Secure Multimodal Storage and Multi-Authority Access Control for Personal Health Records with AI-Driven Analytics

Mr. Indraneel P. Mane

Student, M.Tech CSE (Data Science), D. Y. Patil Agriculture and Technical University, Talsande, Kolhapur, Maharashtra, India

indraneelmane@gmail.com

Mr. Naresh A. Kamble

Assistant Professor, D. Y. Patil Agriculture and Technical University, Talsande, Kolhapur, Maharashtra, India

Abstract: Personal Health Records (PHRs) enable patients to gather and manage their own health information from various healthcare sources, including hospitals and laboratories, but current PHR systems typically treat secure storage, access control, and analytical capability as separate concerns. Single-authority encryption schemes carry a key-escrow problem, fixed access policies create information silos across institutions, and analytics components are rarely designed to operate within an encrypted and access-controlled pipeline. This paper reviews 37 peer-reviewed sources on encrypted cloud storage, multi-authority attribute-based encryption, and AI-based health analytics, and proposes a system that links these three concerns. The proposed architecture combines encryption technologies such as AES, RSA, and attribute-based encryption (ABE) with Multi-Authority Attribute-Based Access Control (MA-ABAC) to give patients, healthcare providers, and laboratories granular control over access to patient information, together with an artificial intelligence (AI) analytics layer for patient-criticality classification and role-specific clinical summarization. The review finds that encrypted storage, distributed access control, and federated or multimodal learning are each fairly well developed as stand-alone services, but no published system combines all three in one deployable architecture. This gap is identified in the paper, which outlines the goals, approach, and evaluation strategy for the system to be built.

Keywords: Personal Health Records; multimodal data; multi-authority access control; attribute-based encryption; artificial intelligence; health analytics.

1. INTRODUCTION

The process of healthcare delivery is moving away from paper records to Electronic Health Records (EHRs) held by institutions and Personal Health Records (PHRs) held by patients. A PHR enables a person to gather, keep, and disseminate health data such as lab results, diagnostic images, and clinical notes from hospitals, clinics, diagnostic centers, and research centers[1]. In contrast to the institution-focused EHR, this patient-centric approach fosters inter-provider collaboration for diagnosis, treatment, and research, without removing the patient from control of their own data[1,2]. Efforts by the government to increase the rate of EHR adoption, including the American Recovery and Reinvestment Act[3], have led to greater institutional EHR adoption, and health information exchange between institutions has been linked to improvements in the quality of patient care received[2], which has given rise to greater interest in patient-controlled alternatives like the PHR.

However, three interrelated challenges must be addressed to achieve effective implementation of PHR systems. First, health information is sensitive, and transferring it to third-party cloud servers makes patients vulnerable to data breaches and unauthorized access - studies indicate that incidents have increased in scale and number over the last decade[4], and defining the security and privacy requirements these systems must meet remains an active research area[5]. Second, PHR data is multimodal: it contains structured numerical data such as laboratory parameters, unstructured textual data such as clinical notes, medical images such as X-rays, MRI, and CT scans, physiological signals such as ECG and EEG recordings, genomic data[6], and streams of data from wearable sensors. Each of these formats requires a different approach to storage, encryption, and analysis, and linking them into a unified picture of a patient's health is challenging[7,8]. Third, access to PHR data must involve more than one organization - for example, a hospital, an independent diagnostic lab, and an insurance provider - and current single-authority access-control schemes handle this poorly: an authority that issues all the encryption keys also has the ability to decrypt all the data, a problem generally termed key escrow[9,10].

Two lines of research bear directly on these challenges. Attribute-Based Encryption (ABE) and its multi-authority extensions have been proposed as a fine-grained access-control mechanism that removes the need for a single trusted authority[10,11]. Separately, AI and machine-learning techniques have been used with multimodal clinical data to assist with anomaly detection, automated summarization, and disease-risk prediction[7]. In most published systems, however, these two lines of work are developed separately, and cybersecurity assessments of the healthcare sector continue to report a gap between the technical capabilities available and the protections actually deployed[12]. This paper is based on a project synopsis, a detailed project report, and a literature review of 37

peer-reviewed sources; it presents the problem, summarizes the literature, and outlines a system that integrates encrypted multimodal storage, multi-authority access control, and AI-based analytics in one package.

2. LITERATURE REVIEW

2.1. Secure storage and encryption of multimodal health data

Symmetric encryption schemes such as AES-256 are efficient for storing bulk data, while asymmetric schemes such as RSA offer good key-exchange support; the two are typically combined to balance computational cost against key-management complexity[13]. In most such approaches, an entire record is either encrypted or not at all, which makes selective sharing difficult in a multi-stakeholder environment. Attribute-Based Encryption (ABE) grew out of earlier work on Fuzzy Identity-Based Encryption (FIBE)[14] and was later extended into ciphertext-policy schemes[15]. ABE does not rely on a user's identity but instead on attributes - for example, cardiologist with hospital privileges - so a patient can specify an access policy and let anyone whose attributes match it decrypt the record, without identifying that person in advance.

Zhang, Liu and Khan (Ref. 9) proposed a scalable solution for PHR sharing using ABE, dividing users into a number of security domains to support large-scale deployment, and proposed a break-glass mechanism that provides access when patient consent cannot be obtained. To limit computing demand on user equipment, decryption was delegated to a semi-trusted cloud server, and the scheme was shown to resist collusion attacks while offering fast encryption and decryption. However, it relies on a single authority for key generation, does not address anonymous authentication, and does not support image-based modalities[9]. Patel, Singh and Reddy (Ref. 16) addressed the data-modeling side of the same problem, proposing a common metadata format for structured, semi-structured, and unstructured medical data, implemented using machine-learning techniques to extract and categorize metadata and made compatible with the interoperable HL7 and FHIR standards. Their evaluation showed efficient retrieval across data types, but the paper does not discuss encryption or access control beyond a high level and does not assess the computational load of handling multimodal data in real time[16].

2.2. Multi-authority access control

Several studies address how the trust concentrated in a single authority can be distributed. Chase (Ref. 11) introduced a multi-authority ABE construction that delegates key-issuing responsibilities to independent attribute authorities - for example, hospitals or an insurance company - which Wang, Liu and Wu (Ref. 10) extended into Multi-Authority Attribute-Based Encryption (MA-ABE). No modification of existing keys or ciphertexts is required, and the reported computational overhead of adding a tenth authority is only 1–2 percent, reflecting near-linear cost as authorities are added. The scheme is secure under standard cryptographic assumptions and resists collusion attacks, but it does not address anonymous authentication, and a user's identity may still be revealed during an access request[10].

Other work covers related aspects of access management without fully addressing the multi-authority problem. In hierarchical ABE, both a user and the user's attributes are automatically escalated on promotion, and revocation does not require existing data to be re-encrypted; however, a single root authority remains, and the hierarchy becomes impractical to maintain beyond five to ten levels[17]. Liang et al. (Ref. 18) presented a revocation mechanism using versioned attribute keys that produces fixed-size ciphertexts, though it does not extend to multiple authorities. Sharma and Kumar (Ref. 19) proposed a ciphertext-policy ABE scheme in which decryption is carried out by the central authority to reduce the computational burden on mobile and wearable devices, reporting faster decryption than earlier ABE schemes while still depending on a single central authority. Anonymous authentication - allowing a user to prove that their attributes satisfy an access policy without revealing their identity - receives little attention across this body of work; a survey of cryptographic solutions for healthcare data sharing in cloud computing by Zhang and Xue reaches a similar conclusion[20].

2.3. AI-driven analytics for multimodal health data

Few studies use clinical data from multiple modalities directly in machine learning. In one study, separate convolutional neural networks for each imaging modality (CT, MRI, and X-ray) were used to predict cardiovascular risk; adding wearable-sensor data such as heart-rate variability, physical activity, and sleep increased the area under the curve from 89.3 percent to 94.2 percent[7]. Recurrent neural networks and Long Short-Term Memory (LSTM) networks, which retain information over longer time series than a standard recurrent network, have been used where data has a strong temporal component - for example, repeated lab measurements indicating worsening kidney function or persistent hypertension.

Federated learning trains such models without transporting patient information to a central server: each participating site trains the model locally, and only the resulting model parameters are aggregated[21]. Singh and Reddy (Ref. 26) showed that a federated LSTM model for mortality prediction achieved an AUC of 91.8 percent - close to a centralized model trained on the same data (AUC = 93.2 percent) - and outperformed models trained at individual hospitals (AUC = 89.1 percent on data from hospitals not represented in training). McMahan et al. identified communication-efficient methods for federated training that use only about 90 percent of the bandwidth of a naive parameter-sharing approach[22], and a survey of over 200 papers by Kairouz et al. found that federated approaches to healthcare-data training almost always generalized better than single-institution training[23]. Earlier work on

phenotyping from electronic health records and on wearable sensors for atrial-fibrillation monitoring[24,25] suggests that this kind of analytics can extend naturally to PHR data made available under proper authorization. Figure 1 summarizes these comparisons.

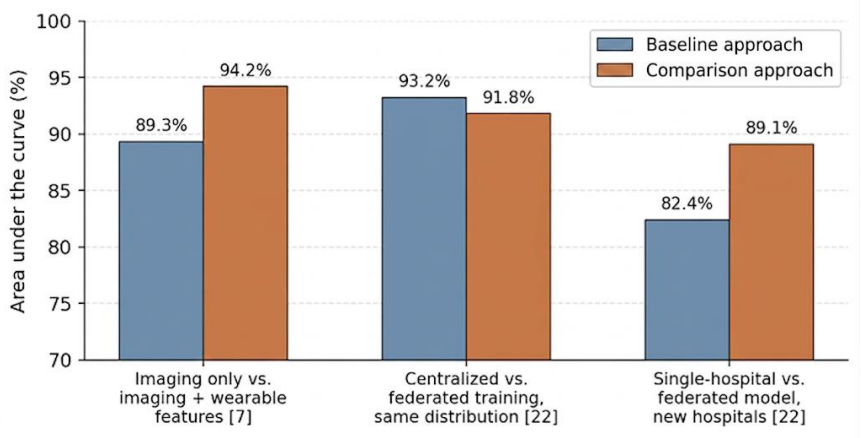

Figure 1. Comparative diagnostic performance reported for multimodal and federated learning approaches, relative to single-modality or single-institution baselines.7,26

2.4. Standardization, usability, and regulatory context

Technical solutions alone are not sufficient to ensure adoption. The healthcare industry has historically been slow to standardize privacy and consent handling; the most recent release of FHIR introduces explicit privacy and consent models[27], but institutional adoption requires investment in system updates and staff training and remains far from widespread. Independent of the particular technology, coordination between organizations, regions, and countries has been identified as a persistent barrier[28]. Usability studies of health information systems[29] and broader human-factors studies of health-system redesign[30] show that, in practice, workflow design is as important as the underlying technology - if not more so - when clinicians' time is limited. Regulatory requirements such as the General Data Protection Regulation[31] and the HIPAA Privacy Rule[32] provide further motivation for the protections outlined above, but meeting regulatory requirements alone does not solve the technical integration challenge this paper addresses.

2.5. Integration challenges and the research gap

The research reviewed above has advanced secure storage, multi-authority access control, and AI-based data analysis largely as separate problems, and the tension between these concerns is not addressed by many published systems. Computing directly on encrypted data without prior decryption remains infeasible at clinical scale: homomorphic encryption allows computation over ciphertexts but introduces several orders of magnitude of computational overhead, making it unsuitable for routine analysis[33], and secure multi-party computation preserves privacy during joint computation but typically requires more network round-trips than real-time clinical use allows[34]. A third, more common approach decrypts data within a discrete, audited environment, shifting the security guarantee from cryptography to operational controls such as access logging and network segmentation - controls that hardware-based secure enclaves attempt to reinforce[35]. Population-level statistics can often be released without exposing individual records; differential privacy is one method used to release aggregate statistics with calibrated statistical noise[36,37].

Table 1 summarizes the strengths and limitations of the main schemes discussed above.

Table 1. Summary of key studies on secure storage, multi-authority access control, and AI-driven analytics for Personal Health Records.

Ref.	Study	Focus	Key strength	Main limitation
9	Zhang, Liu & Khan (2018)	Scalable ABE-based PHR sharing	Security domains; break-glass emergency access; collusion resistant	Single-authority key generation; no anonymous authentication
16	Patel, Singh & Reddy (2022)	Heterogeneous data modeling	Unified representation; interoperable with HL7/FHIR	Encryption and access control specified only at a high level
10	Wang, Liu & Wu (2022)	Multi-authority ABE	Near-linear scaling as authorities are added; provably secure	No anonymous authentication
17	Zhou & Huang (2021)	Hierarchical ABE with revocation	Automatic attribute inheritance; revocation without re-encryption	Single root authority; hierarchy hard to manage at scale

18	Liang et al. (2021)	Versioned-key revocation	Constant-size ciphertexts; sub-millisecond revocation	Does not extend to multiple authorities
19	Sharma & Kumar (2020)	CP-ABE with outsourced decryption	Reduced decryption time on resource-constrained devices	Relies on a single central authority
7	Kumar et al. (2023)	Multimodal CNN for cardiovascular risk	Higher accuracy when imaging is combined with sensor data	No encryption or access-control component evaluated
26	Singh & Reddy (2022)	Federated LSTM for mortality prediction	Comparable accuracy without centralizing data; better generalization	Not evaluated alongside an encrypted storage pipeline

Across the reviewed literature, no single system combines secure multimodal storage, multi-authority access control, anonymous authentication, and AI-driven analytics within one architecture. Individual studies address one or two of these requirements at a time: ABE-based storage with limited modality support[9], heterogeneous data modeling without strong security guarantees[16], multi-authority encryption without anonymous authentication[10], or AI analytics evaluated separately from any encrypted pipeline[7,26]. Proposed reference architectures for decentralized personal health data spaces point toward this kind of integration but remain at a conceptual stage. This gap, identified consistently across the project synopsis, the project report, and the literature review on which this paper draws, motivates the system proposed in the remainder of this paper.

3. PROBLEM STATEMENT

Existing PHR systems generally cannot manage data privacy, secure cross-institutional sharing, and analytics within one consistent framework. Centralized architectures introduce a single point of failure and limit the fine-grained, decentralized access control that distributed healthcare settings require. Most current systems cannot handle multimodal data - including clinical text, medical images, and sensor signals - within a unified storage and security model, which fragments the resulting data and reduces its analytical value[16]. Single-authority encryption schemes retain the key-escrow problem described above, since the authority that issues keys also retains the capacity to decrypt all data it protects[9]. Where multi-authority schemes have been proposed, they generally omit anonymous authentication, leaving a user's identity exposed during an access request[10], and none of the systems reviewed integrates AI-driven analytics operating within a secure, access-controlled environment.

These limitations point to the absence of a framework that supports decentralized, multi-authority access control, preserves data confidentiality and patient privacy, and applies AI-based analytics for prediction and decision support, all within a single system.

4. RESEARCH OBJECTIVES

This study sets out three objectives:

(1) To develop a cloud-based, multimodal Personal Health Record storage system that handles structured text records, medical images, and signal data, with encryption, integrity verification, and retrieval suited to a distributed cloud environment.

(2) To implement a Multi-Authority Access Control mechanism based on distributed attribute-based encryption, supporting fine-grained, role-based, policy-driven permissions, while removing the key-escrow problem and distributing trust among independent authorities.

(3) To integrate AI-driven analytics within the secure framework for early disease prediction, patient-criticality assessment, and personalized treatment recommendations, ensuring that this processing occurs within a privacy-aware, access-controlled environment.

5. PROPOSED SYSTEM

5.1. System architecture

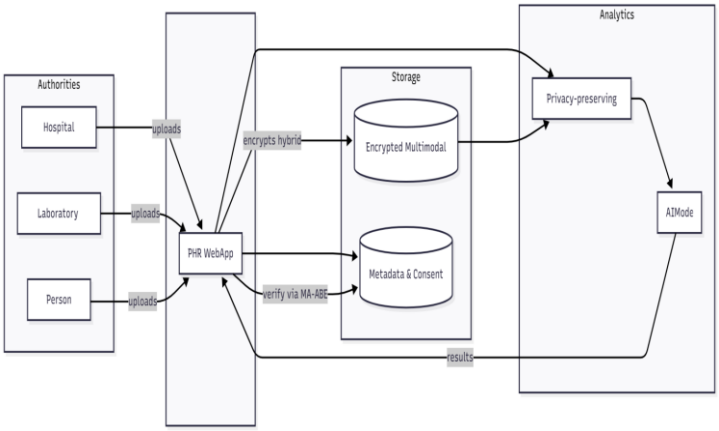

Figure 2. Conceptual data flow from contributing authorities through encrypted multimodal storage to the AI analytics layer, adapted from the project synopsis.

The proposed system is a web-based PHR application organized into five layers, summarized in Table 2 and illustrated in Figures 2 and 3.

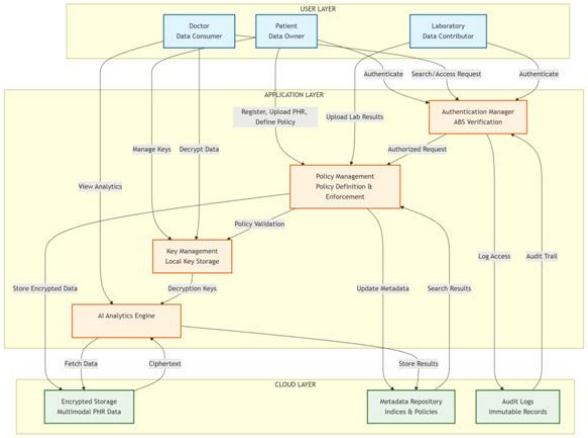

Figure 3. Three-tier layered architecture of the proposed system, adapted from the project report.

Table 2. Layers of the proposed system architecture.

Layer	Components	Responsibilities
Data sources	Patients, hospitals, laboratories, attribute authorities	Multimodal data contribution; attribute issuance
Secure storage	Hybrid AES, RSA, and attribute-based encrypted repository	Confidential persistence of encrypted records
Access control	MA-ABAC policy engine	Fine-grained, hierarchical permission enforcement
AI analytics	Criticality classifier; role-based summarization model	Cohesive analysis of decrypted multimodal data
User interface	Web-based dashboard	Authentication, visualization, role-based interaction

The Data Sources layer collects multimodal data from patients and healthcare providers, including hospitals and laboratories. The Secure Storage layer stores encrypted records using a hybrid scheme that combines AES for bulk encryption with RSA or attribute-based encryption for key management. The Access Control layer applies Multi-Authority Attribute-Based Access Control (MA-ABAC) to issue hierarchical permissions to patients and providers. The AI Analytics layer performs the cohesive analysis described in Section 5.3. The User Interface layer provides a web-based dashboard for authentication and data visualization.

At the center of the access-control design is a Multi-Authority Attribute-Based Encryption scheme. Where a single authority retains the capacity to decrypt all data it protects, MA-ABE distributes that capacity across multiple independent attribute authorities - for

example, a hospital consortium and a national health authority - so that no single entity can compromise the system as a whole. This design choice follows directly from the limitation identified in Section 2.2.

5.2. Operational workflow

The system operates in two phases. To join the system and be able to store attributes, a user registers, and several Attribute Authorities issue private keys corresponding to their role, such as Doctor: Cardiologist or Lab: Pathologist. A patient or a lab generates a record - which may combine text, numeric lab values, and image files - encrypted according to an access policy defined by the lab or the patient, for example "Specialist: Cardiologist AND Role: Doctor" or "Role: Patient AND ID: 123", and uploads it to the cloud together with the access policy. Only the ciphertext is stored on the cloud server, which cannot decrypt it.

In the retrieval phase, a user seeking access to the data - for example, the treating physician - completes an anonymous authentication process using an Attribute-Based Signature, proving possession of attributes issued by an authorized entity without identifying the specific individual. Once authenticated, the cloud server returns the appropriate ciphertext, which the user's client application attempts to decrypt with the attribute keys it holds; decryption succeeds only if the user's attributes satisfy the policy associated with the record. Part of the decryption workload can be delegated to a semi-trusted proxy to reduce computation on the user's own device. When a record is decrypted - or processed within a secure computation environment - the AI analytics module generates output appropriate to the requesting user: a role-based summary and risk flags for a physician, predictive trend information for a patient, and an immediate criticality classification when recent laboratory values call for a clinical response.

5.3. Methodology

The methodology rests on three intertwined activities: secure storage and encryption; multi-authority access control and anonymous authentication; and AI-based analysis for criticality classification and role-based summarization. Two sources of data support development and evaluation. Public repositories, such as the UCI Machine Learning Repository, provide de-identified patient demographics, lab results, vital signs, and diagnoses relevant to training and validating the criticality classifier, along with the structured numerical, unstructured text, and imaging metadata that a PHR system must support. Simulated PHR data is developed for the three end-user roles of patient, doctor, and lab, allowing end-to-end testing of the full workflow - encryption, access-control policies, and analytics - in a controlled setting without using real patient data.

Evaluation follows three lines. Cryptographic analysis examines the computing power required to generate, encrypt, and decrypt keys, the communication overhead between users and the cloud server, resistance to collusion attacks, and the assurance that no single entity can decrypt all stored data. Analytical evaluation assesses the accuracy of criticality classification - implemented with fuzzy logic to accommodate the fuzzy boundaries between health states "Normal," "Less Critical," and "Very Critical" - and the quality of role-based summaries generated by fine-tuned transformer models for different provider types, including cardiologists and general practitioners. System-level evaluation covers end-to-end latency, the latency and accuracy of the AI components across different file sizes and attribute-set complexity, throughput under simultaneous data access, and the storage overhead introduced by encrypting multiple data modalities.

6. CONCLUSION AND FUTURE WORK

This work reviews prior studies on secure cloud storage, multi-authority access control, and AI analytics for Personal Health Records, and proposes an integrated system covering all three. The review shows that each component has reached a mature stage of development on its own: attribute-based encryption has progressed from theoretical concepts to hospital deployments, multi-authority schemes have advanced toward decentralized permission management, and federated and multimodal learning approaches now support clinical prediction without centralizing patient data. However, no system identified in the reviewed literature combines secure multimodal storage, multi-authority access control with anonymous authentication, and AI-driven analytics in a single deployable system - this gap is the contribution the proposed work addresses.

The proposed system presents a unified architecture that shows how multi-authority attribute-based encryption, attribute-based signatures, and AI analytics can interwork within a single application. It directly addresses the integration gap identified in the literature and aims to provide role-specific summaries and criticality alerts that reduce the time clinicians need to review large patient histories.

Several extensions are planned for future work. Building on existing proposals for decentralized health-data architectures, a distributed ledger could record data-access events to create a tamper-resistant audit trail. Federated learning would process data locally at the originating site during AI training rather than consolidating it centrally. Extending the system to support additional modalities such as DICOM imaging, and integrating it with interoperability standards such as FHIR Release 5, would help achieve interoperability with other hospital EHR systems and national health infrastructures.

Acknowledgments

The authors thank the Department of Computer Science and Engineering, D. Y. Patil Agriculture and Technical University, Talsande, Kolhapur, for institutional support during this work. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

REFERENCES

[1] U. Jenter et al., Healthcare Technol. Lett. 11, 112 (2024).

[2] J. H. Grossman and K. L. Kushner, J. Healthcare Finance Adm. 31, 45 (2005).

[3] R. Steinbrook, N. Engl. J. Med. 360, 1057 (2008).

[4] Ponemon Institute, 2017 Cost of Data Breach Study (Ponemon Institute Research Report, 2017).

[5] A. Hosseini et al., BMC Med. Inform. Decis. Mak. 23, 116 (2023).

[6] Genomes Project Consortium, Nature 526, 68 (2015).

[7] A. Kumar et al., IEEE J. Biomed. Health Inform. 27, 1456 (2023).

[8] S. Fong et al., J. Biomed. Inform. 113, 103630 (2021).

[9] R. Zhang, Y. Liu and S. U. Khan, IEEE Trans. Parallel Distrib. Syst. 23, 131 (2018).

[10] G. Wang, Q. Liu and J. Wu, IEEE Trans. Parallel Distrib. Syst. 33, 756 (2022).

[11] M. Chase, Multi-authority attribute based encryption, Theory of Cryptography Conf. (2007), pp. 515–534.

[12] M. Mettler and S. Weis, J. Med. Internet Res. 24, e31801 (2022).

[13] W. Stallings and L. Brown, Computer Security: Principles and Practice, 3rd edn. (Pearson Education, 2015).

[14] A. Sahai and B. Waters, Fuzzy identity-based encryption, Advances in Cryptology–EUROCRYPT 2005 (Springer, 2005), pp. 457–473.

[15] J. Bethencourt, A. Sahai and B. Waters, Ciphertext-policy attribute-based encryption, IEEE Symp. Security and Privacy (2007), pp. 321–334.

[16] S. Patel, R. Singh and K. Reddy, J. Med. Internet Res. 24, e35789 (2022).

[17] L. Zhou and D. Huang, J. Syst. Softw. 163, 110852 (2021).

[18] K. Liang et al., IEEE Access 9, 147258 (2021).

[19] S. Sharma and P. Kumar, IEEE Trans. Inf. Forensics Security 15, 2995 (2020).

[20] R. Zhang and J. Xue, IEEE Access 6, 36503 (2018).

[21] Q. Yang et al., ACM Trans. Intell. Syst. Technol. 10, 1 (2019).

[22] B. McMahan et al., Communication-efficient learning of deep networks from decentralized data, Proc. Int. Conf. Machine Learning (2017), pp. 1273–1282.

[23] P. Kairouz et al., Found. Trends Mach. Learn. 14, 1 (2021).

[24] G. Hripcsak and D. J. Albers, J. Am. Med. Inform. Assoc. 20, e2 (2019).

[25] S. R. Steinhubl et al., J. Am. Coll. Cardiol. 71, 2779 (2018).

[26] R. Singh and K. Reddy, IEEE Access 10, 52365 (2022).

[27] HL7 International, FHIR Release 5: Privacy and Consent Specifications, Version 5.0 (2023).

[28] H. Hyppönen et al., Int. J. Med. Inform. 83, 42 (2014).

[29] P. Chhanabhai and G. Horne, Int. J. Qual. Health Care 27, 338 (2015).

[30] P. Carayon et al., Appl. Ergon. 45, 11 (2015).

[31] Regulation (EU) 2016/679, General Data Protection Regulation (2023).

[32] U.S. Department of Health and Human Services, Summary of the HIPAA Privacy Rule, 45 CFR Parts 160 and 164 (2013).

[33] Z. Brakerski and V. Vaikuntanathan, Fully homomorphic encryption from ring-LWE and security for key dependent messages, Advances in Cryptology–CRYPTO 2011 (Springer, 2011), pp. 505–524.

[34] A. Ben-David et al., Cryptography 4, 12 (2020).

[35] S. Arnautov et al., SCONE: Secure Linux containers with Intel SGX, Proc. 12th USENIX Symp. Operating Systems Design and Implementation (2016), pp. 689–703.

[36] C. Dwork and A. Roth, Found. Trends Theor. Comput. Sci. 9, 211 (2014).

[37] C. Dwork et al., Calibrating noise to sensitivity in private data analysis, Theory of Cryptography Conf. (2006), pp. 265–284.