

Cross-Cloud Data Privacy Protection: Optimizing Collaborative Mechanisms of AI Systems by Integrating Federated Learning and LLMs

Huaiying Luo
College of Computing and Information Science
Cornell University
New York, USA
hl2446@cornell.edu

Cheng Ji
Siebel School of Computing and Data Science
University of Illinois Urbana-Champaign
Champaign, Illinois, USA
chengji5@illinois.edu

Abstract—In the age of cloud computing, data privacy protection has become a major challenge, especially when sharing sensitive data across cloud environments. However, how to optimize collaboration across cloud environments remains an unresolved problem. In this paper, we combine federated learning with large-scale language models to optimize the collaborative mechanism of AI systems. Based on the existing federated learning framework, we introduce a cross-cloud architecture in which federated learning works by aggregating model updates from decentralized nodes without exposing the original data. At the same time, combined with large-scale language models, its powerful context and semantic understanding capabilities are used to improve model training efficiency and decision-making ability. We’ve further innovated by introducing a secure communication layer to ensure the privacy and integrity of model updates and training data. The model enables continuous model adaptation and fine-tuning across different cloud environments while protecting sensitive data. Experimental results show that the proposed method is significantly better than the traditional federated learning model in terms of accuracy, convergence speed and data privacy protection.

Index Terms—Cross-cloud environment, Data privacy protection, Federated learning, Large-scale language models

I. INTRODUCTION

In the field of cross-cloud data privacy protection, the combination of Federated Learning (FL) and Large Language Models (LLMs) has gradually become a hot topic. As data privacy and security concerns become more and more important, traditional centralized data storage and processing methods are increasingly being questioned. Traditional AI training methods often rely on centralizing all data stored on a single server or data center for processing [1].

This centralized data processing model is prone to the risk of data leakage and privacy violations, especially in the face of sensitive data, such as medical and health records, financial information and personal privacy data, and the centralized storage and sharing method can be attacked or misused during data processing [2].

Therefore, how to carry out efficient AI training under the premise of ensuring data privacy and security has become the focus of current research. In order to solve this problem,

federated learning emerged as an emerging distributed learning method, which effectively avoids the risk of data leakage and protects the privacy of data by training the model on the local device and sharing only model updates (such as weights, gradients, etc.) instead of the original data [3].

The key benefits of federated learning are that it reduces data transfer through local computation, ensuring the privacy of each data source while still maintaining the accuracy and performance of the model. Especially in a distributed environment, federated learning can not only conduct data analysis without violating privacy protection regulations, but also has strong scalability and flexibility, and can be widely used in multi-party collaboration scenarios [4, 5].

Notwithstanding the burgeoning advancements in large-scale language models (LLMs) across diverse domains [6–13], particularly within the realms of natural language processing and generative tasks [14–18], a notable caveat exists: the training of these intricate models necessitates substantial computational resources and considerable data processing capabilities. Single-cloud platforms often face computing resource bottlenecks and latency issues, which cannot meet the high-performance requirements of LLMs training. In this context, cross-cloud federated training has become a new solution [19].

Cross-cloud federated training coordinates the computing resources of multiple cloud platforms and makes full use of the advantages of different cloud environments to complete large-scale model training tasks collaboratively. This cross-cloud architecture can distribute data and computing resources across multiple cloud platforms, effectively reducing the computing pressure on a single platform while improving resource utilization and training efficiency.

The advantageous resources of different cloud platforms can be fully integrated, and the latency and bottlenecks caused by insufficient resources on a single platform can be reduced [20], thereby accelerating the process of model training. In addition, cross-cloud federated training not only reduces the cost of training by sharing computing load among multiple platforms, but also avoids excessive centralized processing in the data set, further improving the level of data privacy protection [21].

Compared with a single cloud platform, cross-cloud training can more flexibly adjust resource allocation, further improve training efficiency, and reduce overall costs while ensuring data security.

With cross-cloud training, data privacy protection and efficient computing are guaranteed, especially when dealing with large-scale models, which can effectively solve the problem of computing bottlenecks and latency. The system using cross-cloud architecture can not only improve the collaboration ability of large-scale language models (LLMs) between different cloud platforms.

However, the training process of AI systems [22–27], promotes the research and application of cross-cloud data privacy protection technology based on federated learning. With the continuous improvement of cloud computing infrastructure and the maturity of federated learning frameworks, cross-cloud federated training is expected to become one of the mainstream methods for AI model training.

II. RELATED WORK

Lin et al. [28] point out that although LLMs excel in tasks such as code understanding, high-quality code data often has commercial or sensitive value, limiting its availability in open source AI projects. To solve this problem, the authors propose a governance framework with federated learning as the core, which aims to promote the joint development and maintenance of open-source AI code models while ensuring data privacy and security.

Lazaros et al. [29] emphasize that federated learning not only enables multi-party collaboration while preserving privacy, but also transforms Internet of Things (IoT) systems into more collaborative, privacy-preserving and flexible frameworks. The study provides insight into understanding the role of federated learning in collaborative intelligence.

Yao et al. [30] discussed the challenges of fine-tuning and cued learning in federated settings, analyzed the challenges of model convergence and high communication costs caused by data heterogeneity, and proposed potential directions for future research.

Tao et al. [31] proposed the FLFT approach, which aims to address the challenges of fine-tuning large-scale pre-trained language models in a distributed environment. FLFT incorporates federated learning (FL) strategies that allow multiple participants to share model update information while protecting data privacy, enabling collaborative fine-tuning.

Mawela et al. [32] proposed a web-based federated learning (FL) automation solution that aims to simplify the deployment and management of FL tasks. The system provides a user-friendly web interface with support for FedAvg algorithms, allowing users to configure parameters for FL tasks through an intuitive interface, reducing reliance on programming and network architecture.

Vadisetty et al. [33] proposed an AI-generated privacy protection protocol for cross-cloud data sharing and collaboration. The core innovation of this research lies in the design of a set of protocol frameworks that combine federated learning, differential privacy, dynamic encryption, and context-aware policies, aiming to improve the privacy and security of data collaboration in a multi-cloud environment.

Yang et al. [34] proposed a novel cross-cloud data privacy protection framework that aims to address data privacy and security issues when training large-scale language models (LLMs) in multi-cloud environments. The framework uses a federated learning (FL) approach that combines homomorphic encryption, dynamic model aggregation techniques, and cross-cloud data orchestration solutions to enhance security, efficiency, and scalability.

III. METHODOLOGIES

A. Federated Learning and Large-scale language models

First, let’s review the traditional federated learning architecture. In standard federated learning, multiple edge nodes, such as compute nodes in different cloud environments, train the model locally and then aggregate their updated information (rather than raw data) to a central server. Set the local loss function for each node to Equation 1:

$$\mathcal{L}_i(w_i) = \frac{1}{N_i} \sum_{n=1}^{N_i} \ell(f(x_{i,n}; w_i), y_{i,n}), \quad (1)$$

where w_i is the model parameter of the i -th node, $x_{i,n}$ and $y_{i,n}$ are the features and labels of the n -th sample on node i , respectively, and $f(\cdot)$ is the model prediction function, $\ell(\cdot)$ is the loss function, and N_i is the number of local samples of node i .

The goal of federated learning is to optimize the global model by aggregating local updates, and the global update rule is Equation (2):

$$w^{(t+1)} = \sum_{i=1}^K \frac{N_i}{N} w_i^{(t)}, \quad (2)$$

where $w_i^{(t)}$ is the model parameter of the i -th node after the t -th round of training, N is the total number of samples of all nodes, and K is the number of nodes participating in the training.

The formula aggregates the model updates of each node through the weighted average, so as to obtain the optimization of the global model. In this process, the original data is not directly transmitted or exposed, ensuring the privacy of the data.

In the existing federated learning framework, we combine large-scale language models (LLMs) to enhance the training efficiency and decision-making ability of the model.

Suppose we have a pre-trained model based on LLMs that is capable of extracting contextual and semantic feature information. The way we introduce LLMs in federated learning is by utilizing LLMs for feature augmentation when training a local model on each node. Set the output of LLMs to Equation 3:

$$z_{i,n} = LLM(x_{i,n}), \quad (1)$$

where $z_{i,n}$ is the context feature of the LLM output. These features are then fed into the local model of node i for training, which in turn updates the local parameter w_i . Specifically, the local training objective function is Equation (4):

$$\mathcal{L}_i(w_i) = \frac{1}{N_i} \sum_{n=1}^{N_i} \ell(f(z_{i,n}; w_i), y_{i,n}). \quad (4)$$

By introducing LLMs, we are not only able to leverage the local data of each node for model training, but also improve the model's ability to understand the context and semantics of the data, thereby enhancing the performance of the global model. The addition of LLM enables the model to better capture the deep relationships in the data when dealing with complex tasks.

B. Secure communication layer

To further enhance privacy, we introduce a secure communication layer into the model that encrypts model updates and training data to ensure that sensitive information is not exposed during federated learning.

Specifically, we encrypt the model update information of each node through Homomorphic Encryption. Suppose the cryptographic update information of node i is Equation (5):

$$\hat{w}_i = \text{Enc}(w_i), \quad (5)$$

where $\text{Enc}(\cdot)$ indicates an encryption operation.

The encrypted model update information will be aggregated on the server side, and the global model will be updated to Equation (6) through the decryption operation:

$$w^{(t+1)} = \text{Dec} \left(\sum_{i=1}^K \frac{N_i}{N} \hat{w}_i \right). \quad (6)$$

By introducing homomorphic encryption, we ensure that the data itself is encrypted even during the transmission of model updates, preventing malicious nodes or attackers from accessing the model parameters, thus effectively protecting the user's data privacy.

Finally, we adapted and fine-tuned the model across clouds. Whenever a model is migrated from one cloud node to another, we retrain the model through an adaptation layer to adapt to the data characteristics in the new environment. Set the migrated model to Equation (7):

$$w' = w + \Delta w, \quad (7)$$

where w' is the fine-tuned model parameter, and Δw is the parameter update obtained during the fine-tuning process.

By fine-tuning across clouds, we are able to ensure that the migration of models between different cloud platforms does not result in performance degradation, while maintaining the versatility and adaptability of the models.

IV. EXPERIMENTS

A. Experimental setup

The experiment uses the Google Cloud BigQuery dataset, which contains diverse data from real-world business scenarios, including financial, medical, social media, and geographic information, which has high practical application value. The dataset includes structured and semi-structured data that mimics multiple data formats in the real world, and some of the data involves sensitive information, such as personally identifiable information and transaction records, making it suitable for studying how to protect data privacy across cloud environments.

We will use the following four advanced comparison methods including:

- FedAvg (Federated Averaging): FedAvg is a classic federated learning method that aggregates the updated average of the local model into a global model. It is one of the most commonly used benchmarking methods in federated learning.
- DP-FL (Differential Privacy Federated Learning): DP-FL combines differential privacy technology with federated learning to add noise to the transmitted model updates during model training, thereby further improving data privacy protection capabilities.
- SMC-FL (Secure Multi-Party Computation Federated Learning): SMC-FL combines federated learning and secure multi-party computation (SMC) technology to ensure the security of model updates through cryptographic computing without directly exposing data to participants in the collaborative learning process.
- HE-FL (Homomorphic Encryption Federated Learning): HE-FL uses homomorphic encryption technology to ensure that data is computed and model trained in an encrypted state. This approach protects the privacy of the data, especially when dealing with highly sensitive data, such as medical or financial data.

B. Experimental analysis

Figure 1 illustrates the privacy protection effect of different approaches under different privacy budgets (epsilon). As can be seen in Figure 1, the privacy protection effect of all methods improves as epsilon increases. FedAvg has shown stable privacy protection effect, and gradually improves with the increase of epsilon. DP-FL performs well in terms of privacy protection, but it fluctuates slightly compared to

FedAvg, indicating that there is a certain instability between privacy and accuracy in the differential privacy mechanism.

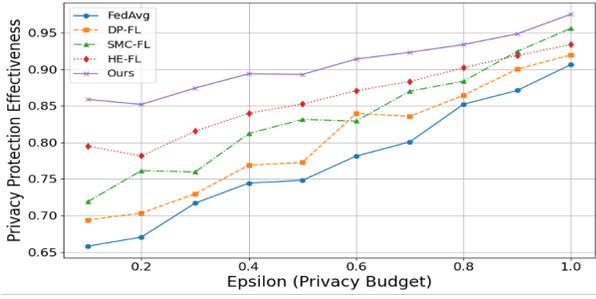


Fig. 1: Differential Privacy Effectiveness Comparison for Different Methods

The curve of SMC-FL is relatively smooth, indicating that secure multi-party computation provides a relatively balanced privacy protection effect. Although HE-FL provides a high privacy protection effect, it is highly volatile, which reflects that homomorphic encryption may affect the performance and stability of the model while improving privacy protection. Our 'Ours' method performed best of all methods, and its privacy protection improved rapidly with the addition of epsilon, showing that our method has a good balance between privacy protection and model accuracy.

As can be seen in Figure 2, the convergence speed (the number of iterations required to reach 85% accuracy) generally decreases for all methods as the number of hidden cells increases. The 'Ours' method has the best convergence speed under all hidden unit configurations, and the number of iterations required is significantly less than that of other methods, indicating that it performs well in training efficiency. In contrast, FedAvg and HE-FL require more training cycles, especially in larger hidden unit configurations, showing slower convergence rates.

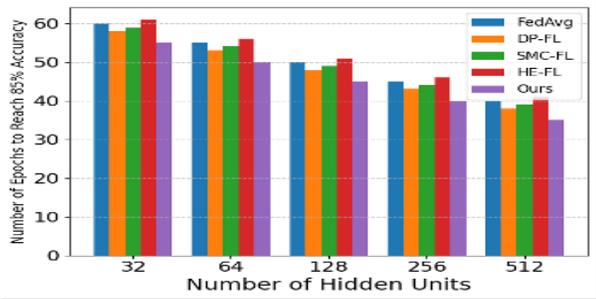


Fig. 2: Convergence Speed Comparison for Different Methods

As can be seen from Figure 3, as the learning rate gradually increases from 0.001 to 0.05, the training time of most methods decreases significantly, indicating that a higher learning rate helps to accelerate model convergence.

In contrast, FedAvg and DP-FL have a longer training time when the learning rate is low, but the convergence speed is

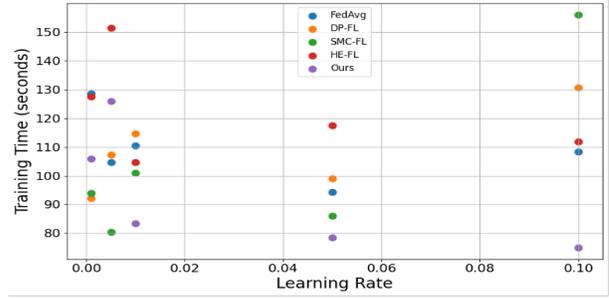


Fig. 3: Training Time Comparison for Different Methods

significantly improved after a moderate increase in the learning rate. SMC-FL and HE-FL were more sensitive to changes in learning rate, and when the learning rate was too high (0.1), the training time increased, reflecting the instability of training.

V. CONCLUSION

In conclusion, by combining federated learning and large-scale language models, the cross-cloud collaborative training framework proposed in this study achieves the highest privacy protection effect in the differential privacy test, with the least iteration rounds, the fastest convergence, and the shortest and most stable training time under different learning rates, which is better than existing methods. As for the future, it can focus on dynamic privacy budget scheduling, adaptive communication compression strategies, and deep integration with trusted execution environment and hardware security module.

REFERENCES

- [1] D. K. Seth, K. K. Ratra, and A. P. Sundareswaran, "AI and generative AI-driven automation for multi-cloud and hybrid cloud architectures: Enhancing security, performance, and operational efficiency," in *2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC)*. IEEE, 2025, pp. 784–793.
- [2] Z. Zhang and B. Liu, "Research on key technologies for cross-cloud federated training of large language models," *Academic Journal of Computing & Information Science*, vol. 7, no. 11, pp. 42–49, 2024.
- [3] F. M. Rasel and B. Peter, "AI-driven frameworks for enhancing cybersecurity in multi-cloud environments," *International Journal of Advanced Engineering Technologies and Innovations*, vol. 1, no. 1, pp. 24–32, 2025.
- [4] R. Gafni, I. Aviv, and D. Haim, "Multi-party secured collaboration architecture from cloud to edge," *Journal of Computer Information Systems*, vol. 64, no. 5, pp. 698–709, 2024.
- [5] Z. Li, S. He, Z. Yang, M. Ryu, K. Kim, and R. Madduri, "Advances in appfl: A comprehensive and extensible federated learning framework," *arXiv preprint arXiv:2409.11585*, 2024.
- [6] Z. Ding, P. Li, Q. Yang, and S. Li, "Enhance image-to-image generation with llava-generated prompts," in *2024 5th International Conference on Information Science, Parallel and Distributed Systems (ISPDS)*. IEEE, 2024, pp. 77–81.
- [7] Q. Deng, Q. Yang, R. Yuan, Y. Huang, Y. Wang, X. Liu, Z. Tian, J. Pan, G. Zhang, H. Lin *et al.*, "Composerx: Multi-agent symbolic music composition with llms," *arXiv preprint arXiv:2404.18081*, 2024.
- [8] Y. Ji, Z. Li, R. Meng, S. Sivarajkumar, Y. Wang, Z. Yu, H. Ji, Y. Han, H. Zeng, and D. He, "RAG-RLRC-LaySum at BioLaySumm: Integrating retrieval-augmented generation and readability control for layman summarization of biomedical texts," in *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, D. Demner-Fushman, S. Ananiadou, M. Miwa,

- K. Roberts, and J. Tsujii, Eds. Bangkok, Thailand: Association for Computational Linguistics, Aug. 2024, pp. 810–817. [Online]. Available: <https://aclanthology.org/2024.bionlp-1.75/>
- [9] Y. Ji, W. Ma, S. Sivarajkumar, H. Zhang, E. M. Sadhu, Z. Li, X. Wu, S. Visweswaran, and Y. Wang, “Mitigating the risk of health inequity exacerbated by large language models,” *npj Digital Medicine*, vol. 8, no. 1, p. 246, 2025. [Online]. Available: <https://doi.org/10.1038/s41746-025-01576-4>
 - [10] Q. Yi, Y. He, J. Wang, X. Song, S. Qian, M. Zhang *et al.*, “SCORE: Story coherence and retrieval enhancement for AI narratives,” *arXiv preprint arXiv:2503.23512*, 2025.
 - [11] Y. Jin, Z. Yang, X. Xu, Y. Zhang, and S. Ji, “Adaptive fault tolerance mechanisms of large language models in cloud computing environments,” *arXiv preprint arXiv:2503.12228*, 2025.
 - [12] H. Xu, X. Wang, and H. Chen, “Towards real-time and personalized code generation,” in *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*, 2024, p. 5568–5569.
 - [13] Z. Yang, Y. Jin, J. Liu, X. Xu, Y. Zhang, and S. Ji, “Research on cloud platform network traffic monitoring and anomaly detection system based on large language models,” *arXiv preprint arXiv:2504.17807*, 2025.
 - [14] Y. Ji, Z. Yu, and Y. Wang, “Assertion detection in clinical natural language processing using large language models,” in *2024 IEEE 12th International Conference on Healthcare Informatics (ICHI)*, 2024, pp. 242–247.
 - [15] J. He, C. I. Kanatsoulis, and A. Ribeiro, “T-GAE: Transferable graph autoencoder for network alignment,” *arXiv e-prints*, pp. arXiv–2310, 2023.
 - [16] J. He, M. D. Ma, J. Fan, D. Roth, W. Wang, and A. Ribeiro, “GIVE: Structured reasoning with knowledge graph inspired veracity extrapolation,” *arXiv preprint arXiv:2410.08475*, 2024.
 - [17] Z. Yang, Y. Jin, and X. Xu, “Hades: Hardware accelerated decoding for efficient speculation in large language models,” *arXiv preprint arXiv:2412.19925*, 2024.
 - [18] D. Liu and Y. Yu, “MT2ST: Adaptive multi-task to single-task learning,” *arXiv preprint arXiv:2406.18038*, 2024. [Online]. Available: <https://arxiv.org/abs/2406.18038>
 - [19] Y. Ramaswamy, V. N. Sankaran, and B. K. M. Sundar, “Advanced cybersecurity strategies in cloud computing: Techniques for data protection and privacy,” *Library Progress International*, vol. 44, no. 3, pp. 2643–2656, 2024.
 - [20] Y. Jin and Z. Yang, “Scalability optimization in cloud-based AI inference services: Strategies for real-time load balancing and automated scaling,” *arXiv preprint arXiv:2504.15296*, 2025.
 - [21] Y. Guo, “Optimization of privacy-aware cloud crowdsourcing resource combinations for product development,” *Expert Systems with Applications*, vol. 227, p. 120176, 2023.
 - [22] J. Bosch, H. H. Olsson, and I. Crnkovic, “Engineering AI systems: A research agenda,” *Artificial intelligence paradigms for smart cyber-physical systems*, pp. 1–19, 2021.
 - [23] P. Li, M. Abouelenien, R. Mihalcea, Z. Ding, Q. Yang, and Y. Zhou, “Deception detection from linguistic and physiological data streams using bimodal convolutional neural networks,” in *2024 5th International Conference on Information Science, Parallel and Distributed Systems (ISPDS)*. IEEE, 2024, pp. 263–267.
 - [24] P. Li, Q. Yang, X. Geng, W. Zhou, Z. Ding, and Y. Nian, “Exploring diverse methods in visual question answering,” in *2024 5th International Conference on Electronic Communication and Artificial Intelligence (ICECAI)*. IEEE, 2024, pp. 681–685.
 - [25] Z. Ding, Z. Lai, S. Li, P. Li, Q. Yang, and E. Wong, “Confidence trigger detection: Accelerating real-time tracking-by-detection systems,” in *2024 5th International Conference on Electronic Communication and Artificial Intelligence (ICECAI)*. IEEE, 2024, pp. 587–592.
 - [26] T. Wang, Q. Yang, R. Wang, D. Sun, J. Li, Y. Chen, Y. Hu, C. Yang, T. Kimura, D. Kara *et al.*, “Fine-grained control of generative data augmentation in IoT sensing,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 32787–32812, 2024.
 - [27] Q. Yang, C. Ji, H. Luo, P. Li, and Z. Ding, “Data augmentation through random style replacement,” *arXiv preprint arXiv:2504.10563*, 2025.
 - [28] Z. Lin, W. Ma, T. Lin, Y. Zheng, J. Ge, J. Wang, J. Klein, T. Bissyande, Y. Liu, and L. Li, “Open-source AI-based se tools: Opportunities and challenges of collaborative software learning,” *ACM Transactions on Software Engineering and Methodology*, 2024, just Accepted.
 - [29] K. Lazaros, D. E. Koumadorakis, A. G. Vrahatis, and S. Kotsiantis, “Federated learning: Navigating the landscape of collaborative intelligence,” *Electronics*, vol. 13, no. 23, p. 4744, 2024.
 - [30] Y. Yao, J. Zhang, J. Wu, C. Huang, Y. Xia, T. Yu, R. Zhang, S. Kim, R. Rossi, A. Li, L. Yao, J. McAuley, Y. Chen, and C. Joe-Wong, “Federated large language models: Current progress and future directions,” *arXiv preprint arXiv:2409.15723*, 2024. [Online]. Available: <https://arxiv.org/abs/2409.15723>
 - [31] Y. Tao and A. Authors], “Flft: A large-scale pre-training model distributed fine-tuning method that integrates federated learning strategies,” *IEEE Access*, 2025, to appear.
 - [32] C. Mawela, C. B. Issaid, and M. Bennis, “A web-based solution for federated learning with LLM-based automation,” *arXiv preprint arXiv:2408.13010*, 2024. [Online]. Available: <https://arxiv.org/abs/2408.13010>
 - [33] R. Vadisetty and A. Polamarasetti, “AI-generated privacy-preserving protocols for cross-cloud data sharing and collaboration,” in *2024 IEEE 4th International Conference on ICT in Business Industry & Government (ICTBIG)*. IEEE, 2024, pp. 1–5.
 - [34] Z. Yang, Y. Jin, Y. Zhang, J. Liu, and X. Xu, “Research on large language model cross-cloud privacy protection and collaborative training based on federated learning,” *arXiv preprint arXiv:2503.12226*, 2025.