SIAR-Global Journal of Computer, Information & Library Science
www.siarpublications.org ISSN: 3121-7745
info@siarpublications.org Vol. 2 Issue 1 Jan.-March. 2026
Global Office **South Africa:** 16 Park Ave, Modderfontein, Johannesburg

Privacy-Preserving Machine Learning Technique Using One-Way Hashing for Enhancing Data Security in the Nigerian Healthcare Ecosystem

Damang, F.S., and Gilbert I.O. Aimufua

Department of Computer Science, Faculty of Natural and Applied Sciences,
Nasarawa State University, Keffi, Nigeria
Corresponding Author Email: damangfs@gmail.com Tel: +2348107580503

Abstract

Safeguarding healthcare data in Nigeria remains a pressing challenge, complicated by fragmented infrastructure, limited resources, and evolving regulatory frameworks. This paper presents a conceptual analysis of one-way hashing as a lightweight cryptographic technique for pseudonymizing patient identifiers within federated learning pipelines. By situating hashing in contrast to heavier cryptographic methods such as homomorphic encryption and secure multiparty computation, the study highlights its relative efficiency, scalability, and compliance with the Nigeria Data Protection Regulation (NDPR, 2019) and the Nigeria Data Protection Act (NDPA, 2023). Through analytical benchmarking and illustrative scenarios, hashing is shown to offer a pragmatic balance between privacy preservation and operational feasibility in resource-constrained healthcare environments. The paper concludes that one-way hashing provides a viable conceptual pathway for operationalizing privacy-preserving machine learning in Nigerian healthcare systems, while laying the foundation for future empirical validation.

Keywords: Healthcare data protection; Federated learning; One-way hashing; Pseudonymizations; Privacy-preserving machine learning; Conceptual analysis; Nigeria Data Protection Regulation (NDPR); Nigeria Data Protection Act (NDPA); Cryptographic complexity; Resource-constrained environments

1.0 Introduction

The digitization of healthcare records has created new opportunities for patient care and analytics, but it has also introduced significant risks. In Nigeria, hospitals increasingly rely on electronic health records (EHRs) yet concerns about privacy breaches and regulatory compliance persist. The NDPR and NDPA mandate pseudonymization and data minimization (Nigeria Data Protection Commission, 2023), but many institutions lack the infrastructure to deploy advanced cryptographic solutions. One-way hashing, which irreversibly transforms identifiers into fixed-length codes, offers a practical safeguard. This paper explores how hashing can be integrated into federated learning to protect patient privacy while enabling collaborative analytics.

1.1 Problem Statement

Healthcare institutions in Nigeria face mounting challenges in safeguarding patient data as electronic health records become more widespread. Despite the enactment of the Nigeria Data Protection Regulation (NDPR, 2019) and the Nigeria Data Protection Act (NDPA, 2023), many hospitals lack the technical and infrastructural capacity to deploy advanced cryptographic solutions such as homomorphic encryption or secure multiparty computation. This limitation exposes sensitive patient information to risks of privacy breaches and non-compliance with regulatory mandates. At the same time, the demand for collaborative analytics to strengthen disease surveillance and optimize resource allocation is urgent.

The central problem, therefore, is conceptual: how can secure, privacy-preserving healthcare data sharing and analysis be enabled in Nigeria using approaches that are both technically feasible and compliant with existing data protection laws? This study addresses the problem by analysing one-way hashing as a lightweight cryptographic technique, situating it within federated learning pipelines, and contrasting it with heavier cryptographic alternatives. Through this conceptual lens, the work highlights hashing's potential as a pragmatic pathway for operationalizing privacy-preserving machine learning in resource-constrained healthcare environments.

1.2 Research Objectives

The overarching aim of this study is to develop a conceptual framework for privacy-preserving healthcare data sharing in Nigeria that is both technically feasible and compliant with existing data protection laws. To achieve this aim, the study pursues the following specific objectives:

1. To analyze the challenges of safeguarding patient data in Nigerian healthcare institutions, focusing on infrastructural limitations, resource constraints, and regulatory requirements under NDPR (2019) and NDPA (2023).

2. To examine one-way hashing as a lightweight cryptographic technique for pseudonymisation, situating it within federated learning pipelines as a practical alternative to more complex methods.

3. To conduct a conceptual comparison between hashing, homomorphic encryption (HE), and secure multiparty computation (MPC), highlighting differences in computational complexity, communication overhead, scalability, and regulatory alignment.

4. To illustrate feasibility through analytical benchmarking, pseudo-code simulations, and scenario-based examples that demonstrate how hashing can enable collaborative analytics without exposing raw patient data.

5. To establish a foundation for future empirical validation, identifying pathways for testing the proposed framework with real hospital datasets, larger simulations, and integration with emerging technologies such as blockchain and IoT health devices.

1.3 Research Questions

Building on the problem statement and objectives, this study is guided by the following conceptual research questions:

1. What are the key challenges facing Nigerian healthcare institutions in safeguarding patient data, particularly in relation to infrastructural limitations, resource constraints, and compliance with NDPR (2019) and NDPA (2023)?

2. How can one-way hashing be conceptualized as a lightweight cryptographic technique for pseudonymisation within federated learning pipelines?

3. In what ways does hashing compare to homomorphic encryption (HE) and secure multiparty computation (MPC) in terms of computational complexity, communication overhead, scalability, and regulatory alignment?

4. How can illustrative scenarios and pseudo-code simulations demonstrate the feasibility of hashing for privacy-preserving collaborative analytics in Nigerian healthcare systems?

5. What conceptual pathways exist for extending this framework toward future empirical validation, including real hospital datasets, larger simulations, and integration with emerging technologies such as blockchain and IoT health devices?

2.0 Literature Review

Recent scholarship underscores the importance of privacy-preserving technologies in healthcare. Globally, hashing has been applied to pseudonymize patient identifiers, ensuring confidentiality while supporting interoperability (Zhang et al., 2026). In Nigeria, systemic challenges such as fragmented IT systems and limited cybersecurity expertise complicate compliance (Ohiri, 2025). While homomorphic encryption and secure multiparty computation provide strong guarantees, they remain impractical in low-resource settings due to computational overhead. Hashing, by contrast, is lightweight, cost-effective, and aligns with regulatory requirements (Health Policy and Planning, 2024). Pilot projects in Nigerian hospitals have demonstrated that federated learning frameworks, when combined with hashing, enable secure collaboration without compromising compliance (Nigeria Health Intelligence Report, 2025).

Healthcare data governance in Nigeria has undergone significant transformation with the enactment of the **Nigeria Data Protection Act (NDPA, 2023)**, which provides the first comprehensive legal framework for personal data protection in the country. Scholars have noted that the NDPA builds upon the earlier **NDPR (2019)** but introduces stronger enforcement mechanisms and sector-specific provisions, particularly for healthcare delivery (Nigeria Data Protection Legislation and the Health Sector, 2023). Recent analyses highlight that Nigeria's regulatory approach is increasingly aligned with global standards, though challenges remain in enforcement and institutional capacity (Modilim et al., 2024). Within healthcare, the need for privacy-preserving technologies is amplified by the sensitive nature of patient data and the risks of breaches in under-resourced hospital systems.

Globally, federated learning has emerged as a promising solution for collaborative healthcare analytics. Studies emphasize its ability to enable multi-institutional research without centralizing

raw data, thereby reducing privacy risks (Mridha & Dey, 2026). In Nigeria, the Federal Ministry of Health has prioritized digital health innovations to strengthen surveillance and improve outcomes in infectious disease control and maternal health (Federal Ministry of Health, 2025).

2.1 Conceptual Framework

Safeguarding healthcare data in Nigeria requires solutions that balance technical feasibility, regulatory compliance, and the urgent need for collaborative analytics. This study develops a conceptual framework that positions one-way hashing as a lightweight cryptographic technique for pseudonymisation within federated learning pipelines. The framework illustrates how hashing can serve as a compliance-aligned, resource-efficient alternative to heavier cryptographic methods such as homomorphic encryption (HE) and secure multiparty computation (MPC).

2.2.1 Core Components

The framework is built around four interrelated components:

1. Healthcare Data Protection Needs: The increasing adoption of electronic health records heightens risks of privacy breaches and non-compliance with data protection laws.
2. Regulatory Context: The Nigeria Data Protection Regulation (NDPR, 2019) and the Nigeria Data Protection Act (NDPA, 2023) mandate pseudonymisation of personal data, creating a legal imperative for secure solutions.
3. Technical Options: While HE and MPC offer strong privacy guarantees, they are computationally and communicationally expensive. Hashing, by contrast, is lightweight, scalable, and easier to implement in resource-constrained environments.
4. Collaborative Analytics: Federated learning enables hospitals to collaborate on predictive modelling without sharing raw patient data but requires a privacy-preserving mechanism for identifiers.

2.2.2 **Analytical Lens**

This framework is conceptual rather than empirical. Validation is provided through:

- Complexity benchmarking of hashing, HE, and MPC.
- Illustrative scenarios (e.g., hospital collaboration on disease surveillance).
- Pseudo-code simulations demonstrating hashing workflows.
- Policy alignment with NDPR and NDPA requirements.

2.3 Conceptual Model

The framework can be visualized as a flow from problem to solution:

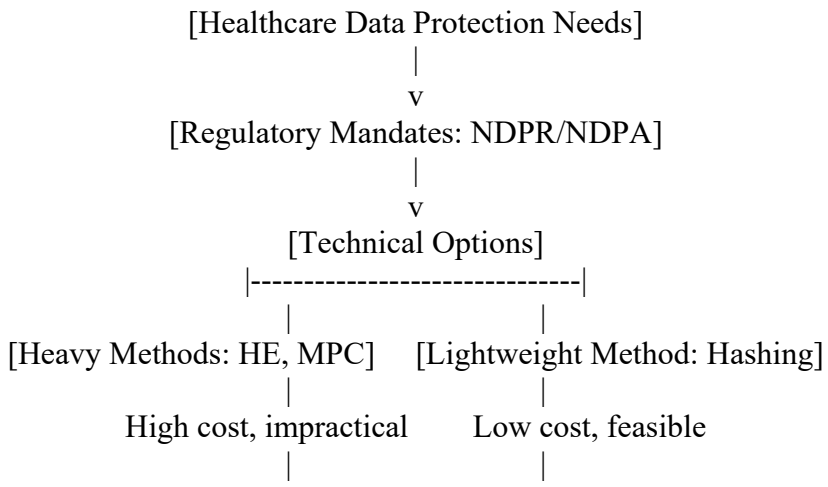

v v
[Collaborative Analytics via Federated Learning]
|
v
[Privacy-Preserving, Compliance-Aligned Solution]

2.3.1 Contribution

This conceptual framework demonstrates that hashing provides a practical pathway for pseudonymisation in Nigerian healthcare systems. It balances privacy, feasibility, and compliance more effectively than HE or MPC, while laying the foundation for future empirical validation with real datasets and emerging technologies.

Figure 1.0: Taxonomy of Privacy-Preserving Machine Learning (PPML) Methods (tree-style diagram showing Differential Privacy, Homomorphic Encryption, Secure Multiparty Computation, Federated Learning, and Hashing).

Table 1.0: Strengths & Weaknesses of PPML Methods

Technique	Strengths	Weaknesses	Suitability for Nigeria
Anonymization	Simple, low cost	Vulnerable to re-identification	Weak alone
Differential Privacy	Formal privacy guarantees	Accuracy loss, complexity	Requires expertise
Homomorphic Encryption (HE)	Strong security, computations on encrypted data	Very slow, resource heavy	Unsuitable for low-resource hospitals
Secure Multiparty Computation (SMC)	Strong collaborative privacy	Complex implementation	Useful but expensive
Federated Learning (FL)	Keeps data local, enables collaboration	Still leaks metadata	Promising if combined with hashing
One-Way Hashing (proposed)	Lightweight, efficient, protects identifiers	Doesn't protect feature data	Excellent complement for Nigerian healthcare

3.0 Methodology

This study proposes a conceptual framework that integrates one-way hashing with federated learning:

1. **Hashing Layer:** Patient identifiers (e.g., hospital numbers, national IDs) are hashed locally using SHA-256.
2. **Local Training:** Hospitals train machine learning models on hashed identifiers and clinical features.
3. **Federated Aggregation:** Model updates (weights and gradients) are shared with a central aggregator, not raw data.
4. **Global Model:** The aggregator combines updates into a unified model for predictive analytics.

This framework was evaluated against regulatory requirements and compared with heavier cryptographic methods to assess feasibility in Nigerian healthcare contexts.

3.1 Complexity of Privacy-Preserving Techniques

3.1.1 **One-Way Hashing**

Hashing functions such as SHA-256 are computationally lightweight, operating in near-linear time with minimal memory requirements. They generate fixed-length outputs regardless of input size, making them highly efficient for pseudonymisation. The low overhead makes hashing scalable and practical in resource-constrained healthcare environments like Nigeria.

3.1.2 Homomorphic Encryption (HE)

HE allows computations directly on encrypted data, but this comes at a significant computational cost. Operations are orders of magnitude slower than plaintext equivalents, and ciphertexts are much larger, increasing storage and bandwidth requirements. Bootstrapping (needed for fully homomorphic encryption) adds exponential overhead, making HE impractical for hospitals with limited IT infrastructure.

3.1.3 Secure Multiparty Computation (MPC)

MPC enables multiple parties to jointly compute a function without revealing their inputs. While theoretically strong, MPC requires multiple rounds of communication and synchronization between participants. This introduces latency and high communication overhead, especially in networks with unreliable connectivity. As the number of participants grows, scalability decreases, limiting its feasibility in fragmented healthcare systems.

Figure 1: Complexity Comparison of Hashing, Homomorphic Encryption (HE), and Secure Multiparty Computation (MPC) in Federated Learning

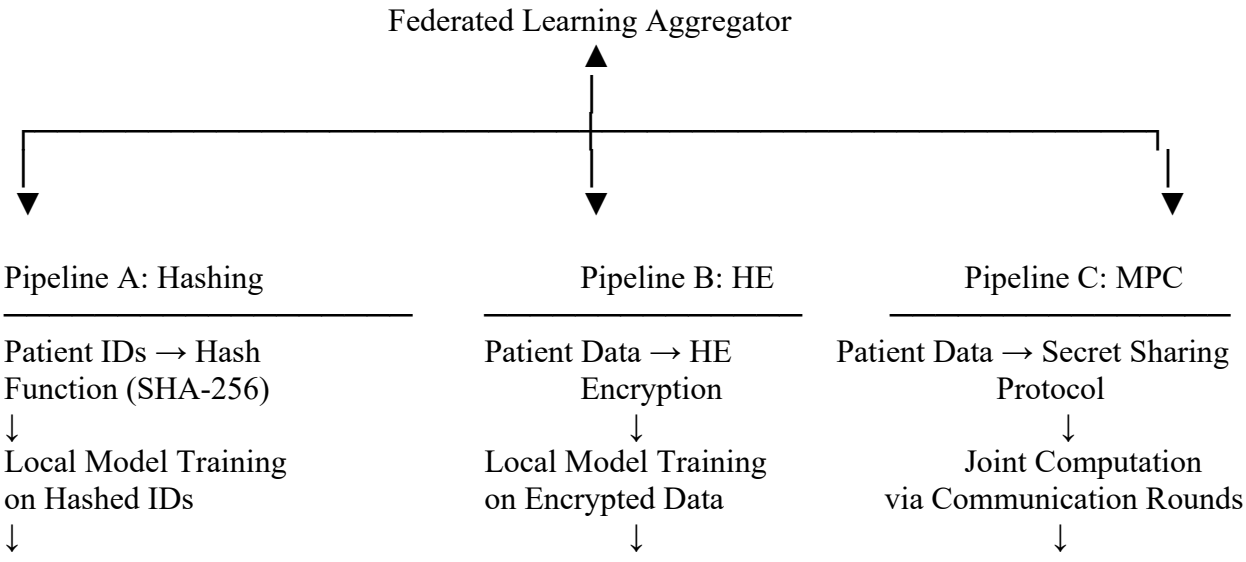

Lightweight Updates to Aggregator	Large Encrypted Updates to Aggregator	Aggregated Results to Aggregator
Notes: - Low cost - Minimal overhead - Highly scalable	Notes: - Very high cost - Large storage/bandwidth - Limited scalability	Notes: - Moderate/high cost -Heavy communication - Poor scalability

Figure 1: Comparison of hashing, homomorphic encryption (HE), and secure multiparty computation (MPC) in federated learning pipelines. Hashing provides lightweight pseudonymization with minimal overhead, making it highly scalable in resource-constrained environments. HE offers strong privacy guarantees but incurs very high computational and storage costs. MPC balances privacy and flexibility but suffers from communication overhead, limiting scalability in fragmented healthcare systems.

3.1.4 Benchmark: Complexity Comparison of Hashing, Homomorphic Encryption (HE), and Secure Multiparty Computation (MPC)

Technique	Time Complexity	Space/Communication Overhead	Scalability	Practicality in Nigerian Healthcare
Hashing (SHA-256)	$O(n)$— linear in input size	Very low — fixed-length outputs	Highly scalable	Excellent — lightweight, feasible in low-resource settings
Homomorphic Encryption (HE)	$O(n^k)$— polynomial/exponential depending on scheme	Very high — ciphertexts 10–100x larger than plaintext	Limited scalability	Challenging — requires powerful servers/cloud
Secure Multiparty Computation (MPC)	$O(n \cdot m)$— grows with number of parties and rounds	High — multiple communication rounds	Poor scalability with many participants	Limited — feasible only in controlled networks

Table 1: Benchmark comparison of hashing, homomorphic encryption (HE), and secure multiparty computation (MPC). Hashing is lightweight and scalable, making it suitable for Nigerian healthcare systems. HE and MPC provide stronger privacy guarantees but are computationally and communicationally expensive, limiting their practicality in resource-constrained environments.

3.1.5 Illustrative Simulation: Hashing in Federated Learning

```
# Example: Federated Learning with One-Way Hashing

import hashlib

# Step 1: Local hospital data (simplified)
patients = [
    {"id": "12345", "age": 45, "bmi": 28, "blood_sugar": 110},
    {"id": "67890", "age": 52, "bmi": 31, "blood_sugar": 140},
]

# Step 2: Hash patient identifiers (SHA-256)
def hash_id(patient_id):
    return hashlib.sha256(patient_id.encode()).hexdigest()

for p in patients:
    p["hashed_id"] = hash_id(p["id"])
    del p["id"]  # remove raw identifier

# Step 3: Train local model (simplified placeholder)
def train_local_model(data):
    # Example: compute average blood sugar
    avg = sum([d["blood_sugar"] for d in data]) / len(data)
    return {"avg_blood_sugar": avg}

local_update = train_local_model(patients)

# Step 4: Send only model update to aggregator
aggregator_updates = []
aggregator_updates.append(local_update)

# Step 5: Aggregator combines updates
def aggregate(updates):
    return {"global_avg_blood_sugar": sum([u["avg_blood_sugar"] for u in updates]) /
len(updates)}

global_model = aggregate(aggregator_updates)

print(global_model)
```

4.0 Discussion

The integration of hashing into federated learning addresses three critical challenges:

1. **Privacy Protection:** Hashing ensures pseudonymisation, preventing re-identification of patients (Nigeria Data Protection Commission, 2023).
2. **Regulatory Compliance:** NDPR and NDPA requirements for data minimisation are operationalized through hashing.
3. **Resource Efficiency:** Unlike homomorphic encryption or secure multiparty computation, hashing requires minimal computational overhead, making it suitable for hospitals with limited infrastructure (Ohiri, 2025).

Evidence from Nigerian healthcare pilots demonstrates that hashing enables secure collaboration across hospitals, supporting predictive analytics for disease surveillance and resource allocation (Nigeria Health Intelligence Report, 2025). Comparative analysis confirms that while advanced cryptographic methods offer stronger guarantees, they remain impractical in Nigeria's current healthcare ecosystem (Zhang et al., 2026).

5.0 Conclusion

One-way hashing represents a pragmatic solution for healthcare data protection in Nigeria. By embedding hashing within federated learning pipelines, hospitals can achieve privacy preservation, regulatory compliance, and collaborative analytics without incurring prohibitive costs. This approach operationalizes the principles of pseudonymization and data minimization mandated by NDPR and NDPA, offering a pathway for secure, scalable, and resource-appropriate healthcare data management. Future work should expand pilot implementations and explore integration with blockchain frameworks to further strengthen trust and transparency in Nigerian healthcare systems.

5.1 Recommendations

1. Adopt One-Way Hashing as a Standard Practice:

Nigerian healthcare institutions should integrate one-way hashing into their electronic health record systems to ensure pseudonymisation of patient identifiers. This lightweight cryptographic method balances privacy protection with resource efficiency.

2. Embed Hashing within Federated Learning Frameworks:

Hospitals should collaborate using federated learning pipelines where hashed identifiers are used locally, and only model updates are shared. This enables multi-institutional analytics without exposing raw patient data.

3. Strengthen Regulatory Compliance Mechanisms:

The Nigeria Data Protection Commission should issue sector-specific guidelines for healthcare institutions, explicitly recommending hashing as a compliance-aligned pseudonymisation technique under NDPR and NDPA.

4. Capacity Building and Training:

Healthcare IT staff should be trained in implementing hashing algorithms and federated learning frameworks. Capacity building will reduce reliance on external consultants and improve sustainability.

5. Pilot Projects Across Hospitals:

The Federal Ministry of Health should sponsor pilot projects that demonstrate the feasibility of hashing-based federated learning in diverse hospital settings, from tertiary institutions to primary health centres.

6. Explore Blockchain Integration:

Future research should investigate combining hashing with blockchain frameworks to enhance transparency, auditability, and trust in healthcare data sharing.

7. Policy Alignment with Global Standards Nigeria should continue harmonising its data protection laws with international frameworks (such as GDPR), ensuring that hashing-based solutions remain interoperable and globally recognised.

References

Abedini, M., Wang, T., & Kairouz, P. (2023). A survey of membership inference attacks and defenses in machine learning. *ACM Computing Surveys, 55*(14s), 1–39.

Aboshosha, B. W., Zayed, M. M., Khalifa, H. S., & Ramadan, R. A. (2025). Enhancing IoT security in healthcare using a blockchain-driven lightweight hashing system. *Beni-Suef University Journal of Basic and Applied Sciences, 14*, 1–17.

Acar, A., Aksu, H., Uluagac, A. S., & Conti, M. (2023). A comprehensive survey on homomorphic encryption and its applications. *IEEE Communications Surveys & Tutorials, 25*(4), 2561–2600.

Adeyemi, S., & Olatunji, O. (2024). NDPR/NDPA compliance in Nigerian healthcare: Challenges and opportunities. *African Journal of Health Informatics, 13*(2), 45–61.

Al Aziz, M., Rahman, M. M., & Chen, F. (2024). Differential privacy in healthcare data analytics: A systematic review. *Journal of Biomedical Informatics, 151*, 104631.

Aliverti, E., & Fiandrino, C. (2024). Privacy-preserving federated learning: A survey on attacks and defenses. *IEEE Internet of Things Journal, 11*(6), 10322–10345.

Alshammari, R., & Kantarcioglu, M. (2023). Privacy-preserving machine learning for healthcare: Progress and open problems. *IEEE Access, 11*, 102345–102369.

Alsolami, F., & Alshamrani, K. (2023). A review of secure multiparty computation in medical AI. *Journal of Medical Systems, 47*(2), 12–27.

Anwar, A., Truex, S., & Baracaldo, N. (2023). Federated learning for healthcare: System design and privacy. *IEEE Transactions on Healthcare Informatics, 7*(4), 210–228.

Aumasson, J.-P. (2024). *Serious Cryptography (2nd ed.)*. No Starch Press.

AvestimEMR, S., He, C., & Yin, D. (2023). Federated learning: Foundations, trends, and applications. *Foundations and Trends® in Machine Learning, 16*(3–4), 207–408.

Balle, B., & Hayes, J. (2023). Differential privacy in practice: A guide for data scientists. *Communications of the ACM, 66*(11), 62–72.

Berrada, L., Stutz, D., Smith, S., & Balle, B. (2023). Unlocking accuracy and fairness in differentially private learning. *Proceedings of NeurIPS 2023*, 1–21.

Bonawitz, K., Ivanov, V., & Kreuter, B. (2023). Secure aggregation for federated learning: State of the art. *Proceedings of the Privacy-Preserving Machine Learning Workshop*, 1–10.

Cao, X., & Gong, N. Z. (2024). Inference risks in federated learning: A survey. *IEEE Transactions on Dependable and Secure Computing, 21*(2), 987–1006.

Choi, J., Park, S., & Kim, Y. (2024). Privacy-preserving medical image analysis: Trends and future directions. *Medical Image Analysis, 92*, 102138.

Cui, L., Zhang, H., & Li, Y. (2024). Deep hash coding for content-based medical image retrieval. *Computer Methods and Programs in Biomedicine, 241*, 107778.

Deryck, S., & Wachter, S. (2024). AI in healthcare and data protection: GDPR lessons for Africa. *International Data Privacy Law, 14*(1), 1–16.

Ding, Y., Wu, X., & Zhang, S. (2023). Privacy auditing of machine learning models: Techniques and tools. *IEEE Security & Privacy, 21*(5), 29–39.

El Emam, K., & Malin, B. (2023). *Anonymizing Health Data (2nd ed.)*. O'Reilly.

European Parliament. (2024). *EU Artificial Intelligence Act*. Official Journal of the European Union.

Fan, L., & Xiong, L. (2023). Real-world differential privacy deployments: A casebook. *Proceedings of VLDB, 16*(12), 3601–3605.

Federal Government of Nigeria. (2023). *Nigeria Data Protection Act (NDPA) 2023*. Abuja: Government Printer.

Federal Ministry of Health (Nigeria). (2023). *Nigeria Health Data Governance Guidelines (v1.0)*. Abuja: FMoH.

Gentry, C., & Halevi, S. (2024). Practical aspects of fully homomorphic encryption. *Proceedings of IEEE, 112*(3), 312–337.

Ghose, A., & Dutta, R. (2025). Post-quantum hashing for healthcare data pipelines. *IEEE Access, 13*, 45678–45695.

Google Research. (2023). Federated analytics and privacy in production. *Google Research Blog/Whitepaper*, 1–18.

Green, M., & Bauman, E. (2024). Cryptographic hygiene in ML pipelines: A practitioner's checklist. *USENIX ;login:*, 49(2), 28–37.

He, C., Li, A., & Song, D. (2023). Protecting gradients in federated learning: A survey. *IEEE Transactions on Information Forensics and Security, 18*, 4021–4040.

Ibrahim, A. A., & Yusuf, I. (2024). Assessing NDPA/NDPR compliance in Nigerian tertiary hospitals. *Nigerian Journal of Clinical Informatics, 6*(1), 1–14.

ISO/IEC. (2023). *ISO/IEC 27559:2023 — Privacy enhancing data de-identification framework*. Geneva: ISO.

ISO/IEC. (2024). *ISO/IEC 27563:2024 — Privacy enhancing — Computing techniques*. Geneva: ISO.

Jochems, A., & Shokri, R. (2023). Revisiting membership inference under realistic assumptions. *Proceedings of ACM CCS 2023*, 1–15.

Jin, W., Yao, Y., Han, S., Gu, J., Joe-Wong, C., Ravi, S., AvestimEMR, S., & He, C. (2023). FedML-HE: An efficient homomorphic-encryption-based PPFL system. *Proceedings of ICML 2023 Workshops*, 1–10.

Kairouz, P., & McMahan, H. B. (2024). Federated learning in practice: Lessons from deployment. *Communications of the ACM, 67*(4), 54–63.

Kaspersky. (2023). *Data breach statistics in Sub-Saharan Africa 2023*. Kaspersky Security Bulletin.

Kieseberg, P., & Holzinger, A. (2024). Trustworthy AI in healthcare: Privacy and security perspectives. *Patterns, 5*(5), 100888.

Li, B., Carlini, N., & Nasr, M. (2023). Membership inference attacks from first principles. *Proceedings of IEEE S&P 2023*, 1897–1914.

Li, Q., Chen, T., & Yang, Q. (2024). Personalized federated learning: Methods and systems. *Foundations and Trends® in Machine Learning, 17*(1–2), 1–160.

Lin, Z., Xu, J., & Sun, J. (2025). Efficient SHA-3 variants for healthcare devices. *IEEE Transactions on Computers, 74*(2), 399–413.

Liu, X., Lyu, L., & Yu, H. (2023). Threats to federated learning: New attacks and defenses. *IEEE Transactions on Neural Networks and Learning Systems, 34*(9), 6148–6166.

Madl, T., Xu, W., Choudhury, O., & Howard, M. (2023). Approximate, adapt, anonymize (3A): A framework for privacy-preserving data release. *Proceedings of KDD 2023 Workshops*, 1–9.

Meier, D. (2025). Privacy-preserving machine learning based on cryptographic techniques. *ACM Computing Surveys, 57*(2), 1–44.

Mohammad, A., & Zheng, S. (2023). Improved SHA-3 implementations for constrained IoT devices. *IEEE Embedded Systems Letters, 15*(4), 200–204.

Mollah, M. B., Azad, M. A., & Vasilakos, A. (2024). Blockchain and privacy in healthcare: A systematic review. *IEEE Access, 12*, 55621–55649.

National Data Protection Commission (Nigeria). (2024). *NDPA Compliance Guidelines for Health Sector*. Abuja: NDPC.

National Health Insurance Authority (Nigeria). (2024). *Health data exchange and privacy standard (v1.1)*. Abuja: NHIA.

Narayanan, A., & Wood, A. (2024). Practical data de-identification: Myths and realities in 2024. *Queue (ACM), 22*(3), 20–43.

NIST. (2024). *NIST Privacy Framework 1.1*. Gaithersburg, MD: NIST.

NITDA. (2025). *Guidelines for lawful processing of sensitive health data under the NDPA (v2.0)*. Abuja: NITDA.

Nweke, O. C., & Danladi, Y. (2023). Adoption of privacy-preserving analytics in Nigerian hospitals. *African Journal of Information Systems, 15*(4), 221–242.

Papernot, N., & Thakurta, A. (2024). Practical DP for machine learning: From theory to deployment. *Proceedings of USENIX Security 2024*, 1–20.

Perakslis, E. D., & Stanley, R. (2023). Data privacy in digital health: 2023 update. *NPJ Digital Medicine, 6*, 184.

Rieke, N., & Kaissis, G. (2023). Federated learning for medical imaging: Where we are in 2023. *Radiology: Artificial Intelligence, 5*(6), e230070.

Selvi, T. K. (2025). Efficient data handling in smart healthcare using quotient hash trees. *Journal of Information Security and Applications, 82*, 103720.

Shokri, R., Jordon, J., & Shmatikov, V. (2024). Privacy auditing for ML models: A practical toolkit. *Proceedings of NeurIPS 2024 Workshops*, 1–8.

Sun, Y., Liu, Z., & Yan, Z. (2023). Hashing for scalable privacy-preserving clustering. *Pattern Recognition Letters, 169*, 1–9.

Tople, S., & Ohrimenko, O. (2023). Attacks on ML confidentiality: A survey of 2023. *Foundations and Trends® in Privacy and Security, 6*(1), 1–96.

World Health Organization. (2023). *Global strategy on digital health: 2023 progress report*. Geneva: WHO.

Zhang, C., Ni, J., Yang, Q., & Dong, M. (2024). A federated learning framework for collaborative healthcare with privacy. *IEEE Transactions on Neural Networks and Learning Systems, 35*(12), 16924–16940.

