

Byzantine-Resilient Federated AI with Low Comm. Cost

Training AI models across distributed, sensitive datasets remains difficult: conventional federated learning often suffers from high communication load and vulnerability to malicious participants. This technology enables privacy-preserving, communication-efficient and robust model training and fine-tuning without sharing raw data, making it highly relevant for healthcare, industrial collaborations and private-data-based AI services.




Privacy-preserving distributed AI training via scalar-based updates

The method replaces transmission of full model gradients with only a small number of scalar update values generated from shared random directions. A central node robustly aggregates these compressed updates, enabling efficient federated training while mitigating disruption from Byzantine or adversarial clients.


- 01 Trains or fine-tunes shared AI models without exposing local datasets
- 02 Maintains convergence even with adversarial participants
- 03 Cuts communication needs to a few scalar updates per training step
- 04 Applicable to healthcare, industrial data spaces and private AI services

REFERENCES:

-  Egger et al., ES-FoMo III, 2025
-  de Sa Delgado Neto et al., arXiv:2406.14362



Compared to standard federated averaging, a similar model accuracy is achieved at several orders of magnitude lower communication cost. Here shown for logistic regression on MNIST. Source: de Sa Delgado Neto et al. 2024.



The method achieves a high accuracy for a wide range of attack scenarios, here shown for fine-tuning the RoBERTa-large LLM on the TREC data set for non-I.I.D. data and 25% Byzantine clients. Source: Egger et al. 2025.




CHALLENGE

Federated learning is attractive for privacy-sensitive settings, but current approaches often require transmitting high-dimensional gradients, creating major bandwidth and compute bottlenecks. At the same time, even a few malicious or faulty participants can impair convergence or derail training. These limitations restrict deployment in real-world multi-party environments with sensitive data and heterogeneous infrastructure.

INNOVATION

The invention enables distributed model training through highly compressed scalar-based updates instead of full gradient exchange. Shared random seeds allow clients and server to reconstruct update directions locally, drastically reducing communication overhead. Robust aggregation protects training against Byzantine participants, while preserving model performance. The approach has been validated on standard machine learning tasks and large language model fine-tuning, demonstrating broad applicability in privacy-sensitive, resource-constrained settings.

01 Basic principles observed · 02 Technology concept formulated · 03 Experimental proof of concept · 04 Technology validated in lab · 05 Technology validated in relevant environment · 06 Technology demonstrated in relevant environment · 07 System prototype demonstrated in operational environment · 08 System complete and qualified



An invention of Technical University of Munich



As one of the largest technology transfer offices in Germany, BayPAT evaluates and commercializes innovations from 28 research institutions. With our extensive experience and expertise in patenting and licensing technologies and our comprehensive range of services, we are your partner of choice. For more information visit www.baypat.de

Bayerische Patentallianz GmbH
Dr. Sindre Haugland
 shaugland@baypat.de
 +49 (0)89 5480177-17
 Prinzregentenstr. 52
 80538 München

