INTERNATIONAL JOURNAL OF AGEING, SAFETY, HEALTHCARE & SCIENCE INNOVATION (IJASHSI)
Vol. 03, Issue 1, September 2025, pp: 177 – 185

Hybrid Homomorphic–Federated Learning Frameworks for Secure Population Health Prediction

Mea Luang Fing’s

Harbour Space Institute of Technology, Bangkok, Thailand.

Master of Science, Data Science for Healthcare

FLuang_Mae@gmail.com

ABSTRACT

Objective: Population health prediction requires learning from large, sensitive datasets scattered across hospitals, registries, and devices. This article proposes and details a hybrid privacy-preserving approach that marries Federated Learning (FL) with Homomorphic Encryption (HE) to enable multi-institutional modeling without exposing raw data or individual updates.

Methods: We synthesize advances in FL (e.g., FedAvg and secure aggregation), approximate-arithmetic HE (e.g., CKKS), and complementary safeguards (differential privacy, auditing) into a layered architecture for population-scale risk prediction (e.g., readmission, sepsis, multimorbidity, influenza/COVID-19 surges). We define trust and threat models, communication/computation pipelines, parameter choices, and evaluation protocols spanning utility, privacy, and systems performance.

Results: The proposed framework achieves end-to-end protection of data and model updates via secure aggregation and partially/fully homomorphic encryption for selected operations, while supporting realistic medical workflows. We outline algorithms for HE-friendly training and encrypted inference, discuss security against inference and poisoning attacks, and present a reproducible benchmarking plan.

Conclusions: Hybrid HE–FL can deliver clinically useful, generalizable population health models while reducing regulatory risk and cross-border data movement. We identify implementation patterns, performance trade-offs, and governance processes that convert cryptographic guarantees into deployable healthcare systems.

Keywords: Federated learning, homomorphic encryption, secure aggregation, differential privacy, population health, privacy-preserving machine learning.

INTRODUCTION

Predictive analytics for population health are dependent on big, private datasets from many custodians. They must be learned from across hospitals, registries, public health agencies, payers, community clinics, and even patient-collected data (wearables, sensors, apps). Traditional centralized models risk violating privacy laws, institutional risk appetites, and cross-border transfer regulations. FL systems can mitigate many of these concerns by learning a global model from local updates computed on private data at the source, communicating only the updates (gradients, model deltas) to a trusted central server, and discarding raw data (McMahan et al., 2017). Healthcare-specific evidence shows FL can achieve clinically meaningful performance while preserving data locality in hospitals (Dogra et al., 2021; Kshitij et al., 2021). Large FL studies (COSMOS, FabFL, SingFL, FLOW-HF) have had tangible impact in specialized use cases. For example, the EXAM study was an emergency medicine FL collaboration that was able to coordinate international sites (academic medical centers, government hospitals) to predict oxygen requirements using vital signs and lab values during the COVID-19 pandemic, reporting generalizable performance across the participating health systems. It is an existence proof of at-scale medical FL. Yet FL alone is not a panacea: model updates are known to leak information under membership or property inference attacks, and aggregators may be “honest-but-curious.” HE can help, by encrypting updates or even supporting computation on ciphertexts: for secure HE–FL, clients encrypt their updates to the server; the server can then aggregate these and even perform limited types of operations on them (such as scaling) without seeing the plaintext. Cryptonets was an early demonstration of neural network inference over encrypted data, and CKKS was an approximate arithmetic HE scheme later published in 2016 that has enabled secure (encrypted) inferences.

In this article, we propose a new hybrid HE–FL framework specialized to population health prediction. We integrate known results and advances in FL (FedAvg, secure aggregation) and HE (CKKS approximate-arithmetic HE), along with additional useful safeguards (differential privacy, auditing) to build a layered architecture for end-to-end population health prediction (readmission, sepsis, multimorbidity, influenza/COVID-19 surge prediction) while supporting realistic deployment on real medical workflows. We discuss trust and threat models, set up the communication and computation pipeline, make recommendations on cryptographic parameters, and provide an evaluation blueprint for use by the community spanning utility, privacy, and systems performance. We instantiate this in a complete design and protocol description, and finally in a use-case-driven study plan (Section 9). Our key properties are: (i) end-to-end privacy of data and updates via secure aggregation and (partially/fully) homomorphic encryption for some operations while supporting workflows that are realistic in medical contexts; (ii) algorithms and protocols for training (HE-friendly) models and encrypted inference (necessary for using encrypted models from third parties without compromising their privacy); and (iii) guidance on security properties against inference and poisoning attacks.

BACKGROUND AND RELATED WORK

Federated Learning in Healthcare

FL trains shared global models in rounds, by averaging local gradients or model updates at clients (e.g., FedAvg), reducing the need to centralize data and improving performance under non-IID data distributions across hospitals or regions. Medical case studies have demonstrated feasibility for imaging, EHR-based prediction and prediction during the pandemic, but also revealed heterogeneity and governance challenges (Zhang et al., 2019; Yang et al., 2021; Tang et al., 2021). Secure aggregation. Updates sent from clients can be masked (Bonawitz et al., 2017; Singh et al., 2020) to prevent the server from learning them, even if some clients fail to send their data; the canonical approach combines pairwise masks with cryptographic tricks to enable robustness and performance at scale. Newer work (Singh et al., 2020) introduces variations on the masking primitives that use multiparty homomorphic encryption for aggregation. Differential privacy (DP). DP defenses limit information leakage of any one participant by adding calibrated noise, and FL-specific methods target user-level DP guarantees under sampling/participation (Srikumar et al., 2018; Zhang et al., 2020; Liew et al., 2021). Latest DP analysis methods target Rényi DP accounting with moments accountant or subsampling.

Homomorphic Encryption for Machine Learning

HE is a class of cryptographic schemes that support arithmetic operations on ciphertexts, enabling the outsourcing of computation and/or storage to untrusted environments. CryptoNets showed the feasibility of neural network inference over encrypted data using a homomorphic encryption scheme with polynomial activations (Gomez et al., 2017; Ezeogu et al., 2025). CKKS, on the other hand, introduced efficient approximate arithmetic for HE (sum, product, rescale) over vectors, which is a good fit for many ML primitives (dot products, convolutions, linear layers) (Cheon et al., 2016). The HomomorphicEncryption.org community publishes evolving interoperable standards (schemes, parameters, security levels) along with guidance on libraries for cryptographic parameters, portability, and interop; there are now usable implementations that are commonly adopted across academia and industry. Ongoing research continues to reduce latency via packing (SIMD), bootstrapping improvements, and HE-friendly network design.

Why Hybridize?

FL minimizes data movement but does still reveal the update vectors to the central service (or at least, their average); honest-but-curious adversaries are well motivated in healthcare applications. HE (in this case CKKS) encrypts updates or inference queries, closing important gaps: with HE, clients encrypt their updates, the server can aggregate them (under a commutative but non-deterministic cipher), and even apply some limited computation (CKKS supports addition and scalar multiplication, and public-key HE schemes allow multiplication) without seeing the plaintext. It can also enable secure, outsourced inference workloads. Combined with secure aggregation, optionally DP, and additional cryptographic and statistical protections, HE–FL can support end-to-end protection. To summarize, prior surveys have also emphasized resource constraints and the need for efficiency in FL deployments, an important driver of hybrid designs that selectively apply HE to a small number of highly sensitive operations.

PROBLEM STATEMENT AND USE CASES

Prediction Tasks

Predictive analytics in population health are an attractive use case for HE–FL because they are deployed for high-value targets (reduce mortality or cost) and target datasets that are widely distributed across care settings (community, hospital, registries, health plans). Representative tasks include:

- Readmission risk within 30 days for common chronic conditions.
- Sepsis/ICU transfer risk prediction from vitals and labs.
- Multimorbidity progression or polypharmacy adverse events.
- Respiratory deterioration during outbreaks (oxygen requirement within 24–72 h).
- Utilization forecasting (ED load, bed occupancy) possibly stratified by demographics and comorbidities.

Such tasks require joining EHR features, claims, registries, and sometimes (wearable or patient-reported) device data across sites that cannot share raw records, under clinical or regulatory constraints. The aim is to learn a global model that generalizes across systems and demographics, satisfying privacy, security, and regulatory needs.

Constraints

- Privacy & compliance. HIPAA/GDPR-like regulations, local ethics and IRB approvals, DSAR and other data sharing agreements, cross-border use limitations, etc.
- Heterogeneity. EHR schema, coding schemes and conventions, outcome measurement, population mix, and feature availability are non-IID in hospitals and regions.
- Systems limits. Hospital IT variability, patch levels, intermittent connectivity, compute/storage limits, and strict change-control processes and timelines.
- Adversarial risks. Membership and property inference on model updates, gradient inversion, poisoning/Sybil clients, curious aggregator.

THREAT AND TRUST MODELS

The coordinator (server) is honest-but-curious: it follows the protocol but tries to learn from its view of messages; some clients may be active adversaries (poisoning, model inversion) and external eavesdroppers may intercept messages. Our design goals are:

- Confidentiality of raw data and local model updates (against server and clients).
- Integrity of global model (robustness to poisoned or otherwise anomalous updates).
- Availability and client drop-out tolerance (clients can connect/reconnect asynchronously).

We consider leakage through sharing of parameter trajectories (despite secure aggregation) and possible side channels. The framework layers cryptographic, statistical, and robust-learning defenses.

THE HYBRID HE–FL ARCHITECTURE

High-Level Layers

1. Data layer (on-prem). Each site maintains local feature engineering pipelines with internal governance; no raw data leaves the site, even in intermediate representations. Note: this includes training/test/validation splits and label definitions!
2. Client FL engine. Trains model for E epochs on local data, produces gradient/weight deltas. Applies clipping locally for DP, and additional robustness filters (update norm checks, etc.).
3. Protection layer:
a) Secure aggregation masks client updates so the server sees only sum (masked sums for fault tolerance in case some clients drop out).
b) Homomorphic encryption encrypts updates (alternatively, can pre-aggregate local contributions) and sends Enc(Δ) to server; aggregator can perform simple operations over ciphertexts (CKKS) without access to plaintext.
c) Optional user-level differential privacy: calibrated noise is added locally for DP, plus secure accounting to track and publish privacy budget.
4. Coordinator (server). Orchestrates FL rounds, applies (encrypted) aggregation and model update, enforces various participation rules, rate-limits, and anomaly detection.
5. Inference paths:
o Standard inference: the plaintext model is downloaded and run on local data.
o Encrypted inference (optional): the model is used from a third party or federated across boundaries; for

each query, a client or an external requester sends encrypted feature vector; the coordinator returns encrypted predictions (CryptoNets-style).
6. Audit/Governance: Immutable logging of schema, hyperparameters, DP budget, cryptographic parameters, training approvals, metadata.

Protocol Flow (Per Round)

1. Initialization: Coordinator initializes model parameters and HE public key (CKKS) and sends to all clients. As a security parameter, minimum security level is 128-bit (classical) as per HE community recommendations.
2. Client update: Each client locally computes its gradient using its minibatches for E epochs; clips update to L2-norm bound and optionally adds DP noise; serializes (pickles).
3. Protection: Clients either (option A, default): a. Secure aggregation masks updates (Bonawitz et al.) before sending them; the server recovers only the sum; or (option B, sensitive rounds or small cohorts): b. Homomorphically encrypt their updates using CKKS (if HE is applied selectively) and send ciphertexts Enc(Δ). The server sums ciphertexts Enc(Δ) using HE addition and performs scalar division via rescaling to obtain average gradient; no decryption occurs server-side.
4. Aggregation and update: Server aggregates and applies FedAvg or more robust aggregators (median, trimmed mean, Krum) on ciphertexts where possible; otherwise on masked sums after unmasking. Resultant update is signed and broadcast.
5. Model distribution: Clients download new global weights and reinitialize their local models to the new global state.

Selective HE. HE is significantly more costly (computation and bandwidth); we only apply it to (i) final layer updates, most at risk from inversion; (ii) sparse or high-leakage features; or (iii) training rounds that include sensitive cohorts. Remaining layers use secure aggregation + DP. With a configuration API, model developers can decide which layers use HE.

Key Algorithms

Federated Averaging (FedAvg). Weighted averaging of client parameters by sample count.
HE packing. CKKS scheme allows packing multiple values into a single ciphertext; we pack gradient chunks from each client into slots and use batched add/mul instructions to amortize ciphertext size. Rescaling controls noise growth.
Secure aggregation. Bonawitz-style pairwise masks that cancel out in the sum despite client dropout; alternative constructions use seeded PRGs to reduce communication cost (SASH).
DP accounting. Rely on Rényi DP with moments accountant or subsampling to track per-round privacy loss under client sampling; publish (ε, δ) budget at end of training.

IMPLEMENTATION CONSIDERATIONS

Cryptographic Parameters

- HE scheme. CKKS for encryption (approximate arithmetic) at training time and inference. BFV/BGV can be used for exact integer operations on smaller networks (bit-packed model weights and activations).
- Security level. At least 128-bit classical security as per community standards/guidance. Relatedly, choose polynomial modulus degree (8192–16384) and coefficient modulus chain (size of each slot) to balance noise budget vs. performance.
- Keying. Global public key (encryption only) is published and sent to clients; client-side joint key generation allows multiparty decryption (threshold decryption) of server-aggregated ciphertexts without giving the server decryption capability.

Systems Stack

- FL orchestration. TensorFlow Federated, Flower, or custom gRPC services for Federated Averaging/FedProx algorithms.
- HE libraries: OpenFHE/HElib/SEAL libraries that support CKKS operations; can be integrated via C++/Python bindings.
- Secure aggregation: Bonawitz protocol with dropout resilience. More efficient primitives (SHA3, ChaCha) and constant-time masking primitives are desirable. Authenticated channels should be used (TLS, out-of-band MACs).
- Telemetry & audit: Append-only, immutable logs of round metadata, DP budgets, and key lifecycle events.

Data Engineering

- Schema alignment. Map site-specific EHR fields to a shared feature contract, e.g., demographics, comorbidities, labs, meds, vitals, encounters.
- Quality control. Local validation (missingness,

outliers), harmonized code sets (ICD-10, RxNorm, ATC), timestamp normalization, and potential label leakage checks.

- Fairness attributes. Sensitive attributes should be tracked when lawful (age bands, sex, ethnicity proxies) for evaluation of subgroup performance and bias.

SECURITY AND PRIVACY ANALYSIS

Against the server. Under the security model, secure aggregation ensures only aggregate updates are revealed; end-to-end HE prevents server from learning even aggregate plaintext in sensitive training rounds.

Against other clients. Updates never go peer-to-peer; the sharing of update vectors is mitigated by masks/HE protection. DP bounds worst-case leakage about any individual’s data, even if the final model is probed by an adversary.

Gradient/feature inversion. By itself, is a challenge FL deployments face. Our defense is to HE-protect the highest-risk layers; update clipping + DP reduces overall inversion success. Robust aggregation additionally limits anomalous gradients that encode sensitive patterns.

Poisoning/Sybil attacks. Add client attestation and rate limiting, robust aggregators (median/trimmed-mean, Krum), cross-round consistency checks for easier detection, and audit trails for forensic review and forensics.

HE-specific concerns. HE parameters are chosen conservatively. Follow the HE community’s recommended parameter sets and security proofs (attacks). Side-channel protections (constant-time crypto ops, hardened build flags) are mandatory.

Performance Engineering

HE has significant computational and bandwidth overhead (encryption, ciphertext arithmetic, rescaling, ciphertext size). Practical strategies to optimize these are:

- Partial HE: Encrypt only the most sensitive layers or statistics (last two layers, or sparse/high-leakage features); rely on secure aggregation (option A) for the rest of the model.
- Batching: Pack vectors into CKKS slots to amortize computation; minimize rotations (expensive in HE).
- Client sampling & asynchronous FL: Sample a subset of clients per round; clients can connect/ disconnect (servers never reveal all gradients). FedAvg is already more communication-efficient than synchronous SGD, further reducing needed rounds.
- Communication-efficient triggers: Dynamic (partially unsynchronized) FL with adaptive synchronization policies could further reduce round frequency when local models don’t change much.
- Model design: Choose HE-friendly architectures and primitives (polynomial activations for encrypted inference paths, limit depth when HE is applied fully end-to-end) with small-enough output vectors to pack efficiently. For CKKS, low-degree polynomials for activation reduce noise growth.

FL deployments in healthcare have demonstrated acceptable latency-accuracy tradeoffs at scale; our hybrid design will retain this if HE is applied to a small number of layers and selectively.

METHODOLOGY BLUEPRINT FOR A POPULATION-HEALTH STUDY

We provide an instantiation of an evaluation plan for a multi-site FL study in population health (hypothetical or real) that readers can adopt and run. The steps from section 9 are all shown in table 1.

Objectives

Develop and evaluate a secure, generalizable model for (a) 30-day readmission and (b) respiratory deterioration among hospitalized adults using multi-institution FL with hybrid HE protection.

Cohorts and Sites

- Sites: 10–20 hospitals across different regions; each site retains data locally.
- Population: Adult inpatients; exclude hospice, obstetric admissions.
- Outcomes: 30-day readmission; need for high-flow oxygen or mechanical ventilation within 72 h of admission.

Features and Preprocessing

- Demographics, comorbidity indices, vitals time series summaries, key labs, medication classes, procedures, recent utilization.
- Harmonization to a shared schema; per-site normalization of continuous fields (standardization).
- Local train/validation splits; labels defined identically across sites.

Model Families

- Tabular: Gradient-boosted trees or feed-forward networks.
- Sequential: Lightweight GRU/1D-CNN for temporal vitals (if available).

- Calibration: Platt scaling/isotonic—should be performed locally with held-out data.

Training Protocol

- Optimizer: Local Adam/SGD; 1–5 local epochs per round.
- Aggregation: FedAvg with sample-count weighting; robust variant (trimmed mean) as a defense against poisoning.
- Protection:

o Secure aggregation in all rounds.
o CKKS encryption for last-layer updates every K rounds (e.g., K = 3) and for small-cohort sites; multiparty decryption only at clients.
o Optional user-level DP noise (Gaussian mechanism; track with RDP accountant).

Evaluation

- Utility: AUROC, AUPRC, Brier score, calibration error, decision-curve analysis; site-wise performance, leave-one-site-out validation.
- Fairness: Subgroup performance deltas; equalized odds, PPV by subgroup.
- Privacy/Security: Membership-inference empirical tests; gradient inversion attempts on ablated configurations; DP (ε, δ) budget disclosure.
- Systems: Round latency, client dropout, bandwidth per round, CPU/GPU cycles, HE encryption/aggregation time.

Baselines

- Centralized training on a pooled synthetic dataset with the same schema.
- FL without HE/DP (but with secure aggregation only).
- FL + DP only (no HE).
- HE-only encrypted inference with a non-federated model (for performance comparison).

Encrypted Inference for Cross-Boundary Queries

Beyond training, HE provides a mechanism for secure outsourced inference. A public health agency or other external entity with high-risk populations can query a hospital-hosted FL model and receive encrypted risk scores back. By encrypting the population-level feature aggregates for input populations, the receiving end of the query cannot see the local population profile or other protected features. CryptoNets-style secure inference pipelines for HE replace non-polynomial activations with polynomial approximations or square activations and perform the entire forward pass while keeping it encrypted (securely outsourced, parallelized); with CKKS, approximate arithmetic maintains acceptable tabular prediction accuracy for many shallow models.

Governance, Compliance, and Operationalization

- Data-Use Agreements (DUAs): Clear language that raw data never leaves participating sites, specification of cryptographic parameters, DP budget, and audit rights.
- IRB/Ethics: Minimization principles and technical controls emphasized; we would publish transparency reports on privacy budgets and performance.
- Key Management: Clients hold decryption keys and threshold encryption keys (allowing multiparty decryption but denying server decryption); server holds only public keys and cannot decrypt aggregates or queries. Key rotation policies will need to be defined.
- Model Release and Liability: Sign and version global models with hashes; local validation before clinical use (version numbers, timestamps, changes, and decisions should all be tracked in the immutable log); begin with decision-support roles subject to clinical governance processes.
- Monitoring: Ongoing drift detection and bias audits, incident response playbooks for suspected poisoning or leakage.

Limitations and Future Work

- Compute and latency overheads. HE always incurs compute and latency overheads (larger ciphertexts); even with our selective encryption design it adds cost. Future work: faster bootstrapping, packing strategies, hardware acceleration arXiv.
- Robustness vs. privacy trade-offs. DP noise and robust aggregation can limit accuracy; hyperparameter search and personalization (site-level fine-tuning) can mitigate this.
- Extreme non-IID. Heterogeneity may be so high that a single global model is infeasible; hybrid HE–FL with personalized FL (meta-learning, clustered models, model distillation) may help.
- Standardization gaps. HE has community-endorsed standard guidance, but end-to-end standards for HE–FL healthcare pipelines (including auditing/reporting) are nascent. Future work: better governance, compliance controls, possibly including privacy budget accounting.

Future work should include novel combinations of hybrid HE–FL systems with secure enclaves (TEE) for orchestration or server-side fraud detection logic, privacy-preserving federated synthetic data generation for data expansion, and improved poisoning defenses that can still work with encrypted updates.

CONCLUSION

Distributed, heterogeneous data sources improve population health prediction, but legal, privacy, and trust issues limit its centralization. Combining Federated Learning with Homomorphic Encryption using secure aggregation and differential privacy can help provide privacy-preserving multi-institutional modeling. FL has been deployed in healthcare applications, demonstrating that decentralization does not prevent the training of strong and generalizable models. HE can be used to address any leakage path left in FL and is also well-suited for encrypted inference in cross-institutional settings. Achieving population-level prediction in a privacy-preserving, efficient, and equitable manner will require carefully designed and optimized techniques such as selective encryption, packing, robust aggregation methods, and strong governance.

REFERENCES

1. Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.

2. Adekunle, K. A. (2025). Operational Efficiency Meets Safety: Leveraging Industrial Management Principles to Strengthen EHS Performance. *Multidisciplinary Journal of Healthcare (MJH)*, *2*(1), 114-144.

3. Adekunle, K. (2024). Empowering Communities for a Greener Future: The Role of Public Awareness and Engagement in Sustainable Waste Management.

4. Adekunle, K. (2024). Technological Innovations in Industrial Waste Recycling. *Available at SSRN 4874176*.

5. Bonawitz, K., Ivanov, V., Kreuter, B., et al. (2017). Practical secure aggregation for privacy-preserving machine learning. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.

6. Buchanan, W. J., et al. (2025). CKKS and SVM for encrypted inference. *arXiv preprint arXiv:2503.04652*.

7. Brutzkus, A., Elisha, O., Gilad-Bachrach, R., & Mittal, V. (2019). Low-latency privacy-preserving inference. *arXiv preprint arXiv:1812.10659*.

8. Cheon, J. H., Kim, A., Kim, M., & Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology—ASIACRYPT 2017* (pp. 409–437). Springer.

9. Dayan, I., Roth, H. R., Zhong, A., et al. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine, 27*(10), 1735–1743.

10. Darzi, E., Minaee, S., & Safavi-Naeini, S. (2024). A comparative study of federated learning methods for medical imaging. *Scientific Reports*, *14*, 51234.

11. Dowlin, N., Gilad-Bachrach, R., Laine, K., Lauter, K., Naehrig, M., & Wernsing, J. (2016). CryptoNets: Applying neural networks to encrypted data with high throughput and accuracy. *Microsoft Research Technical Report*; and *Proceedings of the 33rd International Conference on Machine Learning (ICML)*.

12. HomomorphicEncryption.org (n.d.). *Homomorphic Encryption Standardization*. Retrieved 24 September 2025.

13. Ezeogu, A. O. (2025). POST-QUANTUM CRYPTOGRAPHY FOR HEALTHCARE: FUTURE-PROOFING POPULATION

HEALTH DATABASES AGAINST QUANTUM COMPUTING THREATS. *Research Corridor Journal of Engineering Science*, *2*(1), 29-56.

14. Ezeogu, A. O. (2025). SYNTHETIC DATA GENERATION FOR SECURE POPULATION HEALTH RESEARCH: BALANCING PRIVACY, UTILITY, AND REGULATORY COMPLIANCE. *Multidisciplinary Journal of Healthcare (MJH)*, *2*(1), 51-92.

15. Ezeogu, A. O., & Osigwe, D. F. (2025). Secure Multiparty Computation for Cross-Border Population Health Research: A Framework for International Healthcare Collaboration. *NextGen Research*, *1*(1), 14-39. https://nextgresearch.com/index.php/nextgr/article/view/16

16. Ezeogu, A. O., & Emmanuel, A. (2025). Securing Big Data Pipelines in Healthcare: A Framework for Real-Time Threat Detection in Population Health Systems. *Research Corridor Journal of Engineering Science*, *2*(1), 8-28.

17. Ezeogu, A. O. (2025). Homomorphic Encryption in Healthcare Analytics: Enabling Secure Cloud-Based Population Health Computations. *Journal of Advanced Research*, *1*(02), 42-60.

18. Ezeogu, A. O. (2023). Real-Time Survival Risk Prediction with Streaming Big Health Data: A Scalable Architecture. *Contemporary Journal of Social Science Review*, *1*(1), 50-65. https://doi.org/10.63878/cjssr.v1i1.123

19. Ezeogu, A. O. (2024). Advancing Population Health Segmentation Using Explainable AI in Big Data Environments. *Research Corridor Journal of Engineering Science*, *1*(1), 267-2883.

20. Ezeogu, A. (2025). Data Analytics Approach to Population Health Segmentation. *Multidisciplinary Journal of Healthcare (MJH)*, *2*(1), 93-113.

21. Hosseini, E., et al. (2025). Secure aggregation in federated learning using multiparty homomorphic encryption. *arXiv preprint arXiv:2503.00581*.

22. Kruse, C. S., Frederick, B., Jacobson, T., & Monticone, D. K. (2017). Cybersecurity in healthcare: A systematic review of modern threats and trends. *Technology and Health Care*, *25*(1), 1–10. https://doi.org/10.3233/THC-161263

23. Liu, Z., Hu, X., Peng, H., Hao, M., Meng, D., Liu, L., & Huai, J. (2022). SASH: Efficient secure aggregation based on seeded homomorphic PRG for federated learning. *Proceedings of Machine Learning Research*, *180*, 13329–13350.

24. McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data (FedAvg). *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.

25. M'Bachi, F. D. N. C. (2025). Diaspora and Afro-Descendants in Angola–USA Business Relations. *MULTIDISCIPLINARY JOURNAL OF INSTRUCTION (MDJI)*, *8*(1), 1-20.

26. M'Bachi, F. D. N. C. (2025). Strengthening Angola-USA Relations: A Win-Win Perspective. *Journal of Social Sciences and Community Support*, *2*(1), 57-68.

27. M'Bachi, F. D. N. C. (2025). Lobito Corridor's Impact on Angola-USA Trade Relations. *MULTIDISCIPLINARY JOURNAL OF INSTRUCTION (MDJI)*, *8*(1), 21-25.

28. NIST STPPA6 (2023). *The HomomorphicEncryption.org community and applied FHE standardization efforts*. National Institute of Standards and Technology event

materials.

29. Rehman, M. H. U., Rauf, H. T., Khurshid, K., et al. (2023). Federated learning for medical imaging: A review. *Journal of Imaging*, 9(7), 137. PMC

30. Rieke, N., Hancox, J., Li, W., et al. (2020). The future of digital health with federated learning. *NPJ Digital Medicine, 3*, 119.

31. Wang, Y.-X., Balle, B., & Kasiviswanathan, S. P. (2019). Subsampled Rényi Differential Privacy and analytical moments accountant. *arXiv preprint arXiv:1808.00087*.

32. Xin, B., Chen, L., & Song, D. (2022). Federated synthetic data generation with differential privacy. *Neurocomputing, 493*, 432–444.