge challenges and functions," Internet Research Task Force, RFC 9556, Apr. 2024.

[3] L. Cui and X. Wu, "ALDP-FL for adaptive local differential privacy in federated learning," Sci. Rep., vol. 15, art. 26679, 2025.

[4] European Parliament and Council of the European Union, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)," Official Journal of the European Union, vol. L119, pp. 1-88, May 4, 2016.

[5] State of California, "California Consumer Privacy Act of 2018," California Civil Code, Division 3, Part 4, Title 1.81.5, §§ 1798.100-1798.199, 2018.

[6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Arcas, "Communication-efficient learning of deep networks from decentralized data," in Proc. 20th Int. Conf. Artif. Intell. Statist., vol. 54, 2017, pp. 1273-1282.

[7] D. Wang and S. Guan, "FedFR-ADP: Adaptive differential privacy with feedback regulation for robust model performance in federated learning," Inf. Fusion, vol. 116, art. 102796, 2025.

[8] F. Mo, H. Haddadi, K. Katevas, E. Marin, D. Perino, and N. Kourtellis, "PPFL: Privacy-preserving federated learning with trusted execution environments," Proc. 19th Annu. Int. Conf. Mobile Syst., Appl., Serv., 2021, pp. 94-108.

[9] P. Giannakopoulos, B. van Knippenberg, K. C. Joshi et al., "Key metrics for monitoring performance variability in edge computing applications," EURASIP J. Wireless Commun. Netw., vol. 2025, art. 38, 2025.

[10] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, and Y. Zhou, "A hybrid approach to privacy-preserving federated learning," Proc. 12th ACM Workshop Artif. Intell. Secur., 2019, pp. 1-11.

[11] Y. Ren, Y. Cao, C. Ye et al., "Two-layer accumulated quantized compression for communication-efficient federated learning: TLAQC," Sci. Rep., vol. 13, art. 11658, 2023.

[12] B. Gușiță, A. A. Anton, C. S. Stângaciu, D. Stănescu, L. I. Găină, and M. V. Micea, “Securing IoT edge: a survey on lightweight cryptography, anonymous routing and communication protocol enhancements,” Int. J. Inf. Secur., 2025.

[13] A. Nof, E. Boyle, N. Gilboa, and Y. Ishai, “Secure multiparty computation with sublinear preprocessing,” Advances in Cryptology - EUROCRYPT 2022, IACR, 2022.

[14] S. Naas, L. Rupprecht, D. Kales, and M. Franz, “Open and resilient remote attestation on decentralized platforms,” arXiv preprint arXiv:2402.08908, 2024.

[15] H. Xu, Y. Zhang, and S. Tang, “Pareto set learning for multi-objective reinforcement learning (PSL-MORL),” arXiv preprint arXiv:2501.06773, 2025.

[16] B. Jayaraman and D. Evans, “Evaluating differentially private machine learning in practice,” in Proc. 28th USENIX Security Symp., 2019, pp. 1895-1912.

[17] H. Habib, K. S. Shrestha, J. Jung, and L. F. Cranor, “Evaluating the usability of privacy choice mechanisms,” in Proc. 18th Symp. Usable Privacy and Security (SOUPS), 2022, pp. 1-17.