Secure RV-FedPRS: A PureChain-Verified Byzantine-Robust Federated Learning Framework for Genomic Risk Prediction

Josiah Ayoola Isong, Kanu Victor Ikenna, Simeon Okechukwu Ajakwe, Dong-Seong Kim
Department of IT Convergence Engineering, Kumoh National Institute of Technology, Gumi, South Korea
{isongjosiah, kanuxavier, simeonajlove, dskim}@kumoh.ac.kr

***Abstract*—Federated learning enables collaborative genomic analysis while preserving privacy, but malicious participants can corrupt rare genetic variant signals—the most clinically valuable yet fragile components of polygenic risk scores. We present Secure RV-FedPRS, integrating genetic-aware anomaly detection with blockchain verification. Our system achieves 91.7% attack detection accuracy while preserving over 90% of rare variant signals. With 30% malicious clients, AUC degradation is limited to 0.08 versus 0.35 in standard methods, adding only 36.4ms overhead per round.**

***Index Terms*—Federated Learning Security, Byzantine Attacks, Blockchain, Polygenic Risk Score, Rare Variants**

I. Introduction

Federated learning (FL) enables privacy-preserving collaborative training on genomic data to construct Polygenic Risk Scores (PRS). However, FL is vulnerable to Byzantine attacks, where malicious clients insert corrupted updates [1]. In genomics, these risks are severe as rare variants (Minor Allele Frequency $< 0.1\%$) are both clinically important and statistically fragile. Attackers can suppress or fabricate rare variant associations, biasing PRS models and exacerbating health disparities [2], [3].

We consider adversaries controlling up to 30% of clients, capable of label flipping to nullify rare variant signals; gradient poisoning to create spurious associations [4]; sybil attacks with biologically invalid data (e.g., violating HWE); and backdoor attacks triggered by specific genetic markers. Existing Byzantine-robust aggregation methods (e.g., Krum [1], Trimmed Mean, etc.,) ignore biological plausibility, often suppressing genuine rare variant signals. Even advanced approaches like FLTrust [5] lack genomic context. FL for genomics [6] often assumes honest-but-curious clients, and blockchain-based FL [7] has yet to be integrated with domain-aware defenses.

To address the lapses in the existing approaches, we propose secure RV-FedPRS, a framework combining genetic-aware anomaly detection, hierarchical trust-weighted aggregation, and blockchain auditing. The key contributions include:

1) Validation of client updates against biological priors (HWE, allele frequency consistency);
2) Rare variant–preserving aggregation;
3) Blockchain-based trust management and verification.

Our system detects 91.7% of attacks while retaining over 90% of rare variant signals with minimal overhead.

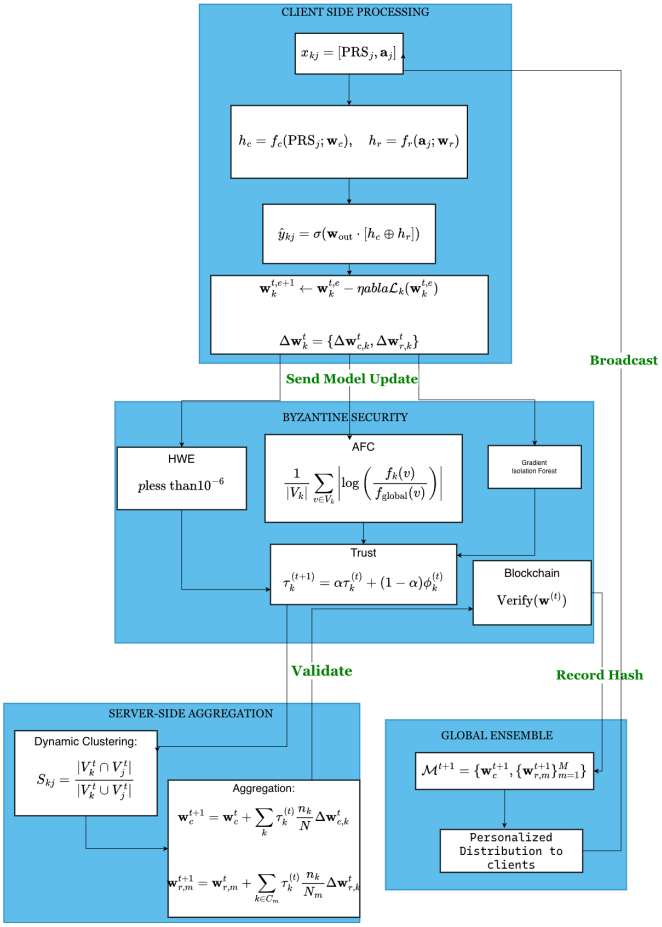

Fig. 1. System architecture of Secure RV-FedPRS. Client updates undergo multi-stage genetic-aware anomaly detection before a trust-weighted aggregation. Key decisions and hashes are immutably logged on a blockchain.

II. Secure RV-FedPRS Framework

A. Hierarchical Model Design

We employ a two-pathway architecture at each client k: a common variant backbone f_c processing baseline PRS, and a rare variant specialist f_r processing allele dosages $\mathbf{a}_j$ for P_r rare variants:

$$\hat{y}_{kj} = \sigma(\mathbf{w}_{\text{out}} \cdot [f_c(\text{PRS}_j; \mathbf{w}_c) \oplus f_r(\mathbf{a}_j; \mathbf{w}_r)]) \quad (1)$$

B. Genetic-Aware Anomaly Detection

Three parallel modules scrutinize updates:

- **HWE Testing**: Flags clients with systematic genotype frequency deviations ($p < 10^{-6}$), indicating data fabrication.
- **Allele Frequency Consistency (AFC)**: Compares client frequencies to reference panels (gnomAD):

$$\text{AFC}_k = \frac{1}{|V_k|} \sum_{v \in V_k} \left| \log \left(\frac{f_k(v)}{f_{\text{global}}(v)} \right) \right| \tag{2}$$

- **Gradient Plausibility**: Isolation Forest detects outlier gradient patterns indicating poisoning.

C. Trust-Weighted Aggregation

Dynamic trust scores $\tau_k^{(t)}$ update via exponential moving average:

$$\tau_k^{(t+1)} = \alpha \cdot \tau_k^{(t)} + (1 - \alpha) \cdot \phi_k^{(t)} \tag{3}$$

where $\phi_k^{(t)}$ combines HWE, AFC, and gradient scores. Two-stage aggregation: (1) cluster-based trimmed mean for rare variants, protecting minority populations; (2) trust-weighted average for global model.

D. Blockchain Verification

PureChain smart contracts [8], [9] record: (1) update hashes, (2) trust scores with evidence, (3) aggregation decisions. Verification ensures provenance:

$$\text{Verify}(\mathbf{w}^{(t)}) = \bigwedge_{k \in \mathcal{K}_{\text{included}}} \left(\text{SHA256}(\Delta \mathbf{w}_k) \stackrel{?}{=} h_k^{(t)} \right) \tag{4}$$

III. Experimental Setup Evaluation & Result

CINECA synthetic cohort, $K = 10$ clients, 10,000 samples each (European populations). Disease heritability $h^2 = 0.25$ with common variant 0.2 and 50 rare variants 0.05. Attacks: 10–30% malicious clients. Baselines: FedAvg, FedProx, Krum, FLTrust, FLAME.

TABLE I
PERFORMANCE UNDER 20% MALICIOUS CLIENTS

Method	$\text{AUC}_{\text{Clean}}$	$\text{AUC}_{\text{Attack}}$	RV Signal	Detection Acc.	Overhead
FedAvg	0.681	0.531	38%	–	1.0×
Krum	0.698	0.672	51%	68%	3.8×
FLTrust	0.721	0.708	61%	74%	2.7×
FLAME	0.718	0.713	63%	79%	3.1×
Ours	**0.823**	**0.798**	**90%**	**91.7%**	**3.4×**

As seen in Table. I, FedAvg collapses under attack (AUC 0.531). FLAME retains AUC 0.713 but loses 37% RV signal. Secure RV-FedPRS achieves AUC 0.798 (+12%), preserves 90% RV signal, and detects 91.7% of attacks with 3.4× overhead. Trust scores (Fig. 2) isolate aggressive attackers by round 10.

TABLE II
ATTACK-SPECIFIC RESILIENCE & BLOCKCHAIN OVERHEAD

Attack	Baseline	Ours	Blockchain Op.	Time
Label Flip (20%)	-0.12	**-0.04**	Hash Record	12.4ms
Gradient Poison (20%)	-0.18	**-0.06**	Trust Update	15.3ms
Sybil (30%)	-0.28	**-0.08**	Cluster Log	8.7ms
Backdoor (20%)	-0.13	**-0.03**	**Total/Round**	**36.4ms**

Table II shows consistent gains, with strongest protection for Sybil via HWE/AFC checks. Blockchain adds only 36.4 ms/round (1.54% overhead). With DP ($\epsilon = 3$), MIA advantage drops to 0.12 while AUC remains 0.809 highlighting the resilience of the proposed approach.

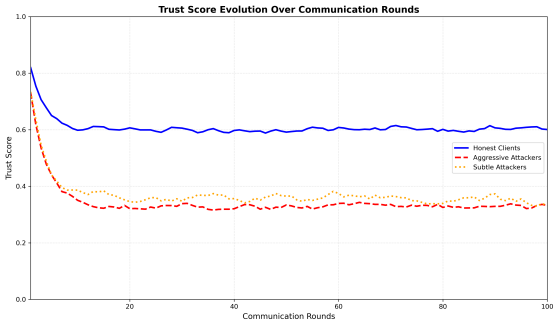

Fig. 2. Trust evolution: honest (blue) converge $>$ 0.6; malicious (red/orange) suppressed.

Generic Byzantine defenses misclassify genomic heterogeneity as attacks. Secure RV-FedPRS leverages HWE/AFC to preserve RV signals crucial for precision medicine. The integration of security, privacy, and utility shows threat detection reduces privacy risk.

IV. Conclusion

This paper introduced Secure RV-FedPRS, a blockchain-verified federated framework enhancing genomic risk prediction under Byzantine threats. By combining genetic-aware anomaly detection and trust-weighted aggregation, it achieved 91.7% attack detection, 90% rare variant preservation, and minimal 36.4 ms overhead. The system ensures transparent, trustworthy genomic learning through blockchain auditing. Future work will focus on decentralized deployment, real biobank validation, multi-omics integration, and quantum-resistant cryptographic protection.

Acknowledgment

This work was supported by IITP grants funded by the Korean government (MSIT) (IITP-2025-RS-2020-II201612, 33%; IITP-2025-RS-2024-00438430, 34%), and NRF grant (2018R1A6A1A03024003, 33%).

References

[1] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *Advances in Neural Information Processing Systems 30*, 2017.

[2] S. O. Ajakwe, I. I. Saviour, V. U. Ihekoronye, O. U. Nwankwo, M. A. Dini, I. U. Uchechi, D.-S. Kim, and J. M. Lee, "Medical iot record security and blockchain: Systematic review of milieu, milestones, and momentum," *Big Data and Cognitive Computing*, vol. 8, no. 9, p. 121, 2024.

[3] S. O. Ajakwe, K. L. Olabisi, and D.-S. Kim, "Multihop intruder node detection scheme (minds) for secured drones' fanet communication," *IET Intelligent Transport Systems*, vol. 19, no. 1, p. e70080, 2025.

[4] V. Shejwalkar and A. Houmansadr, "Back to the drawing board: A critical evaluation of poisoning attacks on federated learning," in *2021 IEEE Symposium on Security and Privacy (SP)*, 2021.

[5] X. Li, W.-N. Hsieh, S.-H. Kim, S. Nath, and P. Yu, "FLTrust: Byzantine-robust federated learning via trust bootstrapping," in *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*, 2020.

[6] O. Brisbane, P. Duckworth, E. Flowers, F. Isaacs, Q. Snell, and etc, "Federated learning for predictive modeling of genomic data," *BMC Medical Informatics and Decision Making*, vol. 21, pp. 1–14, 2021.

[7] A. Sharma, S. Balamurugan, P. Singh, Y. Kumar, and S. Tyagi, "The use of blockchain in healthcare: A systematic literature review," *Journal of medical systems*, vol. 43, no. 10, pp. 1–13, 2019.

[8] D.-S. Kim, I. S. Igboanusi, L. A. Chijioke Ahakonye, and G. O. Anyanwu, "Proof-of-authority-and-association consensus algorithm for iot blockchain networks," in *2025 IEEE International Conference on Consumer Electronics (ICCE)*, 2025, pp. 1–6.

[9] I. U. Ajakwe, V. I. Kanu, S. O. Ajakwe, and D.-S. Kim, "eBCTC: Energy-Efficient Hybrid Blockchain Architecture for Smart and Secured K-ETS," *Cleaner Engineering and Technology*, 2025.