Towards a Meaningful Communication and Model Aggregation in Federated Learning via Genetic Programming

Elia Pacioni[1,2][a], Francisco Fernández De Vega[2][b] and Davide Calvaresi[1][c]
[1] *University of Applied Sciences and Arts of Western Switzerland (HES-SO Valais/Wallis), Rue de l'Industrie 23, Sion, 1950, Switzerland*
[2] *Universidad de Extremadura, Av. Santa Teresa de Jornet, 38, Mérida, 06800, Spain*
{elia.pacioni, davide.calvaresi}@hevs.ch, fcofdez@unex.es

Keywords: Federated Learning, Multi-Agents System, Models Aggregation, Communication Efficiency, Genetic Programming.

Abstract: Federated Learning (FL) enables collaborative training of machine learning models while preserving client data privacy. However, its conventional client-server paradigm presents two key challenges: (i) communication efficiency and (ii) model aggregation optimization. Inefficient communication, often caused by transmitting low-impact updates, results in unnecessary overhead, particularly in bandwidth-constrained environments such as wireless or mobile networks or in scenarios with numerous clients. Furthermore, traditional aggregation strategies lack the adaptability required for stable convergence and optimal performance. This paper emphasizes the distributed nature of FL clients (agents) and advocates for local, autonomous, and intelligent strategies to evaluate the significance of their updates—such as using a "distance" metric relative to the global model. This approach improves communication efficiency by prioritizing impactful updates. Additionally, the paper proposes an adaptive aggregation method leveraging genetic programming and transfer learning to dynamically evolve aggregation equations, optimizing the convergence process. By integrating insights from multi-agent systems, the proposed approach aims to foster more efficient and robust frameworks for decentralized learning.

SCIENCE AND TECHNOLOGY PUBLICATIONS

1 INTRODUCTION

Federated Learning (FL) is a collaborative machine learning paradigm introduced by Google in 2016 (McMahan et al., 2017). FL is designed to train models on decentralized data while preserving privacy. Unlike traditional centralized learning, which requires transferring data to a central server, FL enables training to occur locally on client devices, addressing critical concerns about data confidentiality (Kairouz and et al., 2021). By maintaining data on client devices, FL mitigates privacy risks while facilitating the training of large-scale machine learning models. This paradigm has been significantly adopted in applications including predictive text input (e.g., GBoard (Hard et al., 2019)), speech recognition (e.g., Siri (Granqvist et al., 2020)), healthcare diagnostics (Rieke et al., 2020), and finance (Liu et al., 2020; Long et al., 2020).

[a] https://orcid.org/0000-0002-1557-4870
[b] https://orcid.org/0000-0002-1086-1483
[c] https://orcid.org/0000-0001-9816-7439

FL faces key challenges that limit its scalability and practical adoption. A major issue arises from the heterogeneity of client devices and their non-Independent and Identically Distributed (non-IID) data distributions. This heterogeneity leads to variable update quality (Li et al., 2020; Nie et al., 2022). Traditional aggregation methods such as Federated Averaging (FedAVG) (McMahan et al., 2017), which rely on weighted averaging, struggle under these conditions, impairing global model generalization, convergence speed, and overall performance.

Communication inefficiency is another critical limitation. The standard FL pipeline transmits all client updates to a central server indiscriminately, leading to excessive bandwidth usage, especially in large-scale or resource-constrained environments such as IoT networks (Kontar et al., 2021). Techniques like parameter compression (Konečný et al., 2017), federated dropout (Bouacida et al., 2021), and structured updates (Zhang et al., 2024; Konečný et al., 2017; Wang et al., 2024) offer partial solutions but often fail to adapt effectively to the dynamic and distributed nature of FL.

Pacioni, E., Fernández De Vega, F. and Calvaresi, D.
Towards a Meaningful Communication and Model Aggregation in Federated Learning via Genetic Programming.
DOI: 10.5220/0013380400003890
Paper published under CC license (CC BY-NC-ND 4.0)
In *Proceedings of the 17th International Conference on Agents and Artificial Intelligence (ICAART 2025) - Volume 3*, pages 1427-1431
ISBN: 978-989-758-737-5; ISSN: 2184-433X
Proceedings Copyright © 2025 by SCITEPRESS – Science and Technology Publications, Lda.

To address these limitations, this work proposes a novel paradigm that integrates FL with Multi-Agent Systems (MAS) modeling equipped with an adaptive Genetic Programming (GP)-based aggregation strategy.

GP has been introduced by Koza et al. (Koza, 1992) and is particularly effective in optimizing nonlinear functions and generating solutions dynamically. A prominent application of GP is in solving symbolic regression problems (Augusto and Barbosa, 2000), which involves creating mathematical models that accurately fit a given set of data points. In the context of FL, there is a similar need for a mathematical function—the aggregation function—that can effectively handle diverse data inputs. This resemblance to symbolic regression makes GP a powerful tool for evolving aggregation strategies that are customized to accommodate heterogeneous data distributions and varying client capabilities.

Therefore, we envision MAS empowering clients (agents) to decide when and whether to communicate with the central server, thereby optimizing bandwidth usage and aggregation delays. Moreover, the GP-based strategy dynamically adjusts aggregation equations to better accommodate diverse client data distributions. Combining MAS for autonomy and communication efficiency, and GP-based aggregation for robustness, the proposed framework aims to overcome key bottlenecks in FL, delivering scalable, resilient, and personalized learning systems.

The remainder of this paper is structured as follows: Section 2 outlines the motivations and challenges driving this work. Section 3 details the proposed paradigm, emphasizing the integration of MAS and GP-based aggregation. Section 4 presents a roadmap for implementation and future exploration. Finally, Section 5 discusses the potential impact and concludes the paper.

2 MOTIVATIONS

While FL presents a promising approach for distributed learning, its current implementations face two critical shortcomings: communication inefficiency and limited adaptability in aggregation strategies. These issues are particularly pronounced in non-IID scenarios, where treating clients uniformly often results in inefficiencies (Wang et al., 2020). Furthermore, transmitting all client updates indiscriminately imposes substantial communication costs, particularly in bandwidth-constrained environments. Client selection approaches handle the process server-side, after receiving the weights, this does not help to improve efficiency (Fu et al., 2023). Existing approaches to improve communication efficiency, including parameter compression (Konečný et al., 2017), federated dropout (Bouacida et al., 2021), and structured updates (Zhang et al., 2024; Konečný et al., 2017; Wang et al., 2024), partially mitigate these issues by reducing the volume of transmitted data. However, these methods do not empower clients to perform local selection of updates, resulting in the transmission of irrelevant data that is later discarded by the server—introducing unnecessary overhead. Similarly, static aggregation methods like FedAVG fail to account for the unique characteristics of client data, particularly in non-IID scenarios (Wang et al., 2020).

FedGR (Zeng et al., 2024) introduces an innovation in FL by using a genetic algorithm for dynamic client clustering and a relay strategy to train models within groups. This approach stands out for its ability to reduce the impact of statistical heterogeneity by improving the convergence of the overall model. However, FedGR is limited to a static aggregation strategy within each group and does not consider dynamic customization of aggregation algorithms.

An interesting analysis of aggregation techniques in Federated Learning has recently been produced (Qi et al., 2024) which highlights as major problems (i) statistical heterogeneity of data, (ii) communication bottlenecks, (iii) security and privacy, (iv) model customization, and (v) client evaluation and selection.

This work addresses some of these limitations by integrating MAS to enable autonomous decision-making at both the client and server levels. The local agents assess their "distance" from the global model to determine whether to transmit updates, while the server (aggregating agent) dynamically decides when to distribute aggregated models or skip aggregation rounds. Combined with a GP-based aggregation strategy, this approach can enhance communication efficiency and ensure adaptability to evolving data distributions, capturing the peculiarities of each client

3 NOVEL PARADIGM

The proposed paradigm integrates a MAS-driven communication framework with a GP-based adaptive aggregation strategy, effectively addressing the dual challenges of communication efficiency and aggregation adaptability. Figure 1 illustrates the enhanced FL pipeline, incorporating MAS for communication optimization and GP for dynamic aggregation.

MAS for Communication Efficiency. In the proposed framework, each client operates as an autonomous agent able to assess the relevance of its updates based on a defined distance metric w.r.t. the global model. Agents decide independently whether to transmit updates, thereby reducing unnecessary bandwidth consumption. Concurrently, the central server employs MAS principles to evaluate when to distribute updated global models, ensuring synchronization is both timely and efficient. This decentralized decision-making mechanism enhances scalability and efficiency, particularly in resource-constrained environments.

GP-Based Aggregation. Conventional aggregation techniques (i.e., FedAVG) rely on static equations that fail to account for the heterogeneity of client data, particularly in non-IID scenarios. The proposed GP-based aggregation method dynamically evolves aggregation equations to adapt to the client data distributions. This adaptive strategy balances local data characteristics with global model performance, yielding a more robust solution for non-IID environments. Furthermore, by incorporating transfer learning within the GP-based aggregation framework, the computational overhead of evolving new equations in each round is significantly minimized.

By combining these innovations, the proposed paradigm offers a scalable solution for FL that maintains high performance and supports personalized models, even in complex and heterogeneous data environments.

4 ROAD MAP

This section outlines a roadmap to address key challenges in FL and validate the integration of MAS and GP for scalable, efficient, and decentralized learning.

Phase 1: Developing the GP-Based Aggregation Method – This phase focuses on creating a GP-based technique to evolve aggregation equations tailored to heterogeneous environments dynamically. A key aspect is defining the fitness function. Initially, a single-objective function will be employed, leveraging metrics such as accuracy, precision, or mean Average Precision (mAP), depending on the neural network type. Subsequently, a multi-objective function will be introduced, incorporating factors like convergence speed, model execution time, and model quality to optimize performance comprehensively.

To validate the proposed new aggregation method, a systematic comparison will be made with major aggregation methods in the literature, including FedAvg, Scaffold, MOON, Zeno, Per-FedAvg, FedProx, FedOpt, FedRS, and FedGR. Moreover, a test can be carried out with FedGR in combination with a dynamic aggregation algorithm and compared with the current results. Comparisons will be based on key metrics such as global model accuracy, computational efficiency, robustness to non-IID data, and resilience to malicious clients. This approach will highlight the advantages and limitations of each method, providing a basis for establishing practical guidelines on using different aggregation strategies in specific scenarios.

Phase 2: Transfer Learning for GP Aggregator – Although GP could have a significant computational impact, mainly due to the costs of the fitness function that they have to evaluate each individual (neural network) and to do so, they have to apply the inference process and compute metrics; we believe that the benefits of its ability to dynamically adapt to non-IID features in the data outweigh these costs. However, to reduce computational overhead, transfer learning mechanisms will be integrated into the GP-based aggregation process. This approach involves two strategies: (i) performing a single evolution in the first aggregation round and reusing the same expression tree across subsequent iterations, thereby maximizing resource efficiency, and (ii) transferring the best individual from the previous evolution to form the new population. The latter strategy offers two scenarios: either the transferred individual is supplemented by randomly generated individuals, or the entire population is derived through mutations of the selected individual.

Phase 3: MAS for Communication Efficiency – MAS-based mechanisms will enable autonomous decision-making for clients and servers, optimizing communication and synchronization at run time. This phase also opens avenues for further investigation: (i) refining MAS to manage large-scale FL deployments with thousands of clients and (ii) exploring the integration of evolutionary methods within MAS to customize the distance threshold dynamically.

5 DISCUSSION AND CONCLUSIONS

This paper introduces a novel paradigm for FL, combining MAS-driven communication and aggregation through GP to address challenges in non-IID environments. Enabling autonomous decision-making and dynamic aggregation, we aim to reduce communica-

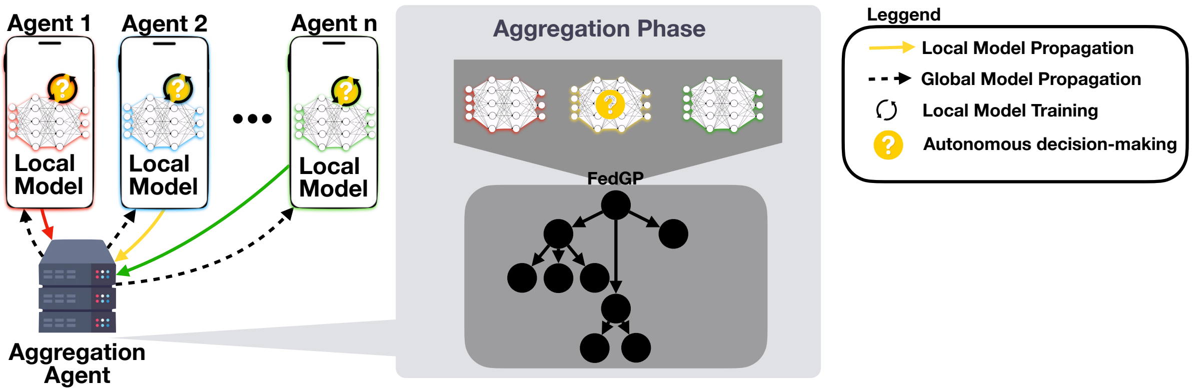

Figure 1: MAS integration (communication efficiency) and GP-based aggregation (model quality).

tion overhead, accelerate convergence, improve the generalization, and enhance personalization.

The impact of this work is threefold: **scientifically**, it advances FL with scalable and adaptable mechanisms; **practically**, it offers robust solutions for applications in healthcare, finance, and mobile systems; **socially**, it promotes improved privacy and fairness in model development.

As a position paper, this work highlights the value of hybridizing AI techniques to optimize FL paradigms, paving the way for higher-quality solutions, for greater optimization, and broader real-world applicability.

ACKNOWLEDGEMENTS

This work was partially supported by the HES-SO RCSO ISNet HARRISON grant (WP2), the Spanish Ministry of Economy and Competitiveness (PID2020-115570GB-C21, PID2023-147409NB-C22), funded by MCIN/AEI/10.13039/501100011033, and the Junta de Extremadura (GR15068).

REFERENCES

Augusto, D. A. and Barbosa, H. J. C. (2000). Symbolic regression via genetic programming. In *Proceedings of the Sixth Brazilian Symposium on Neural Networks*, volume 1, pages 173–178, Rio de Janeiro, Brazil.

Bouacida, N., Hou, J., Zang, H., and Liu, X. (2021). Adaptive federated dropout: Improving communication efficiency and generalization for federated learning. In *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 1–6.

Fu, L., Zhang, H., Gao, G., Zhang, M., and Liu, X. (2023). Client selection in federated learning: Principles, challenges, and opportunities. *IEEE Internet of Things Journal*, 10(24):21811–21819.

Granqvist, F., Seigel, M., van Dalen, R., Cahill, A., Shum, S., and Paulik, M. (2020). Improving on-device speaker verification using federated learning with privacy. *arXiv preprint*.

Hard, A., Rao, K., Mathews, R., Ramaswamy, S., Beaufays, F., Augenstein, S., Eichner, H., Kiddon, C., and Ramage, D. (2019). Federated learning for mobile keyboard prediction. *arXiv preprint*.

Kairouz, P. and et al. (2021). Advances and open problems in federated learning. *Now Foundations and Trends*.

Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., and Bacon, D. (2017). Federated learning: Strategies for improving communication efficiency.

Kontar, R., Shi, N., Yue, X., Chung, S., Byon, E., Chowdhury, M., Jin, J., Kontar, W., Masoud, N., Nouiehed, M., Okwudire, C. E., Raskutti, G., Saigal, R., Singh, K., and Ye, Z.-S. (2021). The internet of federated things (ioft). *IEEE Access*, 9:156071–156113.

Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press.

Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. (2020). On the convergence of fedavg on non-iid data. *arXiv preprint*.

Liu, Y., Ai, Z., Sun, S., Zhang, S., Liu, Z., and Yu, H. (2020). Fedcoin: A peer-to-peer payment system for federated learning. In Yang, Q., Fan, L., and Yu, H., editors, *Federated Learning: Privacy and Incentive*, pages 125–138, Cham. Springer International Publishing.

Long, G., Tan, Y., Jiang, J., and Zhang, C. (2020). Federated learning for open banking. In Yang, Q., Fan, L., and Yu, H., editors, *Federated Learning: Privacy and Incentive*, pages 240–254, Cham. Springer International Publishing.

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., and Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*.

Nie, W., Yu, L., and Jia, Z. (2022). Research on aggregation strategy of federated learning parameters under non-independent and identically distributed conditions. In *Proceedings of the 2022 4th International Conference on Applied Machine Learning (ICAML)*, pages 41–48.

Qi, P., Chiaro, D., Guzzo, A., Ianni, M., Fortino, G., and Piccialli, F. (2024). Model aggregation techniques in federated learning: A comprehensive survey. *Future Generation Computer Systems*, 150:272–293.

Rieke, N., Hancox, J., Li, W., Milletarì, F., Roth, H. R., Albarqouni, S., et al. (2020). The future of digital health with federated learning. *npj Digital Medicine*, 3(1):119.

Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. (2020). Tackling the objective inconsistency problem in heterogeneous federated optimization. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.

Wang, L., Li, J., Chen, W., Wu, Q., and Ding, M. (2024). Communication-efficient model aggregation with layer divergence feedback in federated learning. *IEEE Communications Letters*, 28(10):2293–2297.

Zeng, Y., Zhao, K., Yu, F., Zeng, B., Pang, Z., and Wang, L. (2024). Fedgr: Genetic algorithm and relay strategy based federated learning. In *Proceedings of the 27th International Conference on Computer Supported Cooperative Work in Design (CSCWD)*.

Zhang, W., Zhou, T., Lu, Q., Yuan, Y., Tolba, A., and Said, W. (2024). Fedsl: A communication-efficient federated learning with split layer aggregation. *IEEE Internet of Things Journal*, 11(9):15587–15601.

SCIENCE AND TECHNOLOGY PUBLICATIONS