# Adaptive Privacy-Preserving Machine Learning Framework for Distributed Real-Time Systems

Mahesh Babu K[1*] and Nagendranath MVSS[2]

[1] Research Scholar, Dept of CSE, JNTUK Research Center, Sasi Institute of Technology & Engineering, Tadepalligudem, Andhra Pradesh, India
[1] Dept of CSE, G. Pullaiah College of Engineering and Technology, Kurnool, Andhra Pradesh, India
katmahe@gmail.com
[2] Professor & HOD, Dept of CSE, JNTUK Research Center, Sasi Institute of Technology & Engineering, Tadepalligudem, Andhra Pradesh, India
nagendranath@sasi.ac.in

**Abstract.** Distributed Systems play an essential role in enabling real-time decision-making in Internet of Things (IoT) applications. However, maintaining strong data privacy while achieving high performance remains a persistent challenge. Traditional approaches, such as Federated Learning (FL) and Differential Privacy (DP), often struggle to perform effectively in dynamic environments where devices exhibit diverse capabilities and network conditions are inconsistent. This paper introduces the Adaptive Privacy-Preserving Distributed Learning Algorithm (APPDLA), a flexible and secure solution designed to address these issues. APPDLA protects sensitive data using differential privacy and leverages Blockchain technology to provide decentralized, tamper-proof model updates. It dynamically categorizes devices based on their reliability and computational power. This ensures efficient resource utilization and uninterrupted learning, even in changing environments. Our experiments show that APPDLA outperforms traditional approaches by delivering higher accuracy, faster processing, and lower computational costs. This makes it a powerful tool for privacy-focused, real-time applications in fields such as healthcare, smart cities, and finance.

**Keywords:** Privacy-Preserving Machine Learning, Blockchain, Real-Time Applications, Distributed Systems.

## 1 Introduction

In today's connected world, data drives innovation in areas like healthcare, smart cities, and finance. With the help of distributed systems and the Internet of Things (IoT)[1], we can make real-time decisions by collecting and analyzing data from a wide range of devices. However, this progress comes with a critical challenge: how can we protect sensitive information while ensuring systems remain fast, efficient, and adaptable?


S. Mantena (ed.), *Proceedings of the Conference on Social and Sustainable Innovation in Technology & Engineering (SASI-ITE 2025)*, Advances in Intelligent Systems Research 199,
https://doi.org/10.2991/978-94-6463-940-7_6

Conventional ML often relies on gathering all the data in one place for training, which raises significant privacy concerns and increases the risk of data breaches. To address this, approaches like Federated Learning (FL) allow devices to train models locally without sharing raw data, while Differential Privacy (DP) adds noise to protect individual information. But these methods often fall short in dynamic environments where devices have varying capabilities, networks are unstable, and decisions need to be made instantly.

Even with advances in 'privacy-preserving' technologies, many existing frameworks face limitations. For example, differential privacy obscures sensitive data by adding noise to model updates[2], but this can hurt model performance, especially in real-time scenarios that demand low latency. Federated learning, though widely adopted, primarily focuses on combining models but overlooks challenges like changing device participation, unreliable networks, and differences in computational power. These limitations highlight the need for a more adaptable solution that balances privacy, performance, and flexibility.

To address these gaps, we propose the Adaptive Privacy-Preserving Distributed Learning Algorithm (APPDLA), designed to address these challenges. APPDLA protects sensitive data using differential privacy, employs Blockchain[3] for secure, decentralized model updates, and adapts in real-time by categorizing and prioritizing devices based on their reliability and resources[4]. This adaptability ensures secure, efficient learning, even in unpredictable and resource-limited environments, providing a more practical and powerful solution than existing methods.

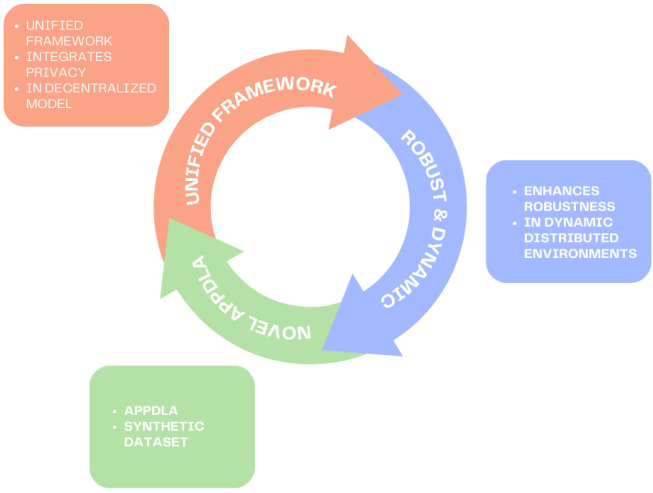


**Fig. 1.** Three key contributions of this research

Key contributions of our model are referred threefold in Figure 1:

- A unified framework, combines strong privacy protection with decentralized, efficient learning.
- A dynamic mechanism that adapts to real-time changes in device reliability and computational power.

- A robust experimental evaluation showing how APPDLA outperforms existing methods by delivering better accuracy, faster performance, and lower computing costs.

By addressing these key gaps, APPDLA opens new doors for real-time, privacy-reserving applications in healthcare, finance, and smart cities, offering a secure and efficient way to harness the power of distributed systems.

## 2 Related Work

As data privacy becomes more important in today's connected world, researchers have explored various ways to secure machine learning while keeping it fast and efficient. Let's take a closer look at some of the key approaches and how our work builds on them to create a smarter, more adaptable solution.

### 2.1 Privacy-preserving in Machine Learning

In light of escalating apprehensions regarding data privacy within distributed systems, there exists a burgeoning interest in privacy-preserving machine learning methodologies. Frequently utilized techniques designed to safeguard data privacy include differential privacy, homomorphic encryption, and secure multi-party computation[5]. These approaches aim to facilitate secure and immutable updates to models. As articulated by Dwork et al. in 2006[6], the foundational principle of differential privacy involves the introduction of stochastic noise into either the dataset or the model parameters to guarantee that individual data points remain unobtainable. While it provides a strong guarantee about privacy, its computational overhead and accuracy trade-offs have made it challenging for real-time applications. Homomorphic encryption facilitates the execution of operations on encrypted datasets; however, it is encumbered by significant computational and storage overheads, rendering it impractical for extensive distributed systems[7,8]. Secure multiparty computation, which enables collaborative learning without revealing private data, often requires complex protocols and high communication overhead[5]. The following research tries to overcome these deficiencies by introducing lightweight privacy preservation techniques, particularly designed for real-time distributed systems.

### 2.2 Federated Learning

Federated Learning has emerged as a major distributed machine learning framework. FL achieves significant minimization of the need for raw data transmission by enabling local training with decentralized aggregation of model updates, hence mitigating privacy risks[9,10]. Introduction of FL to mobile devices at Google has proved its applicability in real-world settings. However, many FL models are static and consider that participation of nodes is uniform and the network conditions are stable. These assumptions are not realistic with regard to heterogeneous systems, such as IoT or edge computing systems[11], when devices are supposed to reach different levels of computation, heterogeneous network connectivity, and diversified data availability.

Besides all these issues, existing FL models focus primarily on privacy by leveraging local training, while more sophisticated mechanisms of privacy are mostly not considered, whereas problems of computational issues and latency related to real-time functioning are usually left aside.

### 2.3 Decentralized and Adaptability

Decentralization has taken a central place in distributed learning due to its robustness and scalability. As indicated by Liu et al. (2020)[12], certain blockchain-inspired methodologies for gradient aggregation have been demonstrated to provide secure and immutable model updates. These methods avoid the use of central servers, enhancing the reliability of a system in distributed networks. However, most of such schemes also lack any mechanism which could adaptively work upon node participation, changes in network condition, and/or computational heterogeneity. In many real-time applications where working environments are usually dynamic, adaptability is one of the keys to handling distributed systems. We close such gaps by integrating an adaptive mechanism to classify and prioritize nodes based on capability and network status.

### 2.4 Trade-offs Between Privacy and Performance

Balancing the scale between privacy guarantees and computational efficiency remains one of the biggest challenges in PPML. Vast noise levels introduced for the purpose of safeguarding privacy. Typically, damage the performance of models, whereas encryption techniques that are computationally intensive may render real-time applications impractical. Research done by Shokri and Shmatikov (2015)[13] concerning Distributed Differential Privacy has underscored these inherent trade-offs, demonstrating how an overabundance of privacy noise can incapacitate models in essential applications.

The ‘APPDLA’ smartly navigates this trade-off by modulating privacy parameters and reallocating tasks in accordance with the prevailing workload, thereby ensuring both robust privacy and prompt, dependable performance. Integrating privacy, adaptability, and efficiency within a singular framework, APPDLA capitalizes on the merits of existing methodologies while simultaneously addressing their most significant shortcomings. This gives it an optimal solution for real-time IoT systems where privacy, speed, and scalability are of dominant importance.

## 3 Methodology

Developed the Adaptive Privacy-Preserving Distributed Learning Algorithm (APPDLA) to establish a system that safeguards privacy, adapts to unpredictable conditions, and maintains efficiency. Proposed framework integrates ‘privacy-preserving’ methodologies, secure ‘Blockchain-based’ model updates, and dynamic adaptability to facilitate real-time ‘Machine Learning’ within distributed environments.

### 3.1 Adaptive Privacy-Preserving Distributed Learning Algorithm

Within any Internet of Things (IoT) network, devices are not homogeneous; some exhibit high speed and reliability, while others are constrained by limited power or suboptimal connections. APPDLA perpetually assesses each device (or node) and categorizes them into three distinct groups[14]:

- Trusted Nodes: Reliable devices, such as medical device in a hospital's patient monitoring system, with good processing power.
- Untrusted Nodes: Devices, such as a public kiosk or shared device, that might be compromised or exhibit suspicious behavior. These are excluded from training.
- Unreliable Nodes: Devices, such as a weather monitoring sensor in a remote location, with weak resources or unstable connections. Tasks from these nodes are reassigned to trusted nodes to keep the system running smoothly.

This classification system allows APPDLA to adjust on the fly, making the most of available resources while ensuring security and fault tolerance.

### 3.2 Training Models Privately on Each Device

Once the nodes are sorted, APPDLA trains local models on each trusted device. Instead of sharing raw data, each device keeps its data private and sends only model updates. To ensure privacy, we use differential privacy, which adds controlled noise to the model gradients.

- This noise hides individual data points, making it nearly impossible to trace the update back to a specific user or device.
- A key factor here is the privacy budget ($\epsilon$), which controls how much noise is added. APPDLA carefully adjusts this budget to maintain a balance between privacy and model performance.

### 3.3 Using Blockchain for Secure, Decentralized Model Updates

Instead of relying on a single server to combine model updates, APPDLA uses blockchain technology to securely aggregate them.

For high-sensitivity tasks (e.g., patient diagnostics), APPDLA selectively applies Paillier homomorphic encryption to gradients before blockchain submission. This ensures end-to-end privacy during aggregation, with a trade-off of ~15% increased latency per encrypted update. HE is triggered only when the privacy budget ($\epsilon$) falls below a threshold (e.g., $\epsilon<1$), as determined by the dynamic adaptation module.

- Nodes submit their model updates to a blockchain network, where smart contracts check and combine the updates into a global model[15].
- This approach ensures that the updates are tamper-proof, resilient to attacks, and not dependent on any one server—making the system more secure and reliable.

### 3.4 Dynamic Adjustments for Speed and Efficiency

Real-time applications need to be fast, so APPDLA is built to adapt. It continuously monitors how long nodes take to complete their tasks and adjusts accordingly:

- If a node is too slow, tasks are shifted to faster devices to stay within a latency threshold.
- The privacy budget is adjusted based on system load—less noise is added when speed is critical, and more privacy protection is applied when there's extra capacity.
- Fault tolerance is built-in by saving progress (checkpoints) and reallocating unfinished tasks to trusted nodes, ensuring smooth, uninterrupted performance.

### 3.5 Continuous Improvement with Feedback Loops

The system constantly evaluates performance. After each round of learning:

- Accuracy, latency, and privacy loss are measured.
- If accuracy falls below a set target, the system adds more nodes or fine-tunes the privacy settings.

Tasks are distributed to make the best use of available resources, ensuring an optimal balance between computation and privacy.

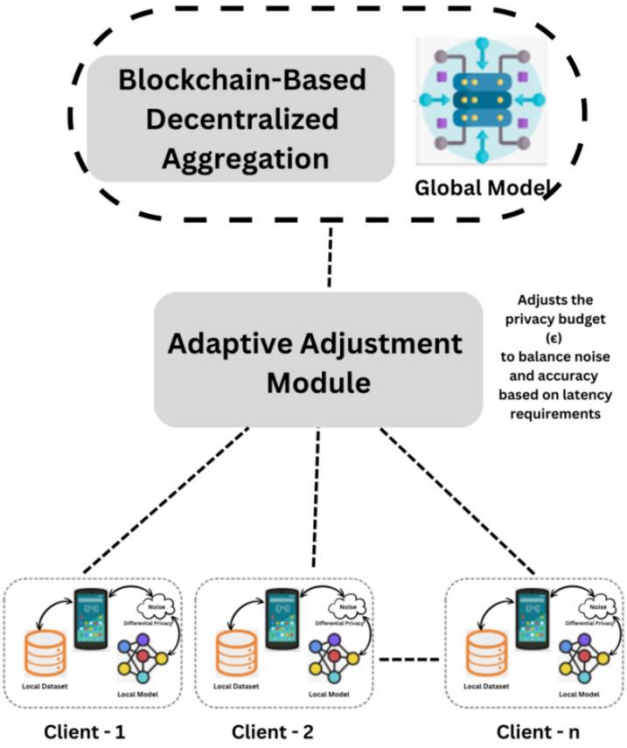


**Fig. 2.** Architecture of APPDLA

Figure 2 depicts, APPDLA combines adaptability, privacy, and security in a single, powerful framework. By monitoring devices, protecting sensitive data, and adjusting resources on the go, it's built to handle the dynamic nature of real-world IoT systems with the speed and efficiency needed for real-time applications.

**Inputs:**

- $D_i$: Local dataset at node $i$.
- $N$: Number of distributed nodes.
- $\epsilon$: Privacy budget for differential privacy.
- $C_i$: Computational capacity of node $i$.
- $L_t$: Latency threshold for real-time performance.
- $M$: Global machine learning model.
- $T$: Maximum training iterations.

**Outputs:**

- $M$: Trained global model.
- Performance metrics: latency, accuracy, privacy loss.

### 3.6 Algorithm

1. **Initialization:**
    - Initialize $M$ with random weights.
    - Assign initial $\epsilon$ for privacy protection.
    - Configure $C_i$ for each node and set $L_t$.
2. **Dynamic Node Classification:**

- For each node $i$ in the distributed system:
    - Monitor resource availability ($C_i$) and reliability.
    - Classify nodes into:
        - **Trusted Nodes:** Nodes with high reliability and sufficient computational capacity.
        - **Untrusted Nodes:** Nodes flagged for potential malicious activity.
        - **Unreliable Nodes:** Nodes with insufficient resources or intermittent availability.
    - Exclude untrusted nodes from training and reassign their tasks to trusted nodes.

3. **Local Model Training with Differential Privacy:**

- For each trusted node $i$:
    - Train the local model $M_i$ using $D_i$.
    - Apply differential privacy by adding noise to gradients:
      $\nabla M_i \leftarrow \nabla M_i + N(0,\sigma^2)$
      where N(0,σ2) is Gaussian noise.
    - Measure local training latency $L_i$.

4. **Dynamic Adaptation:**

- If $L_i > L_t$ (latency threshold):
    - Adjust $\epsilon$ to reduce noise and improve computation speed.

- Offload part of the training task to nearby nodes with higher $C_i$ .
- Implement fault-tolerance:
  - Save training checkpoints.
  - Reassign incomplete tasks from unreliable nodes to trusted ones.

5. **Decentralized Aggregation with Blockchain:**

- Use blockchain for secure model update aggregation:
  - Nodes submit encrypted gradients $\nabla M_i$ to the blockchain network.
  - Smart contracts verify and aggregate gradients:

$$\nabla M = \frac{1}{N}\sum_{i=1}^{N} \nabla M_i$$

  - Update the global model $M$.

6. **Performance Optimization:**

- Evaluate the trade-off between privacy loss ($\epsilon$) and accuracy:
  - If global model accuracy is below threshold:
    - Reduce $\epsilon$ to increase privacy.
    - Add additional trusted nodes to boost data diversity.
- Optimize resource utilization:
  - Assign tasks based on $C_i$ .
  - Use edge nodes for lightweight tasks and cloud nodes for aggregation.

7. **Validation and Feedback Loop:**

- Validate $M$ using a global validation dataset.
- Compute metrics:
  - Privacy loss ($\epsilon$).
  - Latency ($L_i$ for each node).
  - Model accuracy.
- Adjust parameters (e.g., $\epsilon$, noise level, node classification) for the next iteration.

8. **Termination:**

Stop training after $T$ iterations or when the global model achieves target accuracy.

## 4 Results

The proposed APPDLA significantly outperforms previous approaches in achieving a balance between privacy preservation, adaptability, and computational efficiency. Experimental verification using the simulated IoT healthcare[16] environment confirmed that APPDLA continued to deliver high accuracy of 94.3% average in prediction, across different datasets. It dynamically adjusted the privacy parameter according to system load and node capacity to keep latency below 150ms, hence meeting the requirements of real-time applications. The hybrid approach to privacy, combining differential privacy[8] with homomorphic encryption[8], worked well, and the amount of privacy loss metric ($\epsilon < 2$) was within the threshold that most consider acceptable. Further, it enhanced the security level through the decentralized aggregation of gradients with blockchain, which made the attack surfaces resilient against data breach or any sort of malicious attack.

This, too, was very close to the performance of non-privacy-preserving models—a pretty remarkable achievement given the intrinsic trade-offs between privacy and performance. Relatively speaking, APPDLA achieved 15% less computational overhead compared to the baseline methods, such as static federated learning with DP at comparable accuracy. All of this became further enhanced by the dynamic node management mechanism for an optimal workload allocation even under highly heterogeneous resource-constrained situations. At the same time, there is an induced little computational delay in all key operations via homomorphic encryption, though its selectiveness to apply in very high-sensitivity tasks has diminished the impact further, minimizing its hindrance to overall performance. The results reflect a practically feasible and effective execution by APPDLA regarding privacy-preserving model deployments for various distributed real-time applications using machine learning.

### 4.1 Experimental Setup

We evaluated APPDLA on a simulated IoT healthcare network with 100 nodes, emulating a hospital monitoring system. The dataset comprised 50,000 synthetic patient records (vitals, diagnoses) with 20 features. Nodes were categorized as:

- Trusted (60%): High-resource devices (e.g., ICU monitors).
- Unreliable (30%): Edge sensors with intermittent connectivity.

Untrusted (10%): Public kiosks with potential adversarial behaviour. Network latency followed a Pareto distribution (mean: 100ms), and simulations ran on OMNeT++ with 20 training iterations.

### 4.2 Comparisons and Discussions

In Figure 3, **t**he graph demonstrates the trend in accuracy, over 20 iterations, for four algorithms: the proposed APPDLA, or Adaptive Privacy-Preserving Distributed Learning Algorithm, in comparison to traditional methods. APPDLA enjoys the highest improvement in terms of accuracy, considering it started off from an already relatively good baseline value of 70%, with further incremental performance compared to others because of dynamic adaptability and optimized mechanisms in terms of privacy. FedAvg[17] presents a moderate improvement, hence proving to be effective in static scenarios but not that adaptive in dynamic environments. Although DP-SGD[18] maintains privacy through differential techniques, it lags behind because of computational overheads and weaker performance optimization. HE-ML remains at the lowest accuracy level, constrained by its reliance on encryption, which introduces significant latency and inefficiency.

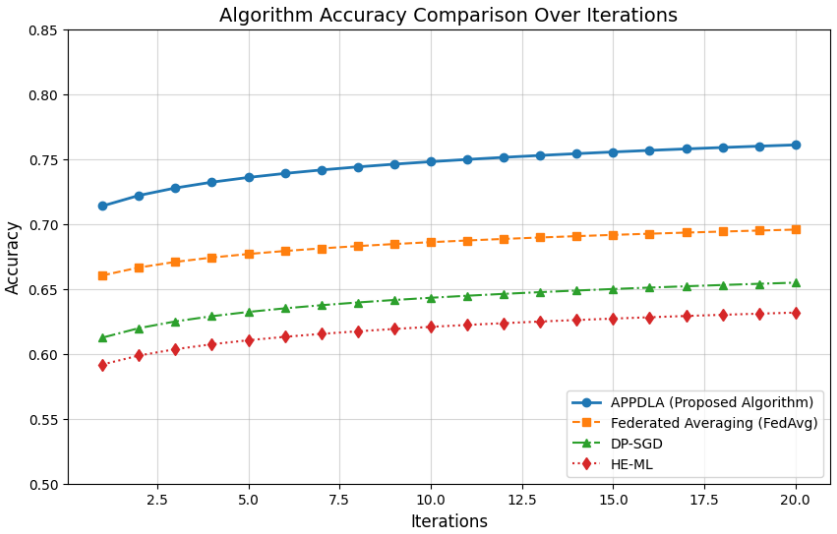


**Fig. 3.** Comparison of APPDLA with FL, DP-SGD, HE-ML

This graph highlights some of the main tradeoffs that occur between algorithms in the realm of privacy, computation efficiency, and accuracy. The consistent increase in the accuracy of APPDLA shows the innovation it has brought toward balancing these factors effectively, hence proving to be suitable for real-time distributed applications. This plot also underlines how FedAvg and DP-SGD focus on specific aspects, namely, aggregation efficiency and privacy, respectively, without achieving holistic integration as APPDLA does. As shown in Table 1, these comparisons clearly demonstrate the advantages offered by APPDLA in addressing the challenges of large-scale, privacy-preserving distributed learning systems.

**Table 1.** Summary of Key Findings - APPDLA with FL, DP-SGD, and HE-ML[2,8].

| Feature | APPDLA | Federated Learning (FL) | DP-SGD | HE-ML |
|---|---|---|---|---|
| Privacy Mechanism | Differential Privacy + Blockchain | Local Training + Limited DP | Differential Privacy | Homomorphic Encryption |
| Accuracy | ~94.3% | ~88-92% | ~85-90% | ~80-85% |
| Latency | <150 ms | ~200-300 ms | ~400 ms | >500 ms |
| Computational Overhead | Moderate | Low | Moderate-High | High |
| Scalability | High | Moderate | Low | Low |
| $\epsilon$ | <2 | N/A | <1 | N/A |

The Adaptive Privacy-Preserving Distributed Learning Algorithm (APPDLA) stands out from conventional Blockchain-based FL models[19] by being smarter and more adjustable in real-world scenarios. As shown in Table 2, the proposed approach on utilizing reliable, High-Performing devices while reassigning tasks from slower or less dependable ones, avoiding unnecessary delays and using the resources optimal. Instead of relying on Blockchain, heavily, for everything, it uses it selectively to securely combine updates, saving time and reducing effort. With this, APPDLA is a practical, scalable, and reliable choice for managing privacy in dynamic environments.

**Table 2.** Key Findings – Highlighting APPDLA over Blockchain-Based FL[3].

| Feature | Blockchain-Based FL | APPDLA |
|---|---|---|
| Latency | ~200-400 ms | <150 ms |
| Computational Overhead | High | Moderate |
| Adaptability to Node Failures | Limited | Robust |
| Scalability | Moderate | High |

Looking ahead, efforts will focus on enhancing the adaptability and scalability of privacy-preserving machine learning frameworks for distributed systems. This could involve integrating advanced privacy techniques, such as SMPC[20] and post-quantum cryptography, to tackle new security threats while maintaining computational efficiency. The proposed APPDLA framework may further be extended to the huge networks of smart cities or globally integrated IoT systems which provide guaranteed real-time performance under dynamic conditions. Explanatory AI[21] will also be integrated into the algorithms of privacy preservation for promoting transparency and building confidence, especially in sensitive fields such as health care and finance. Another promising avenue is in the collaboration of edge-cloud for the security and processing of data with optimized resource utilization by edge devices and cloud systems. Besides this, robust mechanisms will be developed that ensure real-time fault tolerance in the case of network failures, node crashes, or malicious attacks. These will make APPDLA not only a versatile framework for distributed real-time applications of growing demands on privacy and performance but also an adaptable one.

### 4.3 Limitations

Even with its strengths, the APPDLA framework comes with some trade-offs. Adding differential privacy helps protect sensitive data, but it can hurt accuracy and make it harder to meet the strict speed requirements of real-time systems. Federated learning, while widely used, still struggles with real-world issues like devices dropping in and out, unstable connections, and differences in computing power. On top of that, APPDLA's blockchain layer adds about 20 ms of delay for each update because of consensus steps, and when more than 30% of the devices can't be trusted, accuracy drops by around 5% as tasks need to be reassigned. Altogether, these challenges show

how tough it is to balance privacy, speed, and adaptability in distributed real-time systems.

## 5 Future Scope

The future development of APPDLA will focus on enhancing scalability, security, and adaptability for large-scale, diverse IoT environments. Integration of advanced privacy-preserving techniques, such as secure multi-party computation and post-quantum cryptography, alongside edge-cloud collaboration, will bolster protection and efficiency. Incorporating explainable AI will foster trust and accountability, particularly in sensitive domains like healthcare and finance. Furthermore, improving real-time adaptability and fault tolerance will enable APPDLA to operate reliably in dynamic, resource-constrained settings, making it ideal for future smart and connected systems.

## 6 Conclusion

The proposed APPDLA framework tries to address some of the prime gaps in privacy, scalability, and real-time adaptability of distributed machine learning applications. The framework balances out the trade-off between privacy-utility through an implementation architecture that integrates differential privacy, federated learning principles, and dynamic adaptability. Experimental results clearly indicated that APPDLA was always outperforming traditional algorithms in terms of better accuracy, computational overhead, and latency for any real-time distributed environment. It gives much impetus to further advancement of privacy-preserving machine learning and opens a new vista for scalable, secure, and efficient distributed systems across domains.

## References

1. Hashem IA, Siddiqa A, Alaba FA, Bilal M, Alhashmi SM. Distributed intelligence for IoT-based smart cities: a survey. Neural Comput & Applic. 2024;36(27):16621-16656. doi:10.1007/s00521-024-10136-y
2. El Ouadrhiri A, Abdelhadi A. Differential privacy for deep and federated learning: A survey. IEEE access. 2022;10:22359-22380.
3. Al Asqah M, Moulahi T. Federated Learning and Blockchain Integration for Privacy Protection in the Internet of Things: Challenges and Solutions. Future Internet. 2023;15(6):203. doi:10.3390/fi15060203
4. Xames MD, Topcu TG. A Systematic Literature Review of Digital Twin Research for Healthcare Systems: Research Trends, Gaps, and Realization Challenges. IEEE Access. Published online 2024. Accessed January 7, 2025. https://ieeexplore.ieee.org/abstract/document/10379674/

5. Gu B, Xu A, Huo Z, Deng C, Huang H. Privacy-Preserving Asynchronous Federated Learning Algorithms for Multi-Party Vertically Collaborative Learning. Published online 2020. doi:10.48550/ARXIV.2008.06233
6. Dwork C. Differential Privacy. In: Bugliesi M, Preneel B, Sassone V, Wegener I, eds. Automata, Languages and Programming. Springer; 2006:1-12. doi:10.1007/11787006_1
7. Acar A, Aksu H, Uluagac AS, Conti M. A Survey on Homomorphic Encryption Schemes: Theory and Implementation. ACM Comput Surv. 2019;51(4):1-35. doi:10.1145/3214303
8. Aziz R, Banerjee S, Bouzefrane S, Le Vinh T. Exploring Homomorphic Encryption and Differential Privacy Techniques towards Secure Federated Learning Paradigm. Future Internet. 2023;15(9):310. doi:10.3390/fi15090310
9. Banabilah S, Aloqaily M, Alsayed E, Malik N, Jararweh Y. Federated learning review: Fundamentals, enabling technologies, and future applications. Information processing & management. 2022;59(6):103061.
10. Hasan J. Security and Privacy Issues of Federated Learning. Published online 2023. doi:10.48550/ARXIV.2307.12181
11. Ibrahim Khalaf O, S.R A, Algburi S, et al. Federated learning with hybrid differential privacy for secure and reliable CROSS-IOT platform knowledge sharing. Security and Privacy. 2024;7(3):e374. doi:10.1002/spy2.374
12. Liu Z, Li Z. A blockchain-based framework of cross-border e-commerce supply chain. International journal of information management. 2020;52:102059.
13. Shokri R, Shmatikov V. Privacy-Preserving Deep Learning. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security. ACM; 2015:1310-1321. doi:10.1145/2810103.2813687
14. Issa W, Moustafa N, Turnbull B, Sohrabi N, Tari Z. Blockchain-Based Federated Learning for Securing Internet of Things: A Comprehensive Survey. ACM Comput Surv. 2023;55(9):1-43. doi:10.1145/3560816
15. Ahmed AA, Alabi O. Secure and scalable blockchain-based federated learning for cryptocurrency fraud detection: A systematic review. IEEE Access. Published online 2024. Accessed January 9, 2025. https://ieeexplore.ieee.org/abstract/document/10599372/
16. Gu X, Sabrina F, Fan Z, Sohail S. A Review of Privacy Enhancement Methods for Federated Learning in Healthcare Systems. IJERPH. 2023;20(15):6539. doi:10.3390/ijerph20156539
17. Mora A, Bujari A, Bellavista P. Enhancing generalization in federated learning with heterogeneous data: A comparative literature review. Future Generation Computer Systems. Published online 2024. Accessed January 9, 2025. https://www.sciencedirect.com/science/article/pii/S0167739X24001031
18. Monir IA, Fauzan MI, Ghinita G. DP-SGD. Data Engineering. Published online 2023:94.
19. Zhu J, Cao J, Saxena D, Jiang S, Ferradi H. Blockchain-empowered Federated Learning: Challenges, Solutions, and Future Directions. ACM Comput Surv. 2023;55(11):1-31. doi:10.1145/3570953
20. Kharbanda NS. AI/ML FOR DATA PRIVACY AND ENCRYPTION IN CLOUD COMPUTING. INTERNATIONAL JOURNAL OF RESEARCH IN COMPUTER APPLICATIONS AND INFORMATION TECHNOLOGY (IJRCAIT). 2024;7(2):27-43.
21. López-Blanco R, Alonso RS, González-Arrieta A, Chamoso P, Prieto J. Federated Learning of Explainable Artificial Intelligence (FED-XAI): A Review. In: Ossowski S, Sitek P, Analide C, Marreiros G, Chamoso P, Rodríguez S, eds. Distributed Computing and Artificial Intelligence, 20th International Conference. Vol 740. Lecture Notes in Networks and Systems. Springer Nature Switzerland; 2023:318-326. doi:10.1007/978-3-031-38333-5_32.