Vol. 8, No. 2, April 2026

ISI Journal of Information Systems and Informatics

Published By Asosiasi Doktor Sistem Informasi Indonesia

Federated Learning for Privacy-Preserving Sentiment Analysis in Distributed Electronic Health Record Environments: A Systematic Literature Review

Frederick Mlungisi Dandure[1], Belinda Ndlovu[2]

[1,2]Informatics Department, Faculty of Applied Sciences, Master of Science in Big Data, National University of Science and Technology, Bulawayo, Zimbabwe

Received:
January 27, 2026
Revised:
March 10, 2026
Accepted:
April 1, 2026
Published:
April 12, 2026

Corresponding Author:
Author Name*:
Belinda Ndlovu
Email*:
belinda.ndlovu@nust.ac.zw

DOI:
10.63158/journalisi.v8i2.1546

Abstract. Federated learning (FL) has emerged as a privacy-preserving approach for distributed healthcare analytics, yet its application to sentiment analysis of unstructured electronic health record (EHR) narratives remains limited. This systematic review examined the empirical maturity, methodological trends, and governance implications of federated sentiment-aware learning in distributed EHR settings. Following PRISMA 2020, searches were conducted in IEEE Xplore, Scopus, Web of Science, ScienceDirect, and PubMed on January 5, 2026, covering peer-reviewed studies published from January 2021 to January 2026. After screening and eligibility assessment, 29 empirical implementation studies were included in the qualitative synthesis, while conceptual and survey papers were reviewed contextually but excluded from the core analysis. The evidence shows that FL in healthcare is advancing mainly in structured prediction and privacy-preserving infrastructure. By contrast, sentiment-aware learning on unstructured clinical narratives remains at an early stage, with limited implementation and validation. This review distinguishes empirical from conceptual contributions and proposes a governance-aware, literature-derived framework to guide future implementation-focused research.

Keywords: Federated Learning; Electronic Health Records; Clinical Natural Language Processing; Sentiment Analysis; Privacy-Preserving Healthcare Analytics.

1. INTRODUCTION

The increasing use of Electronic Health Records (EHRs) has progressively shifted the healthcare sector toward more data-driven settings [1]. EHRs are extensive digital databases used by healthcare organizations to store patient records, diagnoses, treatments, and clinical notes [2]. Despite these advancements, the widespread adoption of EHR systems has introduced new challenges: patient information is often stored in isolated hospital systems, creating information silos that hinder the large-scale data analysis required for modern research [3].

These challenges have driven growing interest in Artificial Intelligence (AI) and Machine Learning (ML) for healthcare analytics, particularly for improved diagnosis and personalized treatment [4]. In this context, Federated Learning (FL) has emerged as a powerful privacy-preserving approach that enables machine learning model development across multiple institutions without transferring raw data [5]. Institutions develop local models and share only model parameters for central aggregation [6], [7]. FL is especially relevant in healthcare environments where data transfer is strictly regulated, particularly in low- and middle-income countries [7].

Alongside FL, Sentiment Analysis (SA), a component of Natural Language Processing (NLP), provides methods for assessing emotions, attitudes, and subjective impressions conveyed in text [8]. In healthcare, SA has been applied to clinical narratives, nursing records, and patient feedback to extract emotional and psychological dimensions of care. However, most current approaches rely on centralized processing, limiting their applicability in privacy-sensitive, distributed environments [9].

Addressing these limitations requires robust privacy-preserving and governance mechanisms. Homomorphic encryption, secure aggregation, and differential privacy are essential for mitigating threats such as gradient leakage and inference attacks during distributed model training [10]. Blockchain-based EHR systems have also been proposed to improve data stewardship, auditability, and care coordination across institutions [11]. Together, these technical and organizational tools support the development of accountable and trustworthy healthcare AI systems.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Although published work in federated learning, privacy-preserving ML, and healthcare analytics has grown considerably [2], [5], [6], [7], [10], [12], these studies rarely examine the specific challenges of applying FL to sentiment analysis on unstructured clinical EHR narratives. Furthermore, despite ongoing discussions of privacy strategies, few reviews comprehensively evaluate how these methods are implemented and assessed in real-world federated healthcare deployments.

This study addresses a critical gap at the intersection of federated healthcare analytics and clinical NLP. While prior research has advanced FL for structured clinical prediction and privacy-preserving infrastructure, far less attention has been given to sentiment-aware analytics on unstructured clinical narratives, where linguistic variability, privacy sensitivity, and governance complexity intersect. In response, this study conducts a PRISMA 2020-guided systematic literature review to evaluate the empirical maturity, methodological patterns, and governance implications of federated sentiment-aware learning across distributed EHR environments. Drawing on 29 implementation-focused studies, the study contributes by: (i) explicitly distinguishing empirical implementations from conceptual and survey-based work; (ii) identifying sentiment-aware clinical NLP as an under-empiricized and structurally constrained domain; and (iii) synthesizing a governance-aware conceptual framework to guide future deployment-oriented research.

To guide the empirical synthesis, the study is structured around the following research questions:

RQ1: What is the level of empirical maturity of federated learning implementations in healthcare, particularly in terms of deployment readiness, validation practices, and real-world applicability?

RQ2: How are unstructured clinical narratives and sentiment-related information represented and operationalized within federated learning frameworks for healthcare analytics?

RQ3: What technical, organizational, and governance constraints shape the integration of sentiment-aware learning within federated healthcare environments?

RQ4: What architectural and methodological patterns can be synthesized to inform the design of deployable, privacy-preserving federated learning systems for sentiment analysis on electronic health records?

2. METHODOLOGY

This study employed a Systematic Literature Review (SLR) design aligned with the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines [13]. The SLR approach was selected to ensure methodological transparency, reproducibility, and structured synthesis of evidence at the intersection of Federated Learning (FL), privacy-preserving techniques, sentiment analysis, and Electronic Health Records (EHRs). Rather than providing a descriptive overview of prior work, the review was designed to critically evaluate the empirical maturity, methodological transparency, and deployment-readiness of federated sentiment-aware healthcare analytics.

The review followed a structured workflow comprising: (i) search strategy formulation; (ii) database querying; (iii) duplicate removal; (iv) title and abstract screening; (v) full-text eligibility assessment; (vi) quality appraisal; and (vii) thematic synthesis and framework derivation.

2.1. Search Strategy

A structured database search was executed on January 5 2026 across five major bibliographic databases: IEEE Xplore, Scopus, Web of Science, ScienceDirect, and PubMed. These databases were selected to ensure comprehensive coverage of computer science, distributed systems, artificial intelligence, healthcare informatics, and clinical NLP research. The search covered peer-reviewed journal articles and conference proceedings published between January 2021 and January 5 2026, reflecting the period of accelerated development in federated learning architectures and transformer-based NLP systems following the COVID-19 pandemic. Only English-language publications with full-text availability were considered eligible. The following Boolean search string was applied across all databases, with minor syntax adaptations to accommodate database-specific query rules:

(("Federated Learning" OR "Collaborative Learning") AND ("Electronic Health Records" OR "Patient Medical Records") AND ("Privacy-Preserving" OR "Secure Aggregation") AND ("Sentiment Analysis" OR "Opinion Mining") AND ("Distributed Systems" OR "Distributed Computing"))

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Where supported, the query was restricted to title, abstract, and keyword fields to improve precision. Database-specific refinements included phrase searching, Boolean grouping normalization, and field-based filtering such as TITLE-ABS-KEY in Scopus. All retrieved records were exported in RIS or BibTeX format for centralized processing.

2.2. Study Selection Process

All exported records were imported into Mendeley Reference Manager for centralized bibliographic management, automated deduplication, and structured screening. The initial search retrieved 3,671 records. After automated and manual deduplication, 382 duplicate records were removed, leaving 3,289 unique records for evaluation. Title and abstract screening was conducted against predefined inclusion and exclusion criteria, with decisions cross-checked between authors to ensure consistent application. During this stage, 3,260 records were excluded for reasons including: absence of a healthcare application context; centralized ML architectures without federated or privacy-preserving components; lack of sentiment or narrative modelling relevance; and conceptual or editorial contributions without empirical implementation. The remaining 29 studies underwent full-text eligibility assessment, all of which met the inclusion criteria and were retained for qualitative synthesis. The full selection process is documented in the PRISMA flow diagram in Figure 1.

2.3. Inclusion and Exclusion Criteria

Studies were included if they: implemented federated or closely related distributed learning paradigms within a healthcare setting; applied privacy-preserving mechanisms; analyzed structured or unstructured EHR data; and incorporated sentiment analysis, narrative modelling, affective computing, or opinion mining as a primary or clearly defined analytical component. Only peer-reviewed journal articles and conference proceedings within the defined time window were considered. Studies were excluded if they: focused on non-healthcare datasets; relied solely on centralized ML without federated or privacy-preserving elements; lacked empirical validation; or did not meaningfully engage sentiment- or narrative-based analytics. Non-English publications were excluded to ensure consistent evaluation.

Broader conceptual and survey-based contributions identified during screening were examined contextually to support theoretical framing but were excluded from the empirical synthesis to prevent conflation with implementation maturity. Studies were classified as empirical if they reported implementation, experimentation, or evaluation of federated or privacy-preserving models using healthcare data; all others were treated as contextual contributions.

2.4. Quality Assessment

All 29 empirical studies were assessed using a qualitative appraisal rubric aligned with PRISMA 2020 evidence-certainty guidance. A criterion-based approach was adopted over numerical scoring to preserve methodological nuance better. Each study was examined across four dimensions: clarity of dataset description; transparency of federated learning configuration; completeness of privacy mechanism reporting; and rigour of evaluation design and performance metric disclosure. Each dimension was rated as clear, partially clear, or unclear. Studies with higher reporting clarity were treated as stronger evidence during thematic synthesis, though no formal exclusion weighting was applied. Results are summarized in Table 1.

Table 1. Quality Assessment of Included Empirical Studies

Ref	Dataset Description Clarity	Methodological Transparency	Privacy Reporting	Evaluation Rigor
[14], [15], [16], [17]	Clear	Clear	Partially Clear	Clear
[18], [19], [20]	Clear	Clear	Clear	Clear
[21], [22], [23], [24], [25]	Clear	Partially Clear	Clear	Partially Clear
[26], [27], [28], [29], [30]	Clear	Clear	Partially Clear	Clear
[31], [32], [33], [34]	Partially Clear	Clear	Partially Clear	Clear
[35], [36], [37]	Clear	Clear	Clear	Clear

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Ref	Dataset Description Clarity	Methodological Transparency	Privacy Reporting	Evaluation Rigor
[24], [38], [39], [40], [41], [42]	Partially Clear	Partially Clear	Partially Clear	Partially Clear

3. RESULTS AND DISCUSSION

Figure 1 presents the PRISMA 2020 flow diagram summarizing the study selection process. A total of 29 studies met the inclusion criteria and form the primary evidence base for this review.

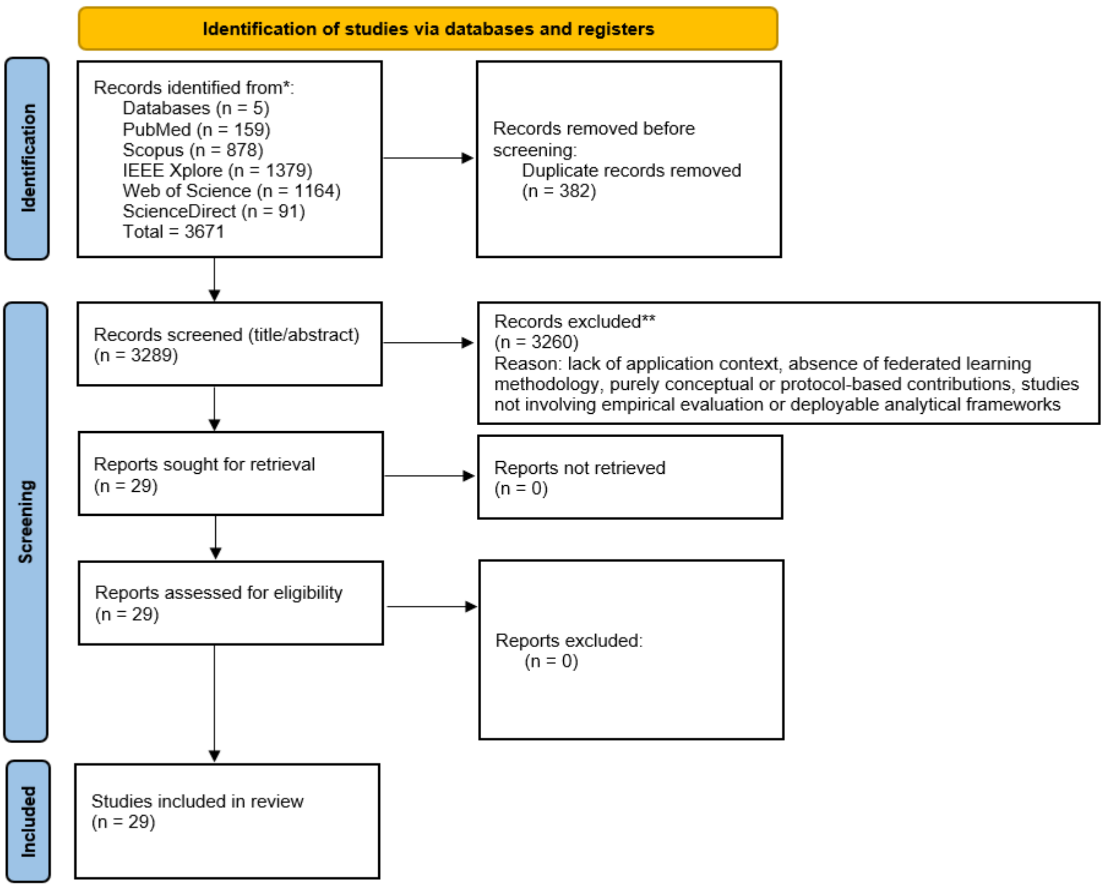

Figure 1. PRISMA 2020 Flow Diagram for Study Selection

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

The included studies were subsequently categorized by study type, publication region, and methodological focus, as summarized in Table 2. Table 2 summarizes the empirical studies included in this review. It presents the application context, the data type, the federated or privacy-preserving techniques applied, the learning algorithms, the evaluation metrics, and the key empirical contributions.

Table 2. Categorized Existing Research on Federated Learning and Privacy-Preserving AI in Healthcare (n = 29)

Ref	Country	Sentiment/Analytical Approach	EHR Application	Privacy-Preserving Mechanism	Opportunities	Challenges
[14]	Zimbabwe	Blockchain-based Record Tracking	Care Coordination	Smart Contracts & Consensus	Decentralized patient record trust	Semantic data interoperability
[15]	Netherlands	Multilingual TWHIN-BERT	Mental Health	Federated Averaging (FedAvg)	Global depression detection	Label distribution skewness
[18]	Netherlands	Gradient Boosting / Random Forest	Pandemic Forecasting	Predictive Learning (LHS)	Real-time data harmonization	Data heterogeneity
[19]	India	Neural Synchronization	Sustainable Health	Vertical Federated Learning (VFL)	Human-centric health systems	High computation overhead
[26]	USA / Intl	EXAM Model (FedAvg)	Global Clinical Analytics	Model Weight Aggregation	Collaborative shared learning	Model generalizability
[27]	Pakistan	SVM / Logistic Regression	Public Health Monitoring	Centralized Sentiment Analysis	Mining social media narratives	High noise in clinical text
[28]	Finland	Generative Adversarial Networks	Voice Diagnostics	Audio Privacy-Preservation	Voice-based medical diagnostics	High data transmission volume
[31]	South Korea	Scalable Privacy Protection	Large-scale FL Networks	Differential Privacy (DP)	Reliable wide-area FL	Privacy-Utility trade-off

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Ref	Country	Sentiment/Analytical Approach	EHR Application	Privacy-Preserving Mechanism	Opportunities	Challenges
[32]	Poland	Fuzzy Logic / LSGDM	Decision Support	Iterative Consensus	Handling uncertain data	Slow convergence speed
[39]	South Korea	Privacy-Preserving ML	Pandemic Digital Response	Anonymization & Encryption	Rapid digital health response	Linkage attack vulnerability
[24]	Italy / Swiss	Peer-to-Peer U-Net	Oncology Analysis	Decentralized Segmentation	Fault-tolerant tumor analysis	High training latency
[35]	Australia	MobileNetV2 / CNN	Rare Disease Diagnosis	Auto-Crafted Feature Extraction	Remote diagnostic accuracy	Sensor signal variability
[40]	Italy	P2P Federated Learning	Hospital Networks	Peer-to-Peer Data Sovereignty	Scalable hospital collaboration	High communication cost
[33]	Pakistan	Multi-Modal CNN	COVID Diagnosis	Edge-based Intelligence	Low-latency intelligence	Hardware constraints
[34]	Netherlands	Manifold Learning / Taskonomy	Task-Based Learning	Heterogeneous Task Learning	Optimized cross-domain learning	Task relatedness complexity
[41]	China	Secure Aggregation	Model Update Protection	Active Adversary Protection	Resilient update validation	Authentication overhead
[20]	Luxembourg	Statistical Quality Inference	Data Log Security	Quality-Aware Aggregation	Malicious participant detection	Log accuracy in secure settings
[23]	Japan	Discrete Fourier Transform	Low-resource IoT	Adaptive Gaussian Clipping	Efficient DP for IoT sensors	Bandwidth limitations
[36]	Australia	Federated Decision Trees	Interpretable Models	Explainable AI (XAI)	High model explainability	Privacy leakage from nodes
[16]	Canada	Transformers / RNN / LSTM	Opinion Mining	Deep Learning Sentiment Analysis	Decentralized opinion mining	Sensitivity to non-IID data
[29]	Bangladesh	Attention Mechanisms	Parkinson's Monitoring	Homomorphic Encryption (HE)	End-to-end data encryption	Computational complexity

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Ref	Country	Sentiment/Analytical Approach	EHR Application	Privacy-Preserving Mechanism	Opportunities	Challenges
[17]	Finland	Personalized FL (FPL)	User Modeling	Data Drift Management	Targeted patient treatments	Client-side drift management
[42]	Ireland/China	PoW / PoS Blockchain	Verifiable Diagnostics	Verifiable Immutable Logs	Diagnostic audit trails	Throughput bottlenecks
[25]	Tunisia	Multi-class CNN	Lung Disease Detection	Blockchain-Enabled Trust	Multi-institutional trust	Weight synchronization delay
[30]	India	zk-SNARKs / AES-256	Trustless Computation	Zero-Knowledge Proofs	Trustless model validation	Proof generation overhead
[37]	China	Threshold HE / BPS-FL	Byzantine Robustness	Threshold Encryption	Poisoning attack resilience	Heavy aggregation latency
[25]	India	SMPC / Edge Aggregation	ICU Monitoring	Secure Multi-Party Computation	Real-time ICU tracking	Secure sum computation latency
[30]	India	ResNet / AutoAlbum	Kidney Disease	Automated Data Augmentation	Robust diagnostic accuracy	Synthetic data validity
[37]	China	Contrastive Learning	Heart Rate Prediction	Clustered Federated Learning	Personalized predictive models	Time-series heterogeneity

3.1. Study Characteristics and Geographical Distribution

Figure 2 illustrates the geographic distribution of the 29 empirical studies. Asia accounts for the largest share (48.3%), driven by large-scale deployments in India, China, and South Korea, where the dual imperative of advancing clinical analytics and complying with stringent data governance frameworks has accelerated federated learning adoption. Europe follows at 31.0%, reflecting strong governance-aware and privacy-centric research traditions. North America, Africa, and Oceania collectively account for only 6.9% of implementations, indicating limited empirical translation despite established digital health infrastructures. African contributions, primarily from Zimbabwe and Tunisia, are highly localized, suggesting emerging but fragmented engagement. The near absence of

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

studies from South and Latin America further underscores a significant geographic imbalance, pointing to disparities in research capacity, regulatory maturity, and infrastructural readiness. Figure 2 highlights the geographic distribution of the 29 empirical studies included in this review.

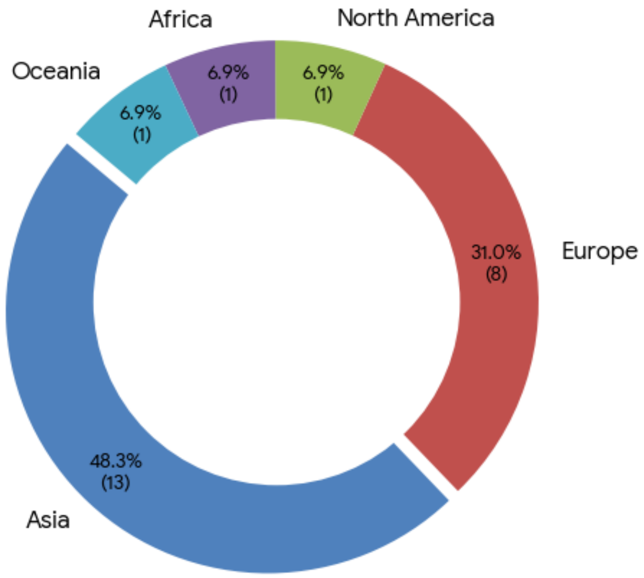

Figure 2. Geographic distribution of included studies (n = 29)

3.2. Sentiment and Analytical Approaches Identified

Table 3 reveals a clear structural imbalance in federated learning applications. Predictive and Decision Models dominate (n=7), reflecting the relative ease of working with structured clinical data where variables are standardized and evaluation metrics well established. Privacy-Preserving Analytics and Blockchain/Distributed Trust (n=6 each) indicate that research has prioritized secure infrastructure and institutional trust alongside predictive performance. Imaging and Signal Analysis (n=4) benefits from mature deep learning pipelines that transfer effectively into federated settings. NLP and Sentiment Analysis (n=3), by contrast, remains the least developed analytical domain alongside Personalised Learning. This underrepresentation is not coincidental: unstructured clinical text introduces vocabulary heterogeneity, site-specific documentation practices, and heightened privacy risks, collectively increasing implementation friction beyond what structured prediction tasks require. Federated learning has therefore progressed most in low-friction domains. At the same time, sentiment-aware analytics remain constrained at the high-friction edge of

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

implementation. Table 3 presents the six major analytical approach areas identified across the included literature.

Table 3. Sentiment and Analytical Approaches in Empirical Federated Healthcare AI (n = 29)

Analytical Approach Area	Count	Reference Studies	Technical Description	Primary Use Case
Predictive & Decision Models	7	[22], [26], [31], [35], [38], [40], [41]	FedAvg, Fuzzy Logic, and decision trees for forecasting.	Disease outcomes and clinical prediction.
Privacy-Preserving Analytics	6	[17], [21], [23], [24], [25], [34]	Scalable protection using DP, HE, and ZKPs.	Secure genetic and patient data exchange.
Blockchain & Distributed Trust	6	[14], [15], [16], [29], [32], [39]	Immutable ledgers and P2P record tracking.	Multi-institutional trust and auditing.
Imaging & Signal Analysis	4	[19], [27], [28], [30]	CNNs, GANs, and feature extraction for signals.	Radiological diagnostics and audio monitoring.
NLP & Sentiment Analysis	3	[18], [20], [22]	Neural text processing using BERT and LSTMs.	Mining clinical narratives and opinions.
Personalized Learning	3	[36], [37], [42]	Clustered and adaptive modelling for local drift.	Localized patient population management.

3.3. EHR Application Categories Identified

General Clinical Prediction and Secure Infrastructure lead (n=6 each), reinforcing their role as the core pillars of current federated healthcare research. Public Health and Pandemic Response and Specialised Diagnostics (n=5 each) reflect COVID-era collaborative momentum and the maturity of imaging-based pipelines. Narrative and User Analytics (n=4) remains comparatively underdeveloped, with several studies relying on partially federated designs, indicating limited real-world deployment readiness. Hospital

Systems and Coordination (n=3), despite being foundational for operational integration, is the least represented category. Collectively, this distribution reveals a structural bias toward low-complexity, easily evaluable domains, while sentiment-aware analytics remain constrained by linguistic variability, privacy sensitivity, and the absence of standardized implementation frameworks. Table 4 summarises the six primary EHR application areas identified from the empirical evidence.

Table 4. Frequency of EHR Applications in Empirical Studies (n = 29)

Application Category	Count	Reference Studies	Primary Focus
General Clinical Prediction & Decision Support	6	[23], [25], [30], [37], [39], [41]	Outcomes for kidney disease, heart rate, and ICU tracking.
Secure Infrastructure & Data Management	6	[17], [21], [22], [34], [40], [42]	Blockchain logging, IoT security, and Byzantine robustness.
Public Health & Pandemic Response	5	[19], [24], [26], [31], [38]	COVID diagnosis and global pandemic forecasting.
Specialized Medical Diagnostics	5	[16], [27], [29], [32], [38]	Oncology, rare diseases, and lung disease detection.
Narrative & User Analytics	4	[18], [20], [35], [36]	Mental health screening and clinical opinion mining.
Hospital Systems & Coordination	3	[14], [15], [28]	Multi-site record tracking and institutional synchronization.

3.4. Privacy-Preserving Mechanisms Observed

Standard model-weight aggregation dominates (15/29), reflecting its simplicity and reliance on data locality; however, it provides no provable privacy guarantee and remains susceptible to inference attacks. Blockchain-based approaches (n=6) emphasize auditability and trust rather than mathematical privacy protection. Differential Privacy, the only mechanism offering quantifiable guarantees, appears in just 4 studies, while Homomorphic Encryption and Secure Aggregation (n=2 each) remain limited by computational overhead. This distribution reflects a pragmatic prioritization of

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

deployability over privacy strength. For sentiment analysis of unstructured clinical narratives, where inference risks are inherently higher, this trade-off carries greater consequences, suggesting that current federated systems are optimized for feasibility rather than robust privacy protection. Table 5 summarises the five primary privacy-preserving mechanism (PPM) categories identified across the 29 implementations.

Table 5. Privacy-Preserving Mechanisms in Empirical Studies (n = 29)

Privacy Mechanism	Count	Reference Studies	Technical Role in Healthcare
Model Weight Aggregation (Standard FL)	15	[18], [19], [20], [22], [23], [25], [26], [27], [28], [31], [34], [35], [37], [41], [42]	Ensures only model parameters are shared rather than raw patient data, providing a baseline privacy layer.
Blockchain & Decentralized Security	6	[14], [15], [16], [29], [34], [39]	Establishes immutable audit trails and decentralized trust to verify model updates and prevent data tampering.
Differential Privacy (DP)	4	[17], [21], [24], [25]	Applies calibrated noise to local model updates, providing mathematically provable protection against patient re-identification.
Homomorphic Encryption (HE)	2	[23], [42]	Enables secure computation directly on encrypted parameters, protecting the model during transmission and aggregation.
Secure Aggregation / SMPC	2	[33], [40]	Uses multi-party protocols to prevent the central server from inspecting individual institutional updates.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

3.5. Opportunities Identified

Privacy-Optimized Data Analytics and Personalized Clinical Modelling lead (n=6 each), directly addressing the dominant challenges of privacy-utility trade-offs and non-IID data heterogeneity. Their prominence signals a field-wide shift toward treating privacy guarantees and site-level adaptation as foundational design requirements rather than optional enhancements. Blockchain and Federated Trust and Edge/Low-Latency Intelligence (n=5 each) form a second tier, highlighting governance and infrastructural readiness as critical enablers of institutional participation. Multimodal Implementation (n=4) and System Standardisation (n=3) are particularly consequential for sentiment-aware analytics: their limited presence underscores the ongoing difficulty of integrating unstructured clinical narratives and the absence of harmonized protocols, both of which constrain scalable federated deployments. Table 6 summarises the six major opportunity areas identified across the 29 empirical studies.

Table 6. Opportunities in Federated and Privacy-Preserving Healthcare AI (n = 29)

Opportunity Area	Study Count	Reference Studies	Primary Technical Achievement
Privacy-Optimized Data Analytics	6	[17], [21], [23], [24], [25], [34]	Mathematically provable anonymity using DP/HE.
Personalized Clinical Modeling	6	[30], [35], [36], [37], [39], [41]	Adapting global models to local site drift.
Blockchain & Federated Trust	5	[14], [16], [29], [40], [42]	Verifiable audit trails and decentralized logs.
Edge & Low-Latency Intelligence	5	[25], [31], [27], [28], [37]	Secure computations on clinical sensors.
Multimodal Implementation	4	[18], [20], [22], [33]	Unified processing of text and sensor data.
System Standardization	3	[15], [26], [38]	Scalable protocols for global analytics.

3.6. Challenges Identified

Non-IID Data and Heterogeneity emerge as the most significant challenge (n=6), reflecting inherent institutional differences in patient populations, clinical practices, and documentation styles. This issue is particularly acute for sentiment analysis on unstructured narratives, where linguistic variability further destabilizes model convergence. A second cluster, comprising Communication Overhead, Cost and Scalability Constraints, and Edge Resource Limitations (n=4 each), highlights the infrastructural burdens of synchronized, privacy-compliant model training across distributed environments. Privacy-Utility Trade-offs, Data Quality and Missingness, and Evaluation Gaps (n=3 each) reveal persistent methodological tensions, particularly the performance degradation introduced by stronger privacy mechanisms and the lack of standardized benchmarks. EHR Interoperability (n=2), though less frequently reported, remains a foundational constraint that is often underrepresented in controlled experimental settings. Collectively, these findings confirm that federated healthcare research is constrained more by deployment complexity than technical feasibility, with sentiment-aware analytics remaining at the periphery of mature implementation. Table 7 summarises the challenge areas identified across the empirical studies.

Table 7. Challenges in Federated and Privacy-Preserving Healthcare AI (n = 29)

Challenge Area	Study Count	Reference Studies	Primary Impact
Non-IID Data & Heterogeneity	6	[14], [15], [20], [26], [34], [37]	Hinders model convergence and creates predictive bias.
Communication Overhead & Latency	4	[24], [27], [39], [41]	Scalability limitations in low-bandwidth medical facilities.
Privacy-Utility Trade-offs	3	[18], [21], [35]	Stronger privacy tools often reduce clinical model accuracy.
Cost & Scalability Constraints	4	[16], [23], [40], [42]	High infrastructure costs limit institutional participation.
Interoperability of EHR Systems	2	[33], [39]	Fragmented standards prevent seamless federated training.
Data Quality & Missingness	3	[22], [25], [38]	Inconsistent records reduce model reliability.
Evaluation & Benchmarking Gaps	3	[19], [28], [30]	Difficulty in replicating and comparing results across studies.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Challenge Area	Study Count	Reference Studies	Primary Impact
Edge Resource Limitations	4	[17], [26], [29], [36]	Power constraints on devices limit edge deployment.

3.7. Discussion

The findings presented across Tables 4-7 provide a structured basis for addressing the four research questions guiding this review. Rather than indicating uniform progress, the evidence reveals differentiated patterns in how federated learning is applied, secured, and operationalized across healthcare contexts. The discussion proceeds by interpreting these patterns in relation to each research question, with specific attention to three interrelated dimensions: the dominance of structured analytical approaches over sentiment-aware NLP, the selection and effectiveness of privacy-preserving mechanisms [48], and the extent to which current implementations support federated sentiment-aware learning on unstructured EHR narratives.

1) **RQ1: What is the level of empirical maturity of federated learning implementations in healthcare, particularly in terms of deployment readiness, validation practices, and real-world applicability?**

Federated learning in healthcare has reached a level of operational maturity, though this maturity is uneven and domain-dependent. Structured clinical prediction tasks have advanced most rapidly, as they align with the core assumptions of federated optimization: consistent feature spaces, well-defined targets, and standardized evaluation protocols that enable stable multi-site validation [21], [28], [41]. Imaging and signal-based applications similarly benefit from mature deep learning pipelines that transfer effectively into distributed environments [26], [36], [38], [43].

Sentiment-aware approaches, by contrast, remain at an earlier stage of empirical maturity. This is not attributable to a lack of methodological capability: transformer-based models have demonstrated technical potential for distributed narrative analysis [15], [44]. Rather, the limitation lies in deployment design. Several studies reviewed here [20], [22] frame their work within sentiment analysis contexts yet rely on centralized or partially centralised architectures, indicating a persistent gap between analytical

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

sophistication and a fully federated system design. This pattern is absent from structured prediction studies such as [21], [34], where federated pipelines are consistently implemented end-to-end, further highlighting the domain-specific nature of implementation maturity. The field has demonstrated proof of capability but not yet consistent real-world readiness for privacy-sensitive, narrative-rich data modalities.

2) RQ2: How are unstructured clinical narratives and sentiment-related information represented and operationalized within federated learning frameworks for healthcare analytics?

Unstructured clinical narratives are not yet fully integrated into federated learning pipelines; they are treated as a secondary or emerging modality within broader healthcare analytics systems. Federated learning is more consistently operationalized around structured prediction tasks and secure infrastructure, where data representation is stable and governance requirements are more easily formalized [25], [30], [41].

When narrative data is explicitly incorporated, processing is typically mediated by transformer-based representations for tasks such as mental health inference and narrative classification [15]. However, as demonstrated by [15], [20], these implementations remain limited in scope and, in some cases, are only partially federated, indicating that free-text analytics within fully distributed environments has not yet been systematically realised. This contrasts sharply with structured modalities: studies such as [23], [37] implement complete federated pipelines with clearly defined privacy mechanisms, reflecting the greater operational tractability of structured data. The handling of narrative data is further constrained by governance considerations, with techniques such as secure aggregation and differential privacy shaping whether and how text data is processed in distributed settings [17], [34], [40]. As a result, unstructured clinical narratives remain conceptually recognized but operationally peripheral within current federated healthcare frameworks.

3) RQ3: What technical, organizational, and governance constraints shape the integration of sentiment-aware learning within federated healthcare environments?

The constraints limiting federated sentiment-aware learning are best understood as interdependent trade-offs spanning privacy, governance, and system feasibility rather than isolated technical barriers. Privacy-preserving mechanisms are selected based on contextual priorities, each introducing distinct operational burdens [45]. Differential Privacy provides formal inference protection but requires careful calibration to avoid degrading model utility in sensitive clinical contexts, as demonstrated by [27], [34]. Secure Aggregation enhances update confidentiality yet does not fully address trust and verification concerns in multi-institutional settings [33], [40].

More robust mechanisms, including Homomorphic Encryption and Zero-Knowledge Proofs, extend privacy assurances but introduce significant computational overhead that limits scalability, as evidenced by [17], [23], [42], [46]. These constraints reflect not merely technical limitations but also organizational realities: resource disparities, trust deficits, and regulatory compliance requirements that vary significantly across healthcare institutions and geographic contexts. For sentiment-aware learning specifically, these constraints are amplified: unstructured clinical narratives carry higher inference risk and greater linguistic variability, increasing the cost of applying already complex privacy mechanisms. The integration of sentiment analytics is therefore not only technically demanding but institutionally constrained, reinforcing the broader pattern that federated healthcare systems currently prioritize feasibility over analytically richer but operationally complex data modalities.

4) RQ4: What architectural and methodological patterns can be synthesized to inform the design of deployable, privacy-preserving federated learning systems for sentiment analysis on electronic health records?

Isolated technical choices do not drive the design of deployable federated sentiment-aware systems, but by recurring architectural patterns that emerge from the interplay between opportunity and constraint. Three interdependent design principles can be inferred from the evidence. First, privacy functions as a participation requirement rather than an optional layer, with mechanisms selectively applied according to institutional risk tolerance and trust conditions [27], [34], [37], [40]. Second, personalization is necessary to address non-IID heterogeneity: clustered and adaptive modelling approaches, as demonstrated by [35], [37], [41], prevent global models from disproportionately favouring data-rich or dominant sites. Third, governance and trust infrastructure are integral to system viability, as evidenced by blockchain and verification-based coordination mechanisms in studies such as [16], [42], [47].

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

Methodologically, deployable systems favour modular and hybrid designs in which sentiment-aware components are embedded within broader clinical analytics pipelines rather than implemented in isolation. Multimodal approaches combining narrative text with structured or sensor data represent this direction, though they remain limited by integration complexity and validation challenges [15], [20]. Edge-based and decentralized architectures further demonstrate the importance of locality and latency in healthcare environments [24], [26], [38], while also introducing resource constraints that directly shape privacy design choices.

These patterns are constrained by the same factors identified in Table 8. Non-IID heterogeneity necessitates personalization [21], [37]; privacy-utility trade-offs shape mechanism selection and calibration [27]; and infrastructural limitations restrict the scalability of computationally intensive solutions such as homomorphic encryption and blockchain coordination [16], [23], [39], [47]. Interoperability gaps and inconsistent evaluation frameworks further limit reproducibility, particularly for sentiment-aware applications where linguistic variability compounds architectural complexity [24], [28], [30]. The most viable architectural pattern emerging from the evidence is therefore not a single optimal configuration, but a balanced integration of privacy-preserving mechanisms, personalized modelling strategies, and governance-aware infrastructure capable of sustaining deployment under heterogeneous and resource-limited conditions [45].

5) Thematic Synthesis to Framework Derivation: Addressing Structural Gaps in Federated Sentiment-Aware Healthcare Systems

The preceding analysis reveals a consistent structural gap: while federated learning in healthcare has advanced in predictive modelling and privacy-preserving infrastructure, the integration of sentiment-aware analytics on unstructured clinical narratives remains fragmented, under-validated, and poorly operationalized [45]. Existing studies address components of this problem in isolation, focusing on model performance, privacy mechanisms, or governance structures, without converging on a cohesive, deployable design that simultaneously accounts for data heterogeneity, privacy risks, and institutional trust requirements.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

The proposed governance-aware conceptual framework was derived inductively to address this fragmentation. Thematic synthesis was conducted by identifying recurring patterns across the 29 studies in four areas: federated architecture design, privacy-preserving strategies, governance mechanisms, and heterogeneity management. These patterns were systematically compared to surface structural consistencies and persistent trade-offs shaping implementation feasibility. The extracted themes were subsequently abstracted into four higher-order socio-technical dimensions that directly informed the layers of the proposed governance-aware conceptual framework, as summarised in Table 8.

Table 8. Mapping of Extracted Themes to Conceptual Framework Layers

Extracted Theme	Key Studies	Framework Layer Informed
Non-IID data heterogeneity and site-level variability	[15], [20], [21], [35], [37]	Local Clinical Data Layer
Differential privacy, homomorphic encryption, secure aggregation	[17], [21], [22], [23], [24]	Privacy Enforcement Layer
FedAvg aggregation, personalized and clustered FL	[21], [36], [37], [41], [42]	Federated Coordination and Aggregation Layer
Blockchain audit trails, governance gaps, interoperability constraints	[14], [16], [29], [33], [39]	Governance and Accountability Layer

This mapping ensures transparency between the empirical evidence base and the framework architecture, directly addressing the structural gaps identified across the research questions.

Thematic synthesis identified four recurring structural dimensions directly mapped to the framework layers: local data handling, privacy enforcement, federated coordination, and governance accountability. Each layer reflects patterns consistently observed across the 29 empirical implementations and documented. Specifically, the dominance of standard FL aggregation and the underrepresentation of formal privacy mechanisms informed the Privacy Enforcement Layer; the non-IID heterogeneity challenge informed the Federated Coordination Layer; and the governance and interoperability gaps informed the Governance and Accountability Layer. This explicit linkage between

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

empirical findings and framework architecture ensures that the proposed design is grounded in evidence rather than imposed as a theoretical abstraction. The framework is therefore presented as a review-derived synthesis intended to guide future implementation-focused research, particularly for advancing privacy-preserving sentiment analysis in distributed and heterogeneous healthcare environments, rather than as a validated system architecture.

6) Proposed conceptual framework

This study, therefore, provides a conceptual framework for privacy-preserving, sentiment-aware federated analytics in distributed EHR environments. The framework organizes the system into four connected layers:

a) Local clinical data layer (data stays on-site): Clinics and hospitals store unstructured narratives and organized records, as well as raw EHR data. Sentiment extraction and other aspects of narrative modelling occur locally to protect private patient data from direct disclosure.

b) Privacy enforcement layer (protection before sharing): Model updates are subject to privacy constraints before leaving a site. Protection decisions (such as adding noise, encrypting updates, or limiting what is communicated) are decided at this layer. This layer's primary theoretical function is to define "safe sharing" in terms of clinical risk and compliance requirements.

c) Federated coordination and aggregation layer (learning without centralizing data): To create shared models, protected changes are combined while maintaining local autonomy. This layer highlights how personalization can be used to prevent systematic underperformance at minority or smaller sites, as well as how global learning is accomplished under heterogeneity.

d) Governance and accountability layer (trust, oversight, and interoperability): The governance system, which establishes who is involved, how permission is managed, how audits are carried out, and how interoperability is preserved, spans all levels. Given the evidence that technological viability by itself does not ensure acceptance, this layer is crucial.

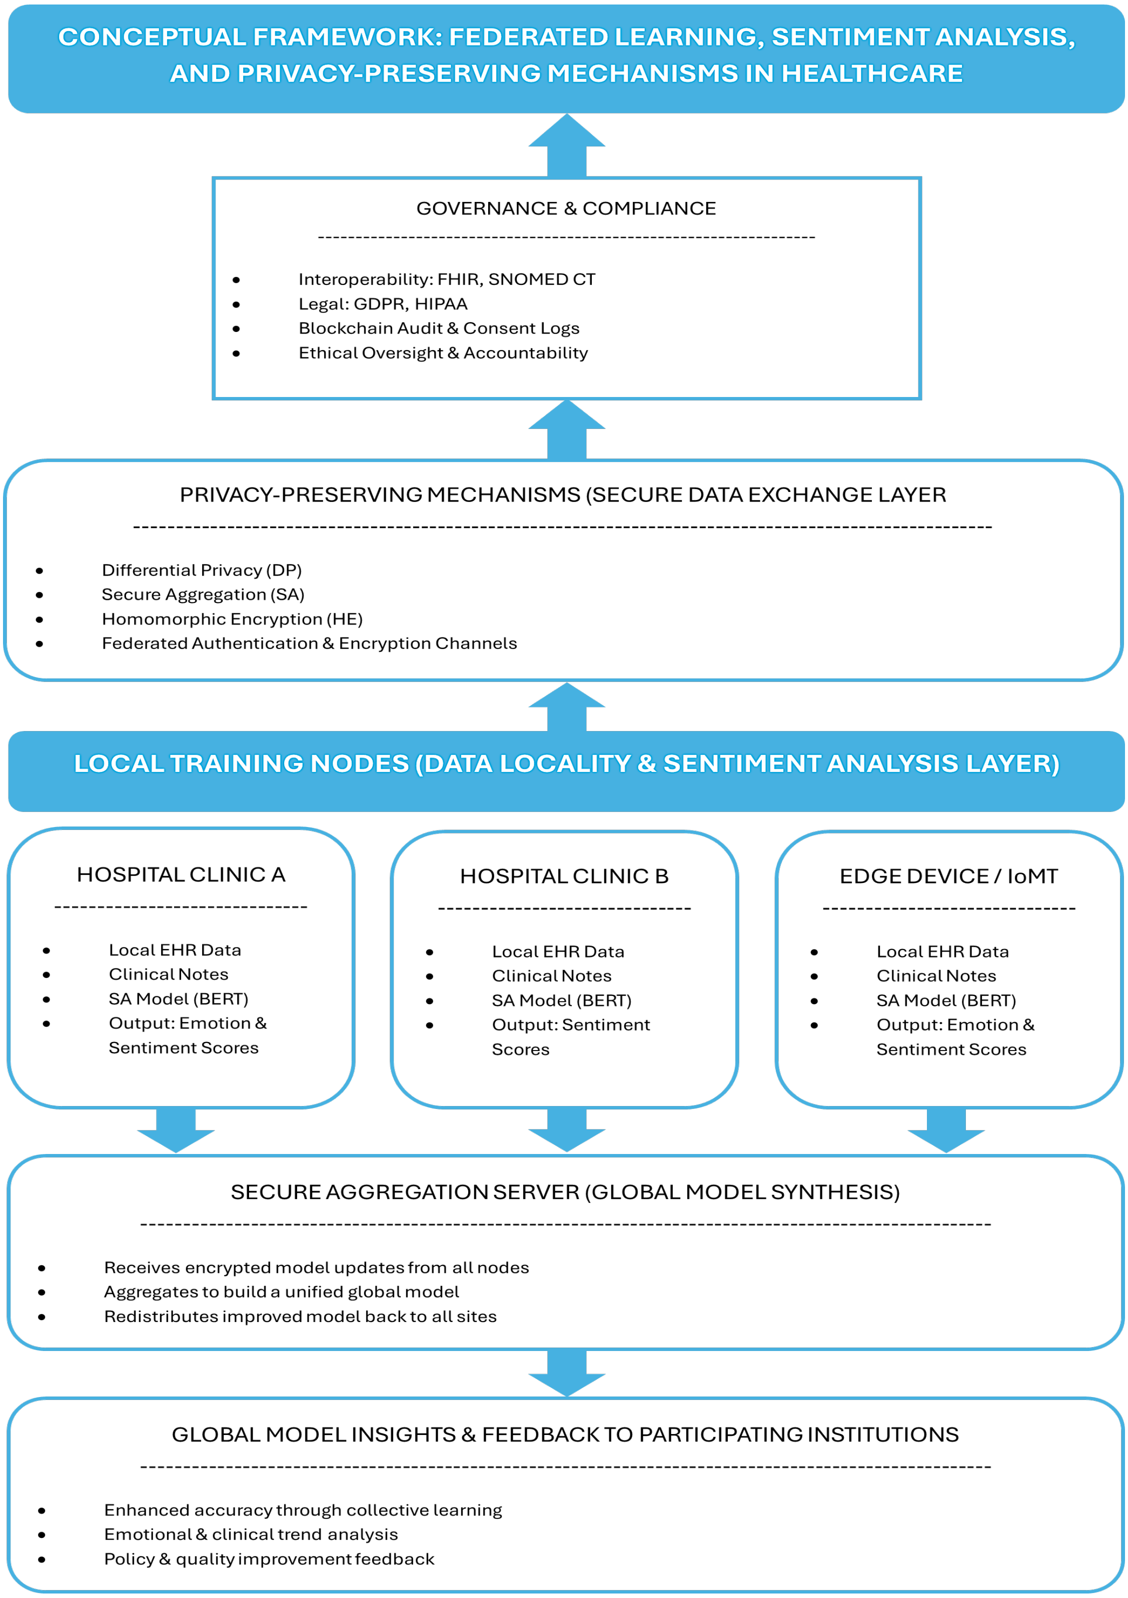

Figure 3. Conceptual Integration of Federated Learning, Sentiment Analysis, and Privacy-Preserving Mechanisms in Distributed Electronic Health Record (EHR) Environments

Unlike implementation-specific architecture, the framework presented here synthesizes recurring structural patterns identified across the reviewed studies. From a conceptual standpoint, the framework makes clear (i) the appropriate context for unstructured narrative analytics, (ii) the need to protect privacy, (iii) the information that is shared during federation, and (iv) the governance requirements for long-term cooperation. This immediately addresses the apparent disparity between the small number of fully functional sentiment-aware federated EHR deployments and promising technological approaches. By explicitly integrating data locality, privacy enforcement, federated coordination, and governance accountability, the framework provides a structured foundation for the next generation of deployable, privacy-aware clinical NLP systems in distributed healthcare environments.

3.8. Implications

This study translates the empirical evidence into practical, theoretical, and governance implications for deploying privacy-preserving federated learning in healthcare, with particular emphasis on sentiment-aware analytics in distributed electronic health record environments. Rather than presenting isolated recommendations, the implications reflect a coherent interpretation of how current capabilities, limitations, and design priorities interact in real-world settings.

1) Practical implications

The findings indicate that federated learning is deployable in healthcare, but only when the problem scope, data modality, and infrastructure readiness are carefully aligned. In practice, successful implementations are concentrated in bounded, structured use cases, suggesting that sentiment-aware analytics on unstructured clinical narratives should be introduced incrementally rather than as first-stage deployments.

Privacy preservation should be treated as a design decision shaped by deployment context rather than theoretical optimality. Lightweight coordination strategies are more feasible in resource-constrained environments. At the same time, computationally intensive privacy mechanisms are likely to remain limited to high-risk or well-resourced settings. This implies that implementers must balance privacy strength with operational feasibility, particularly where infrastructure and expertise are unevenly distributed. A staged deployment pathway emerges as a practical requirement: initial local validation, followed by controlled multi-site experimentation, and only then progression to larger federated networks. This progression is necessary not only to ensure technical robustness, but also to build institutional trust, validate governance processes, and manage operational risk over time

2) Theoretical implications

The findings reinforce the need to conceptualize federated healthcare analytics as a socio-technical learning system rather than a purely algorithmic framework. Variability across clinical sites is not an exception but a defining characteristic, requiring theoretical models that prioritize adaptability, robustness, and context-aware learning. More importantly, the study highlights a persistent fragmentation in how literature treats federated learning, clinical text analytics, and privacy preservation. These components are often developed in parallel, with limited integration into cohesive system-level designs. This fragmentation constrains both theoretical advancement and empirical maturity, particularly for sentiment-aware analytics on unstructured clinical data. Accordingly, there is a clear need for integrative theoretical frameworks that treat distributed learning, narrative understanding, and privacy enforcement as interdependent design dimensions. This need becomes more pronounced as emerging AI paradigms, including large-scale language models, intersect with federated healthcare systems, further increasing both analytical capabilities and privacy risks in distributed environments.

3) Policy and governance implications

The findings underscore that the scalability of federated healthcare systems depends as much on governance clarity as on technical capability. Institutional participation is unlikely in the absence of clearly defined accountability structures, consent models, audit mechanisms, and data stewardship responsibilities, particularly when handling sensitive, unstructured clinical narratives. This implies a transition from isolated pilot implementations to formalized governance frameworks that enable repeatable, scalable collaboration across institutions. Such frameworks must define roles, responsibilities, verification processes, and mechanisms for managing risk and accountability in distributed learning environments. In addition, interoperability should be addressed as a governance priority rather than solely a technical challenge. Differences in data standards, coding practices, and documentation conventions directly affect the feasibility

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

of coordination and scaling. Without alignment at this level, federated systems are likely to remain fragmented and difficult to operationalize at scale.

3.9. Limitations and Future Work

This review should be interpreted with several limitations in mind that affect its generalizability and depth of conclusions. First, the empirical synthesis is based on 29 implementation-focused studies drawn from a broader pool of 42, with the remaining studies providing conceptual and survey-based insights rather than deployable evidence. While these contributions are valuable for theoretical grounding, they also highlight a broader limitation in the field, namely, the relative scarcity of implementation-driven research, particularly for complex use cases such as sentiment analysis of unstructured clinical narratives. Second, the included studies exhibit inconsistent reporting, with strong emphasis on predictive performance but limited transparency into privacy strength, computational overhead, latency, and real-world deployment conditions. This uneven reporting constrains systematic comparison and assessment of their practical viability across diverse healthcare environments. Finally, the review is bound by database selection, language, and timeframe, focusing on studies indexed in major repositories between 2021 and early 2026. As such, relevant contributions outside these sources or in non-English contexts may not have been captured, introducing potential publication and selection bias.

The findings of this review point to several directions for advancing federated learning in healthcare, particularly for sentiment-aware analytics on unstructured clinical data. First, there is a clear need for more implementation-focused studies that move beyond structured data and imaging to fully operationalize federated learning on clinical narratives and patient-generated text. Such work should reflect real-world conditions, including heterogeneous documentation practices, site-specific variability, and stringent privacy constraints, rather than relying on controlled or simplified experimental settings. Second, future research would benefit from greater standardization in reporting and evaluation. Consistent disclosure of privacy assumptions, operational constraints (such as communication overhead and computational cost), and robustness to non-identically distributed data is essential for enabling meaningful comparisons and cumulative knowledge building across studies. The development of shared benchmarking frameworks and evaluation protocols would further strengthen the field's methodological coherence. Third, empirical validation in resource-constrained healthcare environments remains a critical priority. Infrastructure limitations continue to shape what is practically deployable, and without targeted investigation in such contexts, federated solutions risk remaining confined to well-resourced settings. Closely related to this is the need for reproducible experimental testbeds and benchmark datasets tailored to federated clinical NLP and sentiment-oriented tasks, which would support both validation and scalability. Finally, future work should explore integrating emerging AI paradigms, including large language models and foundation models, into federated clinical NLP pipelines. While these approaches offer significant potential to advance narrative understanding and decision support, they also introduce new challenges in privacy preservation, computational costs, and governance that must be addressed within the constraints of distributed healthcare environments.

4. CONCLUSION

This study provides a systematic, empirically grounded assessment of federated learning in healthcare, revealing a clear imbalance between the maturity of privacy-preserving infrastructure and the limited deployment of sentiment-aware analytics for unstructured clinical data. While federated approaches have advanced in structured domains, their application to narrative-rich, privacy-sensitive EHR contexts remains underdeveloped. Drawing on 29 implementation-focused studies, this study contributes a governance-aware conceptual framework integrating privacy, heterogeneity, and coordination as interdependent design considerations, serving as a blueprint to guide more coherent and deployable system designs. Progress in this area requires a shift from proof-of-concept studies toward implementation-focused research reflecting real-world clinical conditions. Greater alignment between methodological innovation and deployment feasibility will be essential for scalable, trustworthy federated sentiment-aware systems. These conclusions should be interpreted within the study's scope, and broader validation is required before generalizing findings to diverse healthcare settings.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

REFERENCES

[1] R. Kumar, A. Singh, A. S. A. Kassar, M. I. Humaida, S. Joshi, and M. Sharma, "Leveraging artificial intelligence to achieve sustainable public healthcare services in Saudi Arabia: A systematic literature review of critical success factors," *CMES-Comput Model Eng Sci*, vol. 142, pp. 1289–1349, 2025.

[2] Y. Liu, U. R. Acharya, and J. H. Tan, "Preserving privacy in healthcare: A systematic review of deep learning approaches for synthetic data generation," *Comput Methods Programs Biomed*, vol. 260, p. 108571, 2025.

[3] R. Chahar and A. K. Dubey, "Machine intelligence for mental health diagnosis: A systematic review of methods, algorithms, and key challenges," *Comput Mater Contin*, vol. 86, pp. 1–65, 2025.

[4] F. Alafari, M. Driss, and A. Cherif, "Advances in natural language processing for healthcare: A comprehensive review of techniques, applications, and future directions," *Comput Sci Rev*, vol. 56, p. 100725, 2025.

[5] K. A. S. Kumar, L. Nelson, and B. R. Jibinsingh, "Systematic review of privacy-preserving federated learning in decentralized healthcare systems," *Frankl Open*, vol. 13, p. 100440, 2025.

[6] M. I. Arias, L. Cadavid, and J. D. Velásquez, "Advancing healthcare analytics: A thematic review of machine learning, health informatics, and real-world data applications," *J Biomed Inf.*, vol. 171, p. 104934, 2025.

[7] H. N. C. Neto, J. Hribar, I. Dusparic, D. M. F. Mattos, and N. C. Fernandes, "A survey on securing federated learning: Analysis of applications, attacks, challenges, and trends," *IEEE Access*, vol. 11, pp. 41928–41953, 2023.

[8] H. Chamboko and B. Ndlovu, "Twitter (X) Sentiment Analysis on Monkeypox: A Systematic Literature Review," *Int J Inform. Dev*, vol. 14, no. 2, pp. 629–639, 2025, doi: 10.14421/ijid.2025.5196.

[9] C. Shiranthika, P. Saeedi, and I. V. Bajić, "Decentralized learning in healthcare: A review of emerging techniques," *IEEE Access*, vol. 11, pp. 54188–54209, 2023.

[10] M. U. Afzal, A. A. Abdellatif, M. Zubair, M. Q. Mehmood, and Y. Massoud, "Privacy and security in distributed learning: A review of challenges, solutions, and open research issues," *IEEE Access*, vol. 11, pp. 114562–114581, 2023.

[11] U. Zukaib, X. Cui, M. Hassan, S. Harris, H. J. Hadi, and C. Zheng, "Blockchain and machine learning in EHR security: A systematic review," *IEEE Access*, vol. 11, pp. 130230–130256, 2023.

[12] A. Tariq and others, "Trustworthy federated learning: A comprehensive review, architecture, key challenges, and future research prospects," *IEEE Open J Commun Soc*, vol. 5, pp. 4920–4998, 2024.

[13] M. J. Page and others, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, p. n71, 2021.

[14] R. Sibanda, B. Ndlovu, S. Dube, and K. Maguraushe, "Towards Health 4.0: Blockchain-based electronic health record for care coordination," in *Proc. 19th Eur. Conf. Innov. Entrep.*, 2024. doi: 10.34190/ecie.19.1.2606.

[15] S. S. Khalil, N. S. Tawfik, and M. Spruit, "Federated learning for privacy-preserving depression detection with multilingual language models in social media posts," *Patterns*, vol. 5, p. 100990, 2024.

[16] R. Myrzashova and others, "Safeguarding patient data-sharing: Blockchain-enabled federated learning," *IEEE Trans Sustain Comput*, vol. 10, pp. 176–189, 2025.

[17] S. B. Babu and K. R. Jothi, "A secure framework for privacy-preserving analytics in healthcare," *IEEE Access*, vol. 13, pp. 8439–8455, 2025.

[18] A. Ankolekar and others, "Using artificial intelligence and predictive modelling to enable learning healthcare systems for pandemic preparedness," *Comput Struct Biotechnol J*, vol. 24, pp. 412–419, 2024.

[19] A. Sarkar, "Secured human-centric sustainable healthcare in industry 5.0 through artificial neural synchronization-guided blockchain and vertical federated learning," *Appl Soft Comput*, vol. 175, p. 113040, 2025.

[20] D. Gholamiangonabadi and K. Grolinger, "Federated learning for sentiment analysis in presence of non-IID data," *IEEE Access*, vol. 12, pp. 128049–128060, 2024.

[21] I. Dayan and others, "Federated learning for predicting clinical outcomes in patients with COVID-19," *Nat Med*, vol. 27, pp. 1735–1743, 2021.

[22] F. Rustam and others, "A performance comparison of supervised machine learning models for COVID-19 tweets sentiment analysis," *PLoS ONE*, vol. 16, p. e0245909, 2021.

[23] S. A. Tanim and others, "Secure federated learning for Parkinson's disease," *IEEE Access*, vol. 12, pp. 127309–127327, 2024.

[24] A. Qayyum and others, "Collaborative federated learning for healthcare: Multi-modal COVID-19 diagnosis at the edge," *IEEE Open J Comput Soc*, vol. 3, pp. 172–184, 2022.

[25] A. Maurya and others, "Federated learning for privacy-preserving severity classification," *IEEE Access*, vol. 13, pp. 102339–102358, 2025.

[26] A. Woubie and T. Bäckström, "Federated learning for privacy-preserving speaker recognition," *IEEE Access*, vol. 9, pp. 149477–149485, 2021.

[27] C. Baek, S. Kim, D. Nam, and J. Park, "Enhancing differential privacy for federated learning at scale," *IEEE Access*, vol. 9, pp. 148090–148103, 2021.

[28] D. Połap, "Fuzzy consensus with federated learning method in medical systems," *IEEE Access*, vol. 9, pp. 150383–150392, 2021.

[29] R. Saidi and others, "A privacy-enhanced framework for chest disease classification," *IEEE Access*, vol. 13, pp. 185493–185512, 2025.

[30] A. Batra, P. Chatterjee, and J. Chaki, "Federated learning with AutoAlbum for kidney disease detection," *IEEE Access*, vol. 13, pp. 151033–151054, 2025.

[31] A. Majeed and S. O. Hwang, "A comprehensive analysis of privacy protection techniques developed for COVID-19 pandemic," *IEEE Access*, vol. 9, pp. 164159–164187, 2021.

[32] B. C. Tedeschini and others, "Decentralized federated learning for healthcare networks: A case study on tumor segmentation," *IEEE Access*, vol. 10, pp. 8693–8708, 2022.

[33] B. Pejó and G. Biczók, "Quality inference in federated learning with secure aggregation," *IEEE Trans Big Data*, vol. 9, pp. 1430–1437, 2023.

[34] M. A. Hidayat, Y. Nakamura, and Y. Arakawa, "Enhancing efficiency in privacy-preserving federated learning for healthcare," *IEEE Access*, vol. 12, pp. 88445–88457, 2024.

[35] H. Jamali-Rad, M. Abdizadeh, and A. Singh, "Federated learning with taskonomy for non-IID data," *IEEE Trans Neural Netw Learn Syst*, vol. 34, pp. 8719–8730, 2023.

[36] M. T. Hosain and others, "Privacy preserving machine learning with federated personalized learning," *IEEE Open J Comput Soc*, vol. 5, pp. 694–704, 2024.

[37] Y. Yin and others, "Prototypical clustered federated learning for heart rate prediction," *Front Inf Technol Electron Eng*, vol. 26, pp. 1896–1912, 2025.

ISI Journal of Information Systems and Informatics
Published By Asosiasi Doktor Sistem Informasi Indonesia

[38] T. Ngo and others, "Federated deep learning for the diagnosis of cerebellar ataxia: Privacy preservation and auto-crafted feature extractor," *IEEE Trans Neural Syst Rehabil Eng*, vol. 30, pp. 803–811, 2022.

[39] A. Giuseppi and others, "Decentralized federated learning for hospital networks with application to COVID-19 detection," *IEEE Access*, vol. 10, pp. 92681–92691, 2022.

[40] H. Tang, M. Li, and S. Ren, "ESE: Efficient security enhancement method for the secure aggregation protocol in federated learning," *Chin. J Electron*, vol. 32, pp. 542–555, 2023.

[41] S. R. Heiyanthuduwage and others, "Decision trees in federated learning: Current state and future opportunities," *IEEE Access*, vol. 12, pp. 127943–127965, 2024.

[42] J. Yu and others, "Blockchain-based privacy-preserving federated learning," *Big Data Min Anal*, vol. 8, pp. 189–213, 2025.

[43] A. Klemetti and others, "Systematic literature review on cost-efficient deep learning," *IEEE Access*, vol. 11, pp. 90158–90180, 2023.

[44] J. Machemedze and B. Ndlovu, "Transformer-based models for electronic health records," *J Appl Inf. Comput*, vol. 10, no. 1, pp. 90–105, 2026, doi: 10.30871/jaic.v10i1.11893.

[45] K. L. Chizengwe and B. Ndlovu, "A systematic review of agentic AI for threat detection and mitigation in 5G networks," *J Inf Syst Inform.*, vol. 8, no. 1, pp. 315–354, 2026, doi: 10.63158/journalisi.v8i1.1382.

[46] M. Moyo and B. Ndlovu, "Towards cloud-based electronic health records in healthcare systems: Security, scalability, and migration strategies," *J Inf Syst Inform.*, vol. 8, no. 1, pp. 739–780, 2026, doi: 10.63158/journalisi.v8i1.1431.

[47] N. Nezhadsistani, N. S. Moayedian, and B. Stiller, "Blockchain-enabled federated learning in healthcare," *IEEE Access*, vol. 13, pp. 119922–119945, 2025.