

Cryptographic Sparsification for Privacy-Preserving Collaborative Learning in Zero-Trust Edge Architectures

Nagaraja. K.V.

Associate professor, Dept of Computer science, Government First Grade Women's College, Tumkur.

ABSTRACT

Although decentralized machine learning paradigms safeguard private user data by exchanging model updates rather than raw records, advanced gradient inversion techniques have shown that sensitive source inputs can still be reconstructed with high fidelity. Standard cryptographic countermeasures, such as Homomorphic Encryption (HE), prevent untrusted server nodes from accessing gradients directly, yet they introduce severe compute and transmission bottlenecks—frequently inflating payload dimensions by 10x to 50x. To mitigate this constraint, we present CryptoSparsity, a novel framework tailored for zero-trust edge environments. CryptoSparsity combines dynamic threshold-based gradient pruning with additive homomorphic operations. By evaluating the relative magnitude of model updates during each training step, the architecture eliminates redundant parameters (filtering up to 95% of non-essential weight variations) prior to encrypting the remaining residual vectors using an efficient Paillier cryptosystem. Evaluated across 150 physical edge nodes (comprising NVIDIA Jetson devices and Raspberry Pi 4 systems) training Deep Neural Networks, CryptoSparsity decreased total encrypted packet size by 88.4% and cut client-side encryption processing time by 82.1%, while maintaining robust defense against gradient-based data recovery attacks.

Keywords: Privacy Preservation, Additive Encryption, Gradient Pruning, Zero-Trust Architecture, Reconstruction Mitigation, Edge Intelligence.

1. Context & Motivation

Decentralized execution frameworks allow distributed edge endpoints to collaboratively train shared models without centralizing proprietary datasets. However, recent vulnerability findings indicate that transmitting unencrypted gradient vectors leaves the system susceptible to severe security risks.

Two primary privacy and operational challenges affect existing edge deployments:

- 1. Gradient Inversion Risks:** Adversaries can analyze exposed gradient updates to reconstruct original training samples with high accuracy through optimization-driven inversion methods.
- 2. Cryptographic Payload Inflation:** Conventional security primitives, such as Fully Homomorphic Encryption (FHE), yield heavy ciphertexts that saturate constrained edge network channels.

1.1 Adversarial Model and Zero-Trust Principles

We consider an 'honest-but-curious' central coordinator that executes protocol routines as specified but continuously analyzes model updates to extract underlying client data. Additionally, malicious network eavesdroppers may intercept client-to-server traffic streams.

Standard Unencrypted Architecture (Vulnerable):

Edge Node $\rightarrow$ [Plaintext Gradient Vector ∇W] $\rightarrow$ Central Server (Inversion Attack Reconstructs Inputs)

Dense Homomorphic Encryption Scheme (Bandwidth Limited):

Edge Node $\rightarrow$ [Encrypt Entire Gradient Matrix $\text{Enc}(\nabla W)$] $\rightarrow$ [Severe 50x Payload Expansion] $\rightarrow$ Aggregator

CryptoSparsity Framework (Proposed Method):

Edge Node $\rightarrow$ [Dynamic 95% Magnitude Pruning] $\rightarrow$ [Encrypt Top 5% Residuals] $\rightarrow$ Compact Transmission $\rightarrow$ Aggregator

1.2 Principal Contributions

- **Dynamic Residual Sparsification:** An adaptive magnitude-based filtering mechanism that prunes up to 95% of non-critical gradient elements prior to encryption.
- **Homomorphic Parameter Aggregation:** An encrypted zero-trust aggregation protocol leveraging additive homomorphic keys, allowing direct parameter combination over sparse ciphertexts.
- **Residual Error Compensation:** An error-tracking accumulation queue that retains non-transmitted gradient residues across iterations to ensure steady convergence.

2. Mathematical & Cryptographic Formulation

Let client k maintain local model weights w_t at training iteration t . The local loss gradient is defined as:

$$g_k^{(t)} = \nabla F_k(w_t)$$

To enforce communication efficiency and strong privacy, CryptoSparsity applies a top- k selection function S_p where p signifies the retention ratio (e.g., $p = 0.05$):

$$v_k^{(t)} = S_p(g_k^{(t)} + e_k^{(t-1)})$$

In this expression, $e_k^{(t-1)}$ represents the error residual vector carried over from unselected updates in previous steps. The revised error residual is maintained as:

$$e_k^{(t)} = e_k^{(t-1)} + g_k^{(t)} - v_k^{(t)}$$

2.1 Homomorphic Encryption on Sparsified Indexes

The non-zero residual components are encrypted via a public key pk under the Paillier Additive Homomorphic Scheme:

$$c_k^{(t)} = \text{Enc}_{\{pk\}}(v_k^{(t)})$$

Exploiting homomorphic properties, the aggregator performs parameter aggregation directly in the encrypted space without decrypting intermediate client updates:

$$\text{Enc}_{\{pk\}}(v_{\text{global}}^{(t)}) = \text{Prod}_{\{k=1\}^K} \text{Enc}_{\{pk\}}(v_k^{(t)})$$

3. Experimental Evaluation

3.1 Testbed Setup

The framework was validated on a physical deployment of 150 edge devices (NVIDIA Jetson Xavier modules and Raspberry Pi 4 nodes) linked over simulated cellular/5G channels.

3.2 Performance Benchmark

Security Strategy	Payload (MB)	Round	Encryption Latency (s)	Inversion Defense Rate	Target (%)	Accuracy
Unencrypted Baseline (Plaintext)	48.5		0.00	0% (Vulnerable)	93.2%	
Differential Privacy ($\epsilon=2.0$)	48.5		0.12	68%	88.4%	
Dense Paillier HE	485.0		14.80	100% (Secure)	93.1%	
CryptoSparsity (Proposed Framework)	56.2		2.65	100% (Secure)	92.8%	

Experimental findings confirm that CryptoSparsity cuts the severe transmission overhead of standard homomorphic encryption by 88.4% while ensuring complete protection against gradient inversion attacks.

4. CONCLUSION

CryptoSparsity successfully integrates dynamic gradient sparsification with additive homomorphic encryption to deliver zero-trust privacy for edge computing environments. By filtering non-critical parameters prior to cryptographic processing, the platform makes homomorphic encryption practical for resource-limited edge networks.

REFERENCES

1. Zhu, L., et al. (2019). Deep Leakage from Gradients. NeurIPS.
2. Paillier, P. (1999). Public-Key Cryptosystems Based on Composite Degree Residuosity Classes. EUROCRYPT.
3. Acar, A., et al. (2018). A Survey on Homomorphic Encryption Schemes. ACM Computing Surveys.
4. Aivodji, U., et al. (2020). GAG: Gradient Aggregation Guard for Privacy-Preserving Federated Learning. IEEE TDSC.