Study of selective Fully Homomorphic Encryption approaches for large scale Federated Learning

Supervisors: Nesrine Kaaniche and Akram Bendoukha

Context

Privacy protection is a central challenge in modern machine learning applications dealing with sensitive data (healthcare, finance, energy, etc.). Fully Homomorphic Encryption (FHE) offers a strong solution by enabling computations directly on encrypted data [1], ensuring that the server never sees the raw data. This paradigm has been applied to federated learning (FL) [5], allowing encrypted model updates to be aggregated securely. However, FHE introduces significant computational and communication overheads.

Indeed, each model update must be transformed into ciphertexts, which are several orders of magnitude larger than their plaintext representation. In addition, the aggregation over encrypted values requires costly homomorphic additions and, depending on the scheme, rotations or relinearization. Then, with modern deep learning models containing hundreds of millions or billions of parameters, a single round of FL would require transmitting hundreds or thousands of ciphertexts per client, making the communication overhead prohibitive.

This motivates research into *selective encryption* [2, 3, 4], where only a subset of parameters (or the most privacy-sensitive portions of the model update) are encrypted, while the rest are sent in plaintext. Selective encryption promises to find a balance between strong privacy guarantees and the practical feasibility of large-scale federated learning.

Objectives

This project will focus on understanding and modeling the privacy/efficiency trade-off of selective encryption for federated learning. It aims at:

- analyzing which parts of the update (layers, coordinates) carry the most private information, using metrics such as gradient norms, or privacy leakage scores.
- exploring different parameter selection policies: random masking, top-k sensitive, layer-wise encryption, and adversarially optimized masks, and evaluating their effectiveness.
- measuring how privacy leakage (via reconstruction or membership inference attacks) decreases as the fraction of encrypted parameters increases and identifying the optimal fraction that gives the best cost–privacy ratio.
- providing design guidelines
- evaluating the scalability on scenarios with various parameters (For instance FL for LLM foundation models) and multiple clients per round.

Deliverables:

- Visualization of the privacy–efficiency trade-off curve, showing how selective encryption improves scalability compared to full encryption.
- Practical guidelines for selective encryption in real-world FL deployments with very large models.

References :

1. Gentry, C. *Fully Homomorphic Encryption Using Ideal Lattices.* STOC 2009.
2. Jiang, Z. et al. *MaskCrypt: Privacy-preserving Federated Learning via Selective Homomorphic Encryption.* arXiv 2023.
3. Song, L. et al. *Selective Parameter Encryption for Efficient and Private Federated Learning.* NeurIPS Workshop 2023.
4. Li, J. et al. *SenseCrypt: Sensitivity-Aware Selective Homomorphic Encryption in Cross-Device Federated Learning.* IEEE TIFS 2024
5. Bonawitz, K. et al. *Practical Secure Aggregation for Privacy-Preserving Machine Learning.* CCS 2017.