Title : Trustworthy and Sustainable Federated Learning Based on Negotiation"

Supervisors : Genoveva Vargas, Nadia Bennani, Chirine Ghedira

Federated learning on long-COVID data enables hospitals and other healthcare institutions to collaboratively train machine learning models without sharing raw data, which remains stored locally. However, participants are often reluctant to share their locally trained models when other members of the federation are not fully trusted, or when the consortium fails to meet certain trustworthiness or robustness requirements.

To address this challenge, in (Ardagna et al., 2024; Bena et al., 2025), we proposed a trust-based model in which each participant can express its own trust conditions. At each training round, only the nodes that satisfy the mutual trust constraints are selected to contribute to the aggregation of the global model. Nodes with high trust levels contribute gradients with standard weighting, whereas lower-trust nodes are subject to tighter gradient clipping, stronger differential privacy noise, or are temporarily restricted to evaluation roles until re-verified. This dynamic selection process is managed through a negotiation-based protocol.

We have developed an initial prototype of this system. The objective of this project is to build upon this work to develop a production-grade platform that implements our negotiation-based trust management protocol for federated learning (FL), and to deliver a complete demonstrator using a long-COVID use case.

We will extend the trust negotiation mechanism by integrating a resource-dispatching module capable of handling node arrivals and departures during FL rounds. This module will enforce responsible computing criteria —such as data sovereignty and governance constraints—as well as sustainability objectives, including energy efficiency and carbon footprint considerations (Vargas-Solar et al., 2024; Vargas-Solar et al., 2025). The outcomes of the negotiation process will guide dispatching decisions, including node admission, assigned roles, model weighting, and privacy budget allocation. Additionally, the system will support graceful rebalancing under participant churn, ensuring that the federation operates fairly, reliably, and in an energy-aware manner at scale.

References

[1] Vargas-Solar, G., Zechinelli-Martini, J.-L., Ardagna, C. A., Bena, N., Bennani, N., Catania, B., Espinosa-Oviedo, J. A., & Ghedira-Guégan, C. (2025). Techno/Ecofeminism in action: Fair and responsible resource allocation for sustainable data science pipelines. In Proceedings of the Workshops of the EDBT/ICDT 2025.
[2] Ardagna, C. A., & Bena, N. (2024). Revisiting trust management in the data economy: A road map. IEEE Internet Computing, 28.
[3] Vargas-Solar, G., Bennani, N., Espinosa-Oviedo, J. A., Mauri, A., Zechinelli-Martini, J.-L., Catania, B., Ardagna, C. A., & Bena, N. (2024). Decolonizing federated learning: Designing fair and responsible resource allocation. In Proc. IEEE/ACS AICCSA.
[4] Bena, N., Vargas-Solar, G., Bennani, N., Grecchi, N., Ghedira-Guégan, C., & Ardagna, C. A. (2025). Trust negotiation in dynamic service-based applications. Preprint.