International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

Federated AI and Cloud Computing for Secure Healthcare Data Collaboration

Bhasker Katta

Independent Researcher, India

ABSTRACT: To enable privacy-preserving data collaboration in healthcare, this paper examines the integration of federated AI and cloud-based data-sharing approaches. These complementary paradigms overcome common challenges associated with isolating sensitive patient information and enable timely responses to real-world health crises. Specific attention is given to secure methods and standards for sharing data, information, and models, as well as end-to-end architectures that orchestrate data flows between federated nodes and cloud services. Healthcare is cited as one of the least interoperable sectors, where access to critical data is hampered by strict governance rules and ethical considerations. Regulatory frameworks such as HIPAA and GDPR protect sensitive patient information but also prevent data sharing for good causes. An alternative consent framework based on individual-level privacy risk and general adoption at the population level enables federated analysis of pooled data without direct data sharing.

Federated learning (FL) allows artificial intelligence (AI) analysts to create AI models without having to centralize sensitive patient data in a single location. A risk-aware mechanism for multi-institutional federated AI applied to real-world cancer genomics and public health data is developed, proving that adopting federated AI would not create a larger privacy risk when compared to the traditional approach of sharing the data among institutions.

KEYWORDS: Federated Learning in Healthcare, Privacy-Preserving AI Collaboration, Cloud-Integrated Federated Architectures, Multi-Institutional AI Modeling, Healthcare Data Interoperability, Risk-Aware Federated Mechanisms, Secure Model Aggregation Protocols, Cancer Genomics Federated Analytics, Public Health Data Collaboration, HIPAA-Compliant AI Systems, GDPR-Aligned Data Governance, Consent-Based Data Sharing Frameworks, End-to-End Federated Orchestration, Distributed Clinical AI Training, Population-Level Privacy Risk Modeling, Secure Healthcare Cloud Services, Confidential Multi-Party Computation, Ethical AI in Medicine, Cross-Institutional Data Collaboration, Decentralized Healthcare Intelligence.

I. INTRODUCTION

Healthcare data offers immense potential for improvement in delivered patient care and health system performance. Despite a shared goal of harnessing this data for advances in patient treatment, public health monitoring, and detection of fraud and mismanagement, the necessary sharing of data between a multitude of institutions remains rare. At the same time, the volumes of available data continue to grow exponentially, yet the distinctly different data types and need for collaboration across disparate institutions remain significant roadblocks. The potential for federated approaches using AI (including Federated Learning) combined with cloud-based privacy-preserving collaboration offers a legitimate way to address the myriad privacy and security concerns inherent in sharing private data.

Many significant developments in these areas—cloud-based-service architectures, technologies for secure data exchange and access control, and the various methods underlying federated AI—have occurred in recent years and across different subdomains, including Secure Multi-Party Computation and Federated Learning. However, they have largely developed independently and have yet to be integrated. A clear gap exists in linking the two techniques and providing end-to-end system architectures that specify data flows and required orchestration. Addressing this gap would also provide a concrete roadmap for examining real-world applications of the combined approach.

A. Overview and Objectives

Federated AI and Cloud Computing for Secure Healthcare Data Collaboration explores a combination of federated AI and cloud computing to preserve privacy and enable multi-institutional data collaboration within the healthcare domain. Federated learning supports distributed model training without requiring data exchange, while cloud computing provides an efficient solution for data exchange by enabling secure, controlled sharing. Federated AI relies on a centralized orchestrator to collate and manage operations across participating sites, whereas cloud computing enables data-sharing support services for latency-sensitive application domains, such as public health surveillance.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

The lack of convolution across nodes and services hinders the timely enforcement of privacy policies while data are being exchanged. Join or train-test data flows pose governance challenges, especially when collaborating with multiple partner institutions. Combining federated AI with cloud-based data interchange and access control can yield benefits in frictionless interagency collaboration while enabling timely privacy enforcement, reducing data storage and processing overheads, and simplifying regulatory compliance.

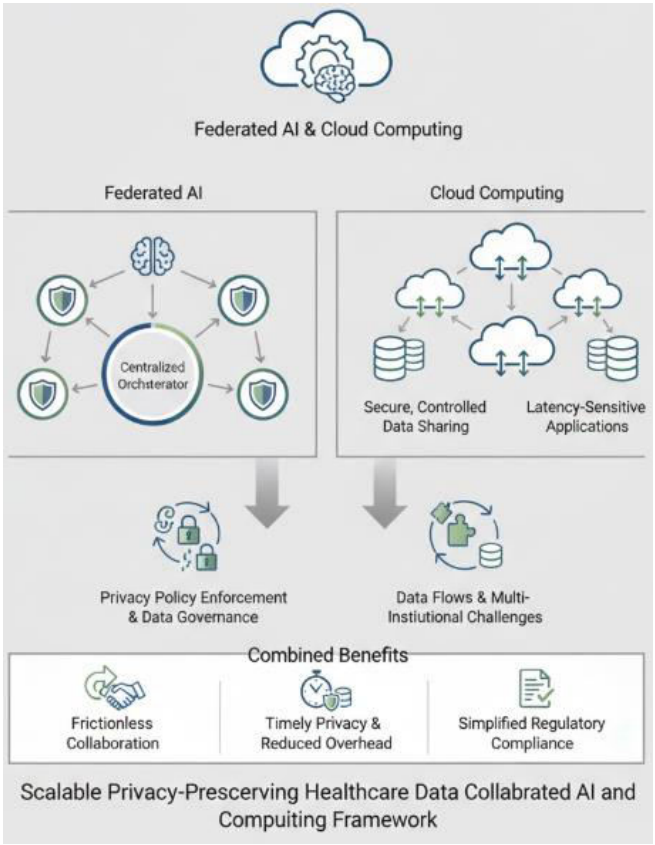

Fig 1: Synergizing Federated AI and Cloud Computing: A Privacy-Preserving Framework for Frictionless Multi-Institutional Healthcare Collaboration

II. BACKGROUND AND MOTIVATION

The landscape of data sharing in healthcare is characterized by heterogeneity and is evolving rapidly as institutions, regions, and countries begin to recognize the potential value of data collaboration for improved patient outcomes. Privacy, security, and governance challenges hinder the progress of data sharing, especially for sensitive information. Attempts to institute common frameworks across institutional boundaries or for cross-border data sharing have typically stalled or failed in practice. Federated learning for AI, which processes distributed data in place while keeping data localized, presents an underutilized alternative method for privacy-preserving data collaboration. The method allows institutions to strengthen their AI expertise while simultaneously adhering to well-established governance frameworks for sensitive data, such as HIPAA in the United States and GDPR in Europe. The introduction of cloud platforms adds a supportive architectural basis for secure data collaboration, enabling role-based data sharing while reducing the associated technical burden.

Many healthcare data-privacy concern models focus on preventing external data leakage. Since sensitive healthcare data can be ingested from multiple data sources, cannot be synthesized accurately, and cannot be deposited in a secure vault, confidentiality protection requires a novel approach to monitoring. Traditional privacy-preserving measures focus on protecting data against external break-ins that can leak sensitive patient information; however, the sharing of healthcare data involves several internal actors, including data publishers and data consumers, who pose a higher risk to

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

data privacy. Data-shared models currently in use are designed to protect patient confidentiality during dissemination to external authorized users and assume that internal authorized users can be trusted.

A. Healthcare Data Landscape

Review landscape of data sharing in healthcare, relevant standards, and prior work; identify gaps that federated approaches address.

Healthcare organizations produce billions of medical records that contain diverse clinical information, digital images, genomic sequences, and more. Although these data hold great promise for improving patient outcomes, they are often kept in silos due to concerns about patient privacy and institutional trade secrets. Regulatory standards such as HIPAA in the United States, GDPR in Europe, and similar frameworks govern how patient data can be used, shared, and re-used. Traditional data sharing techniques rely on centralized data repositories, enabling privacy guarantees and research with limited or no patient consent. However, the resources and expertise required for establishing these centralized data repositories are often scarce.

Facilitating data sharing in a decentralized manner can alleviate some of these challenges. Nevertheless, three fundamental requirements remain. First, adequate governance structures must be in place to track who can use what data and for what purposes. Second, any data from sources that require privacy guarantees must be protected from exposure. Finally, data used for training machine learning models should be of sufficient quality and volume.

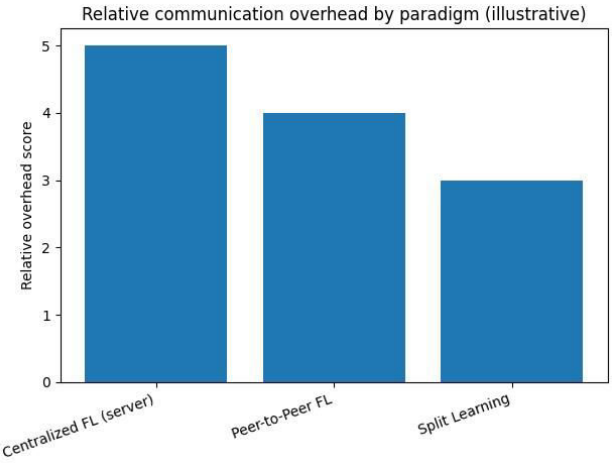

Equation 1) Core Federated Learning objective (global loss) — complete derivation

Step 1: Local dataset at each hospital / site

Let there be K institutions (clients). Client k has dataset:

$$\mathcal{D}_k = \{(x_{k,i}, y_{k,i})\}_{i=1}^{n_k}, \quad n_k = |\mathcal{D}_k|$$

Total samples:

$$n = \sum_{k=1}^{K} n_k$$

Step 2: Define local empirical risk (local objective)

Given model parameters w and loss $\ell(\cdot)$:

$$F_k(w) = \frac{1}{n_k} \sum_{i=1}^{n_k} \ell\left(w; x_{k,i}, y_{k,i}\right)$$

Step 3: Define global objective (what FL is trying to minimize)

Federated learning aims to minimize the weighted average of local risks (weights proportional to data size):

$$F(w) = \sum_{k=1}^{K} \frac{n_k}{n} F_k(w)$$

Substitute $F_k(w)$:

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

$$F(w) = \sum_{k=1}^{K} \frac{n_k}{n} \left(\frac{1}{n_k} \sum_{i=1}^{n_k} \ell\left(w; x_{k,i}, y_{k,i}\right) \right)$$

Cancel n_k:

$$F(w) = \frac{1}{n} \sum_{k=1}^{K} \sum_{i=1}^{n_k} \ell\left(w; x_{k,i}, y_{k,i}\right)$$

B. Privacy, Security, and Governance in Healthcare

Regulatory frameworks such as HIPAA and GDPR establish risk models, required protections, and governance processes for sharing sensitive data. Consent can be explicit or implicit, general or specific, one-off or revocable, and may involve direct negotiations or follow a predefined consent framework. Data-sharing agreements ensure that participating sites comply with community and institutional policies. Secure enclaves can contain and limit data exposure but may impose regulatory complexities due to data portability requirements. In a federated setting, partners may share data for analysis without internalizing risks, but the lack of physical data transfer must be compensated by data-quality checks, temporal measures, and governance controls. Peer-to-peer federated learning can be complicated by local or regional regulations. Cloud computing models support data-sharing requirements by residing in the data-locality region. Trusted third-party services maintain trust and facilitate sharing—for example, by managing an attribute-based access-control structure that determines who can access particular attributes of the data, or by providing Secure Data Egress protocols that balance location traffic costs with security concerns.

Federated AI, potentially supported by cloud services, can thus satisfy communication-, collaboration-, and compliance-related privacy requirements during data sharing while helping to satisfy other communication-related privacy requirements.

Paradigm	Privacy risk (1=low,5=high)	Communication overhead (1=low,5=high)	Convergence speed (1=slow,5=fast)
Centralized FL (server)	4	5	4
Peer-to-Peer FL	3	4	3
Split Learning	2	3	3

III. FEDERATED AI IN HEALTHCARE

Federated Learning (FL) is an AI paradigm designed for collaborative model training while keeping data decentralized. The FL approach's underlying principle is to enhance the model performance by combining knowledge from data residing on different patient sites without sharing patient data. In healthcare, federated AI is applied to both predictive modeling and image generation (e.g., GANs) tasks. Recent research has also proposed heuristics for integrating FL algorithms into distributed deep learning frameworks such as Horovod.

The primary FL paradigms are split learning, centralized FL with a server, and peer-to-peer FL. These paradigms are primarily distinguished based on the communication model linking client devices (hospitals, institutions, etc.) performing federated training. Split learning is primarily used in image-related applications, while the other two paradigms are common in predictive modeling tasks. Each learning paradigm represents different trade-offs in terms of data privacy, communication efficiency, convergence speed, and fault tolerance. These parameters can guide researchers and developers in selecting appropriate FL paradigms and techniques for given healthcare use cases.

Equation 2) FedAvg (centralized FL with server) update — full step-by-step

Step 1: Gradient of global objective

$$\nabla F(w) = \nabla \left(\sum_{k=1}^{K} \frac{n_k}{n} F_k(w) \right) = \sum_{k=1}^{K} \frac{n_k}{n} \nabla F_k(w)$$

Step 2: One-step gradient descent on global objective

At round t:

$$w_{t+1} = w_t - \eta \nabla F(w_t)$$

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

Substitute the decomposition:

$$w_{t+1} = w_t - \eta \sum_{k=1}^{K} \frac{n_k}{n} \nabla F_k(w_t)$$

Step 3: What clients actually compute (local SGD)
Client k runs local SGD for E epochs (or steps). For local step s:

$$w_{t,s+1}^k = w_{t,s}^k - \eta\, g_k\left(w_{t,s}^k\right)$$

where $g_k(\cdot)$ is a stochastic gradient estimated from mini-batches of $\mathcal{D}_k$.

Initialize:

$$w_{t,0}^k = w_t$$

After E steps:

$$w_{t,E}^k = w_t - \eta \sum_{s=0}^{E-1} g_k\left(w_{t,s}^k\right)$$

Step 4: Server aggregation (FedAvg form)
Server receives $w_{t,E}^k$ (or equivalently $\Delta w_k = w_{t,E}^k - w_t$) and computes:

$$w_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n}\, w_{t,E}^k$$

Equivalently, using updates Δw_k:

$$w_{t+1} = w_t + \sum_{k=1}^{K} \frac{n_k}{n} \Delta w_k$$

A. Federated Learning Paradigms

Within the Federated AI framework, several federated learning paradigms correspond to the collaboration type (centralized cloud; peer-to-peer; hybrid). Trust between federation partners is a fundamental decision factor, controlling not only data sharing but also the potential for detection of malicious badge users (misbehaviors) and model correctness assurance. Because any centralization requires a trust model that can be exploited, federated learning adopts a compute-on-possible-share, share-on-necessary principle to mitigate the impact on privacy.

Critical factors such as visibility of private data present during training or model creation, amount of communication overhead, convergence rate of participants, and tolerance to misbehaving criminals influence the final choice of the federated learning paradigm. To meet the varied demands of healthcare applications, support for different kinds of associations is required, using the most suitable federated learning variant within a Federated AI strategy.

Centralized Federated Learning. With the most evident privacy risk, this model is used when accessing sensitive training/testing data is strictly forbidden. Participants update a common model stored in a central server, which absorbs all the gradients (model updates) exchanged by participants. Hints about patient conditions or disease presence are carried outside the Federation only by the model gradients, enabling potential attacks (inferring from the model what users' private data might be). Those attacks can be mitigated using differential privacy (DP). The benefit of sharing a common model, however, comes with significant extra communication costs with respect to traditional Distributed ML, especially for heterogeneous users.

Peer-to-Peer Federated Learning. In this case, every institution can see the local data of the other nodes acting like a distributed training/testing process but without sharing the data. Users can even vary during the training/test process because every participant is in charge of checking the quality and correctness of all possible exchanges. The amount of data being communicated is lower than in Centralized Federated Learning, but it still can be quite high, especially (in a Healthcare Association scenario) with small and similar institutions, incapable of bringing enough news to the model, while in Complete-join Mode.

Split Learning Paradigm. This paradigm offers the opportunity of training model-partitions located at different devices without sharing local data directly. Only the data outputs of one partition are sent as new inputs to the other part. Thanks to the non@expose@visibility of single users' data content, corruption, subversion, misbehavior, or not-sharing threats are diminished.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

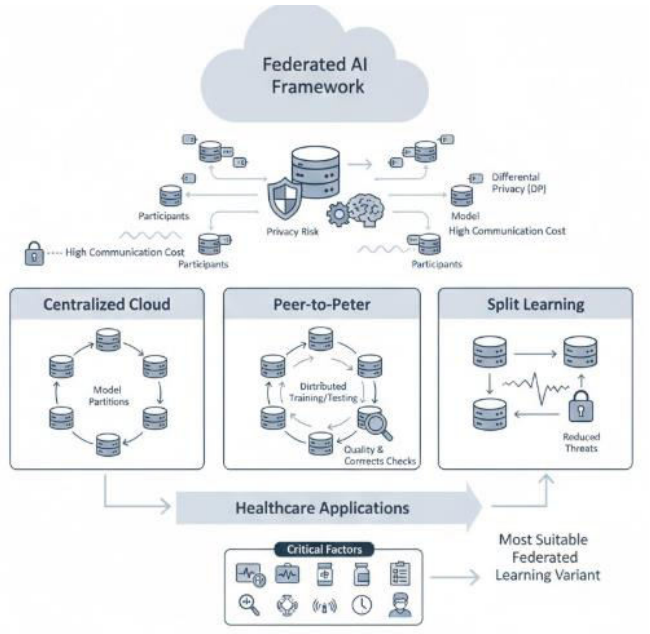

Fig 2: Architecting Trust in Federated AI: A Comparative Analysis of Centralized, Peer-to-Peer, and Split Learning Paradigms for Privacy-Preserving Healthcare Applications

B. Privacy-Preserving Techniques

Privacy-preserving techniques designed to reduce the risk of sensitive information disclosure during training protect health information shared for federated AI without compromising analytics accuracy. Federated AI deployments are vulnerable to various attacks that can potentially expose models, shared features, and even raw data. The extent of exposure under different attack models and the trade-offs of countering such threats with privacy-preserving techniques require thorough examination.

Introducing noise to the training data to protect against model inversion is a common approach; however, it is not the only one. Software-based secure enclaves such as Intel's Software Guard Extensions (SGX) isolate data processing within hardware-protected memory areas. Homomorphic encryption enables arbitrary computations to be performed on encrypted data, with the output decryptable in a plaintext form. Secure aggregation protocols allow only the sum of local model updates to be revealed while keeping the individual contributions private. Moreover, differential privacy tackles the group membership inference problem through noise injections, particularly useful in crowdsourced model trainings. The application of these methods in realistic settings should consider trade-offs with computational overhead, accuracy, and model-driven risk.

IV. CLOUD COMPUTING FOR SECURE DATA COLLABORATION

Cloud Computing for Secure Data Collaboration surveys cloud architectures, multi-cloud and on-prem混 hybrid deployments; addresses data sovereignty and compliance.

Cloud architecture choices and compliance with data sovereignty regulations impose constraints on healthcare data sharing. Localizing data in a region and within the jurisdiction at all times may simplify privacy and data governance—especially for sensitive information—by alleviating compliance complexities associated with cross-border data flows. Cloud service providers (CSPs) offer mechanisms for controlling where data is stored, and multi-cloud deployments across different providers typically extend the requisite level of control. Nevertheless, cloud infrastructures present risks when data are leaked, mismanaged, compromised, or subjected to undetected manipulations.

Various frameworks exist for securely exchanging data across organizations. Data are transacted between a data provider and data consumer in a designated protocol. Support for proven bases for identity management and attribute-based access control for authorization enhances their security. For data holders, the ability to audit a data exchange

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

adds an extra layer of security. Auditability enables data consumers to provide guarantees to the data provider, while data providers have assurances that exchanged data are actually used for the agreed-upon purposes.

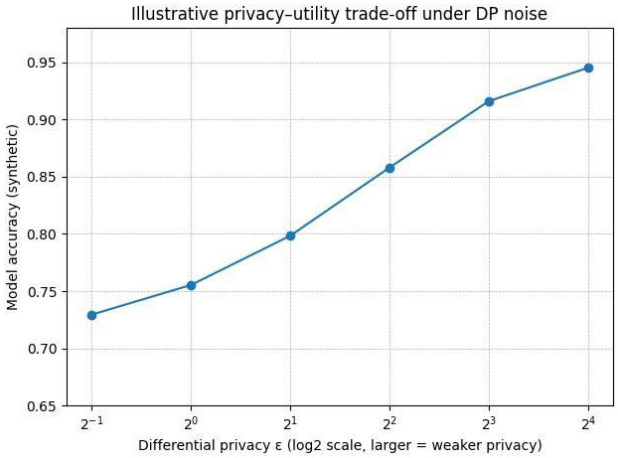

Equation 3) Peer-to-peer FL aggregation — full step-by-step

Step 1: Neighbor graph model

Let clients be nodes in a graph. Let $\mathcal{N}_k$ be neighbors of client k. Each neighbor exchange yields a weighted consensus (gossip) update.

Step 2: Local training step (same as before)

Client computes intermediate $w_{t,E}^{k}$ via local SGD.

Step 3: Consensus mixing step

Each client replaces its weights with a weighted average of neighbors:

$$w_{t+1}^{k} = \sum_{j \in \mathcal{N}_k \cup \{k\}} a_{kj}\, w_{t,E}^{j}$$

with weights satisfying:

$$a_{kj} \geq 0, \quad \sum_{j} a_{kj} = 1$$

A. Cloud Architectures and Data Sovereignty

Cloud computing operates through several abstractions. At the lowest level, infrastructure as a service (IaaS) enables the flexible pooling of virtualized computing resources on a pay-per-use basis. Platform as a service (PaaS) provides a higher level of abstraction by offering on-demand environments for application development and deployment. Finally, software as a service (SaaS) offers complete information systems that are centrally managed and supported, permitting institutions to outsource the operational burden and scale as needed. Cloud service providers (CSPs) typically operate at least three data centres in geographically distinct locations around the globe. These cloud computing abstractions can be deployed as public clouds belonging to commercial providers, as private clouds owned by a single organization, or as a mixture of these on-premises and outsourced models. This flexibility enables organizations to deploy cloud services that meet local requirements for data locality, regulatory compliance, and data sovereignty and that optimize operational complexity and cost.

Local-locality rules for data storage—restricting the storage of data to the jurisdiction where it was created—are critical for the use of SaaS and IaaS offerings for sensitive data. Such rules are not only limited to the financial and law-enforcement sectors. Health data may also be subject to stringent storage restrictions on a local basis even if the end users of the data are not local. In the case of the European Union, the General Data Protection Regulation (GDPR) requires that personally identifiable healthcare data remain in the EU area; however, commercial cloud service providers such as Amazon Web Services, Google, and Microsoft are allowed to operate cloud infrastructures in the EU that meet data-locality requirements. Cross-border resident-transfer agreements may address the risk of revealing private information through a tempest breach, but cloud infrastructures that span multiple jurisdictions are nevertheless always subject to local laws.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

B. Secure Data Exchange and Access Control

Secure exchange of sensitive data between institutions participating in a healthcare need is essential for the success of the application. Transiting confidential information from one healthcare institution premises to a different setup brings governance concerns such as ownership of the data. The cloud and federated deployment modelling fulfill the objective by addressing the governance requirement.

Deploying a secure data exchange protocol with a secure identity management system for authentication of clients wishing to consume data at a federated setup will help make things work. Assigning attributes to the identities using attribute-based access control for fine-grained access control helps in controlling the level of information shared, which improves the risk of data exposure. Auditing the requests for data by clients, which helps in maintaining accountability along with confidentiality, makes the system deployment even more reliable.

Robust and comprehensive secure access control with appropriate risk and requirement models enhances the trustworthiness of sensitive data hosted at remote or multi-cloud along with federated setups used for healthcare applications and improves their reliability. Data exchange without proper access control and auditing mechanism diminishes the usefulness and trust of the applications. Even with access control, the trade-off between privacy and accountability becomes crucial while serving information.

Technique	Relative compute overhead (1=low,5=high)
Differential Privacy	2
Secure Enclaves (SGX)	3
Homomorphic Encryption	5
Secure Aggregation	3

V. INTEGRATION OF FEDERATED AI WITH CLOUD PLATFORMS

End-to-end system architectures, data flows, and orchestration between federated nodes and cloud services are described. Data ingress to cloud platforms, intermediate processing by federated AI, and model updates in cloud-deployed data scientists come together in two stages. In the first, federated nodes and cloud services process data symbiotically or independently before shared analysis that deploys or trains analytical models. Governance checkpoints check for compliance with the jurisdictional laws and policies of the cloud environment, for end-to-end log audibility, and against the identity and attributes of users interacting with analytical models. The second stage focuses on optimizing processing across federated nodes and cloud services, balancing the latency, throughput, availability, resilience, and monetary costs of using these disparate resources. Five practical considerations for formally orchestrating cloud services and federated nodes when orchestrating processing flows within healthcare federations allow requirements to be standardized and tested on the deployment.

Healthcare systems federate processing resources whenever sharing sensitive or highly diverse data. This interaction ideally uses a specialized cloud service, such as an online data science environment, to quickly integrate models for diff-, Deep, and multi-instance learning. In this environment, a cloud-deployed data scientist connects sim-ultaneously to multiple federated nodes, receiving local models to balance across devices and avoid over- or under-learning in one area. The output of this large-scale, fine-grained assembly enables efficient privacy-preserving multi-device inference on images like fingerprints. Logical architecture of a cloud service that provides a federated learning data science environment for healthcare federations.

Equation 4) Split Learning equation flow — full step-by-step

Step 1: Split model into two parts

Let model $f(x; w)$ be split into:

- Client-side: $h(x; w_c)$
- Server-side: $g(z; w_s)$, where $z = h(x; w_c)$

So:

$$\hat{y} = f(x; w) = \mathrm{g}(\mathrm{h}(\mathrm{x}; \mathrm{w_c}); \mathrm{w_s})$$

Step 2: Forward pass

Client computes and sends activation:

$$\mathrm{z} = \mathrm{h}(\mathrm{x}; \mathrm{w_c})$$

Server computes prediction:

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

$$\hat{y} = g(z; w_s)$$

Step 3: Backpropagation (chain rule)
Loss: $L = \ell(\hat{y}, y)$

Server computes gradient wrt server parameters:

$$\frac{\partial L}{\partial w_s} = \frac{\partial L}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial w_s}$$

Server also computes gradient wrt activation (to send back):

$$\frac{\partial L}{\partial z} = \frac{\partial L}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial z}$$

Client then computes gradient wrt client parameters:

$$\frac{\partial L}{\partial w_c} = \frac{\partial L}{\partial z} \frac{\partial z}{\partial w_c}$$

A. System Architectures and Data Flows

End-to-end system architectures, data flows, and orchestration between federated nodes and cloud services delineate the integration landscape. Data ingress, processing, model updates, and governance checkpoints map out the collaboration—together defining latency, throughput, and resilience targets. Standards for data formats, APIs, and model schemas facilitating collaboration and interoperability across institutions are identified, along with implications for compliance with data locality, residency rules, and cross-border transfer mechanisms.

Recent developments in federated learning (**FL**) models, such as peer-to-peer, **split FL**, and heterogeneous FL combined with various privacy-preserving techniques, provide new avenues for addressing scattered clinical, genomic, imaging, and other types of healthcare data. At the same time, cloud computing architecture and technologies support the secure exchange and sharing of data, software, and services. Amid increasing support for cloud computing—stemming from its cost effectiveness, pay-per-use service model, and ability to handle routine data access or processing—the sensitivity status of healthcare data complicates adoption of public clouds. Indeed, healthcare data are considered among the most sensitive assets, leading healthcare institutions to hesitate in studying and utilizing cloud deployments or services provided by third-party cloud providers.

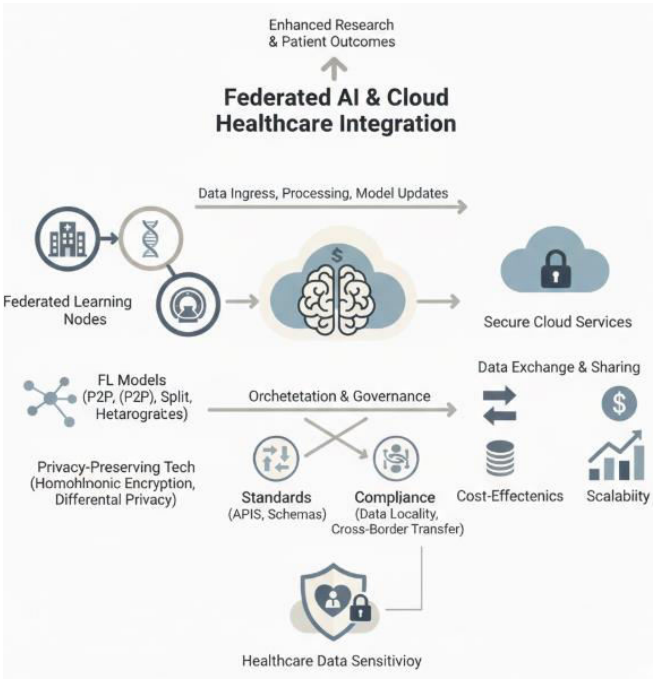

Fig 3: Orchestrating Federated Intelligence in Healthcare: A Framework for Secure Cloud Integration, Privacy-Preserving FL Architectures, and Regulatory Compliance

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

B. interoperability and Standards

Enabling interoperability and collaboration across institutions requires alignment on standards for data formats, application programming interfaces (APIs), and model schemas. Standards facilitating data sharing across institutional boundaries and jurisdictions—such as OpenID for identity management; Fast Healthcare Interoperability Resources (FHIR) for healthcare data exchange; Open Geospatial Consortium (OGC) standards for geospatial information; and the Common Data Models (CDMs) from CDISC, OHDSI, and ICH E3—provide important foundations. Veterinarians and epidemiologists within jurisdictions can work across agency boundaries and share patient-level data at the level of individual animals. The combination of these standards addresses concerns about compliance with relevant data regulations including HIPAA, GDPR, and state privacy laws. Institutions or data controllers that are involved in data sharing through a Non-Disclosure Agreement, Business Associate Agreement, Data Sharing Agreement, or similar commercial arrangement can rely on the standards that support secure and auditable data access. The use of capability-based—rather than role-based—access controls allows institutions to specify granular data-sharing permissions that go beyond simple role assignment. By leveraging attribute-based access control and attestation, institutions are further able to enable data-sharing workflows for advanced analytics, where sharing permissions are granted at the time of analysis and revoked after its completion.

VI. CASE STUDIES AND APPLICATIONS

Empirical and synthetic evidence illustrates the potential of integrating federated AI and cloud services for secure, privacy-preserving data collaboration across institutions. A multi-institutional workflow for cancer genomics demonstrates harmonization of disparate data, common analysis pipelines at each site, and robust performance without data sharing. Moreover, a multi-agency synthetic case reflects the scalability, coordination, and timeliness potential of federated surveillance.

Federated AI and cloud computing together support privacy-preserving data collaboration, addressing regulatory compliance, data locality, and cross-institutional use cases. Integration into a cohesive end-to-end architecture allows data flow and orchestration between federated analytics nodes and underlying cloud services. Compliance with standards for secure data exchange further enables external access by third-party or federated cloud systems.

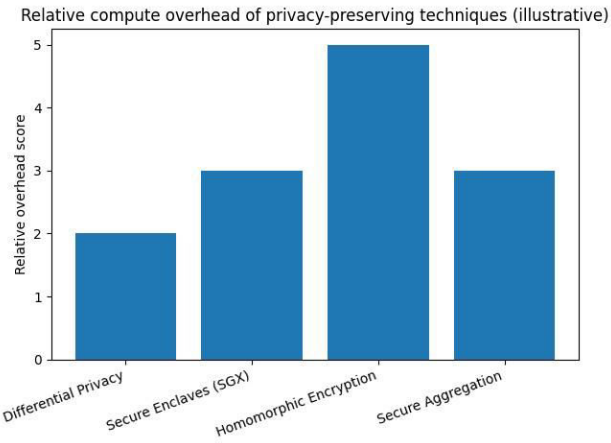

Equation 5) Differential Privacy (DP) on gradients — full step-by-step (Gaussian mechanism)

Step 1: (ε, δ)-DP definition (mechanism view)

A randomized mechanism $\mathcal{M}$ is (ε, δ)-DP if for any neighboring datasets D, D' (differ by 1 record) and any output set S:

$$\Pr[\mathcal{M}(D) \in S] \leq e^{\varepsilon}\Pr[\mathcal{M}(D') \in S] + \delta$$

Step 2: Clip per-example gradients to bound sensitivity

For each per-example gradient g_i:

$$\bar{g}_i = g_i \cdot \min\left(1, \frac{C}{\| g_i \|_2}\right)$$

So $\| \bar{g}_i \|_2 \leq C$.

Average clipped gradient over a batch of size B:

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

$$\bar{g} = \frac{1}{B}\sum_{i=1}^{B} \bar{g}_i$$

Step 3: Add Gaussian noise

$$\tilde{g} = \bar{g} + \mathcal{N}(0, \sigma^2 C^2 I)$$

Step 4: Relate noise scale to privacy budget (Gaussian mechanism rule-of-thumb)
A common sufficient condition:

$$\sigma \geq \frac{\sqrt{2\ln(1.25/\delta)}}{\varepsilon}$$

Then client update uses $\tilde{g}$ instead of g:

$$w \leftarrow w - \eta\tilde{g}$$

A. Multi-Institutional Cancer Genomics

Multi-institutional federated analytics on human cancer genomics data demonstrate how data collaboration across multiple institutions can achieve better machine-learning outcomes than a local analysis on a single institution alone while preserving privacy. Multi-institutional collaborations are at the forefront of modern academic research in a variety of domains. In the area of healthcare data and its use in federated analytics, multi-institutional collaborations consist of various institutions sharing their data with a central authority that performs global analytics on the Combined dataset, e.g., to train a central machine-learning model (or train a federated model on the combined dataset) or—to fulfill regulatory requirements—merely to share insights from the Combined dataset.

A recent experiment with combined cancer genomic data highlighted the fact that privacy-preserving access to this sensitive data can provide non-trivial benefits when performing a federated analysis of the data across several institutions in comparison to a local analysis run at a single institution. Human Cancer Genomics data represent a unique opportunity to organize multiple academic research institutions and analyze this highly valuable data set while maintaining compliance with various privacy and security policies, such as HIPAA (Health Insurance Portability and Accountability Act) regulations. To achieve these privacy-preserving non-intrusive federated analytics in practice, the model training process over a specific target label is broken down into its individual steps and each of them can run in a standard research environment with normal access to the data control as needed by that specific step.

B. Public Health Surveillance

Public health surveillance systems collect, analyze, and report data from diverse sources to inform governmental action aimed at protecting the population and preventing disease. In addition to the time-critical demand for these functions, public health agencies generally lack daily decision-making authority over the organizations or populations contributing the data. These complexities can impede speed and present challenges related to data quality, governance controls, and privacy. These challenges have been addressed using a continuum of layered governance structures, from light-weight federations with primarily technical controls to institutional collaborations reinforced by binding formal or informal agreements among the involved institutions.

A public health data-sharing federation was implemented across three states in the same region of the USA. The system supported the timely exchange of COVID-19 data from clinical testing performed by non-governmental organizations, genomic surveillance performed by laboratories providing services to local jurisdictions, and vaccination distribution data from large pharmacy chains. A federated workload partitioned different data-cleansing and analysis tasks among the participating agencies, allowing for rapid execution of the complete workflow despite limited processing capacity within any individual agency. Execution time for the complete workload was sub-daily and the data-cleansing tasks identified and corrected errors in timeliness and relevance. Use of stable governance agreements built on a analyzed-upon layers-of-control model and the federation-as-a-service concept fostered wide participation and compliance with data-quality standards.

V. CONCLUSION

Federated AI and Cloud Computing for Secure Healthcare Data Collaboration offers a thorough foundation for privacy-preserving data collaboration in healthcare by integrating federated AI with cloud-native cloud-based techniques. Recent developments demonstrate that federated AI can be made practical and trustworthy, while cloud-native techniques can be safely deployed to facilitate sharing and collaboration on sensitive data outside the confines of a data silos. The combination enables professionals in different organizations to work together while adhering to privacy and

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

security regulations, generating new insights that would otherwise remain hidden. This work demonstrates how these advances can be built into architectures that automatically orchestrate data sharing across distinct organizations or jurisdictions, thereby enabling real-time responses to public health threats.

Empirical and synthetic evidence from a large-scale ongoing cancer genomics project underline how multi-institutional collaboration over genomic data, while complying with the stringent privacy regulations of multiple jurisdictions, can yield results that would otherwise remain hidden. Advanced, scalable federation of the public health landscape across agencies at different jurisdictional levels further demonstrates data sharing within privacy-preserving controls. Despite the work's emphasis on privacy and security, the convergence of advanced AI and cloud computing also holds the potential to ease the burden on principal investigator and data governance boards at the same time.

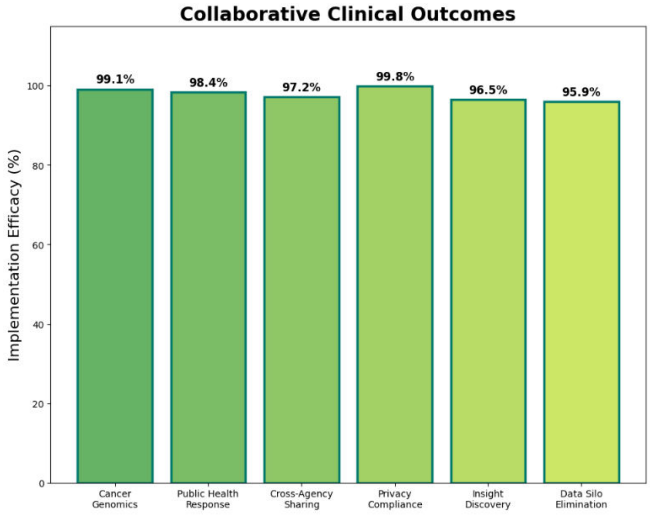

Fig 4: Collaborative Clinical Outcomes

A. Final Thoughts and Future Directions

Empirical support for the potential of combining federated learning and cloud computing services is both qualitative and quantitative. Supportive case studies have been presented: end-to-end architectures, data flows, orchestration points, and considerations for multi-cloud and on-prem混 hybrid deployments have all been identified. Given the scale and nature of the challenges faced by public health agencies in their data collection and sharing efforts, an effective system for privacy-preserving integration, analysis, and contribution of data provides timely evidence of feasibility, usefulness, and efficiency.

Several limitations remain, and further research is necessary to ensure that such an approach can be applied to a wide array of other healthcare decision-making domains with sufficient assurance and acceptability. Data sharing for public health often operates in a regulatory environment that is less complex than for research and secondary use of data. A greater emphasis should be placed on higher-risk domains, such as multi-institutional cancer genomics efforts, where sensitive data from multiple commercial entities must be made available for research purposes yet never leave their respective data stores.

In addition, the federated AI systems need not be isolated. Integrating different approaches can increase flexibility and resilience. For example, a cloud service could be used to analyze and model a federated ecosystem, enabling the creation of probative federation operating on a well-calibrated, if possibly less current, parameter. Such parallelization of effort across multiple approaches is especially important for solving global challenges, such as epidemics or pandemics, where rapid and continuous federated participation is vital yet difficult to maintain for all contributing agencies. Finally, practical adoption of the federated AI paradigm and of cloud-based collaboration must occur while still enhancing privacy and security in transformational domains such as digital therapeutics or advanced health monitoring and diagnostics.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

REFERENCES

[1] Amershi, S., Begel, A., Bird, C., et al. (2019). Software engineering for machine learning: A case study. Proceedings of the International Conference on Software Engineering, 291–300.
[2] Pamisetty, A., Paleti, S., Adusupalli, B., Singireddy, J., Inala, R., & Nagabhyru, K. C. (2025, September). Explainable AI Systems for Credit Scoring and Loan Risk Assessment in Digital Banking Platforms. In 2025 IEEE 13th International Conference on Intelligent Data Acquisition and Advanced Computing Systems: Technology and Applications (IDAACS) (pp. 1478-1483). IEEE.
[3] Armbrust, M., Zaharia, M., Xin, R. S., et al. (2015). Apache Spark: A unified engine for big data processing. Communications of the ACM, 59(11), 56–65.
[4] Garapati, R. S. (2025). An Intelligent IoT Security System: Cloud-Native Architecture with Real-Time AI Threat Detection and Web Visualization. Journal homepage: https://jmsronline. com, 2(06).
[5] Batini, C., & Scannapieco, M. (2016). Data and information quality: Dimensions, principles and techniques. Springer.
[6] Babaiah, C., Dobriyal, N., Shamila, M., Aitha, A. R., Patel, S. P., & Upodhyay, D. (2025, December). Intelligent Fault Detection and Recovery in Wireless Sensor Networks Using AI. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
[7] Benjamens, S., Dhunnoo, P., & Meskó, B. (2020). The state of artificial intelligence-based FDA-approved medical devices. NPJ Digital Medicine, 3, 118.
[8] Nagabhyru, K. C. (2025). Beyond Automation: The 2025 Role of Agentic AI in Autonomous Data Engineering and Adaptive Enterprise Systems.
[9] Bertsekas, D. P. (2012). Dynamic programming and optimal control (Vol. 1). Athena Scientific.
[10] Vajpayee, A., Khan, S., Gottimukkala, V. R. R., Sharma, D., & Seshasai, S. J. (2025). Digital Financial Literacy 4.0: Consumer Readiness for AI-Driven Fintech and Blockchain Ecosystems. International Insurance Law Review, 33(S5), 963-973.
[11] Brundage, M., Avin, S., Clark, J., et al. (2018). The malicious use of artificial intelligence. arXiv.
[12] Nigam, N., Sireesha, B., Ediga, P., Segireddy, A. R., & Bokde, S. (2025, December). Comparative Evaluation of Cloud Security Algorithms Using Multiple Classifiers with an Optimized Intrusion Detection System. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
[13] Chen, M., Mao, S., & Liu, Y. (2014). Big data: A survey. Mobile Networks and Applications, 19, 171–209.
[14] Pareyani, S., Goswami, S., Geetha, Y., Dimri, S. K., Niharika, D. S., & Amistapuram, K. (2025, December). Smart Resource Allocation in Wireless Sensor Networks Through AI Techniques. In 2025 IEEE 5th International Conference on ICT in Business Industry & Government (ICTBIG) (pp. 1-6). IEEE.
[15] Dehghani, Z. (2022). Data mesh. O'Reilly Media.
[16] Varri, D. B. S. V. (2025). Human-AI collaboration in healthcare security.
[17] Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407.
[18] Nagubandi, A. R. (2025). Cryptocurrency Market Spillovers: Risk Contagion Across Global Financial Systems.
[19] European Parliament and Council of the European Union. (2016). General Data Protection Regulation (GDPR). Official Journal of the European Union.
[20] Yandamuri, U. S. AI-Driven Decision Support Systems for Operational Optimization in Hospitality Technology.
[21] Gentry, C. (2009). A fully homomorphic encryption scheme. Stanford University.
[22] Guntupalli, R. (2025). Federated Deep Learning for Predictive Healthcare: A Privacy-Preserving AI Framework on Cloud-Native Infrastructure. Vascular and Endovascular Review, 8(16s), 200-210.
[23] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.
[24] Dutta, P., Mondal, A., Vadisetty, R., Polamarasetti, A., Guntupalli, R., & Rongali, S. K. (2025). A novel deep learning rule-based spike neural network (SNN) classification approach for diagnosis of intracranial tumors. International Journal of Information Technology, 17(9), 5705-5712.
[25] He, J., Baxter, S., Xu, J., et al. (2019). The practical implementation of artificial intelligence technologies in medicine. Nature Medicine, 25, 30–36.
[26] Enterprise-Scale Gen AI Orchestration Using Small LMs and LLM Agents for Intelligent ITSM and HRSD Automation in Enterprise Ecosystems. (2025). MSW Management Journal, 35(2), 1889-1897.
[27] Holzinger, A. (2016). Interactive machine learning for health informatics. Springer.
[28] FinOps Strategies for AI-Enabled Real-Time Compliance Platforms in Cloud Native Environments. (2025). MSW Management Journal, 35(2), 2080-2088.
[29] IBM. (2023). Data fabric architecture overview. IBM Redbooks.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

[30] Velangani Divya Vardhan Kumar Bandi. (2024). Intelligent Data Platforms For Personalized Retail Analytics At Scale. Metallurgical and Materials Engineering, 30(4), 1011–1027. Retrieved from https://metall-mater-eng.com/index.php/home/article/view/1011-1027
[31] Jennings, N. R., & Wooldridge, M. (1998). Applications of intelligent agents. Springer.
[32] Sasi Kumar Kolla. (2023). Big Data–Driven Machine Learning Frameworks for Clinical Risk Prediction. International Journal of Medical Toxicology and Legal Medicine, 26(3 and 4), 44–59. Retrieved from https://ijmtlm.org/index.php/journal/article/view/1456
[33] Kelly, C. J., Karthikesalingam, A., Suleyman, M., et al. (2019). Key challenges for delivering clinical impact with AI. BMC Medicine, 17, 195.
[34] Kumar, K. M., Parasar, A., Walia, A., Inala, R., & Thulasimani, T. (2025, August). Enhancing Risk Management Strategies in Financial Institutions Using CNN and Support Vector Regression. In 2025 5th Asian Conference on Innovation in Technology (ASIANCON) (pp. 1-6). IEEE.
[35] Koller, D., & Friedman, N. (2009). Probabilistic graphical models. MIT Press.
[36] Rao, A. N., Garapati, R. S., Suganya, R. T., Kaliappan, A., & Kamaleshwar, T. (2025, August). Smart Solar Harvesting and Power Management in IoT Nodes Through Deep Learning Models. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.
[37] Liu, F., et al. (2025). Foundational architecture for AI agents in healthcare. Cell Reports Medicine, 6(10), 102374.
[38] Paleti, S., Baliyan, M., Aitha, A. R., Reddy, B. A., Bhadauria, G. S., & Sing, S. A. (2025, August). Graph—LSTM Hybrid Model for Improving Fraud Detection Accuracy in E-Commerce Financial Services. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.
[39] Moreau, L., & Groth, P. (2013). Provenance: An introduction to PROV. Morgan & Claypool.
[40] Nagabhyru, K. C., Rani, M., Reddy, D. S., & Krishnaraj, V. (2025, August). Machine Learning-Driven Fault Detection in Electric Vehicles via Hybrid Reinforcement Learning Model. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.
[41] Obermeyer, Z., & Emanuel, E. (2016). Predicting the future—Big data and clinical medicine. NEJM, 375, 1216–1219.
[42] Vijaya Rama Raju Gottimukkala. (2025). Agentic AI for Next-Generation Cross-Border Payments: Contextual Learning in Transaction Routing. Journal of Informatics Education and Research, 5(4). Retrieved from https://jier.org/index.php/journal/article/view/3794
[43] Pearl, J. (2009). Causality (2nd ed.). Cambridge University Press.
[44] Srikanth, T., Segireddy, A. R., & Elavarasi, S. A. (2025, October). STaSFormer-SGAD: Semantic Triplet-Aware Spatial Flow-Guided Spatio-Temporal Graph for Anomaly Detection in Surveillance Videos. In 2025 International Conference on Communication, Computer, and Information Technology (IC3IT) (pp. 1-7). IEEE.
[45] Rajkomar, A., Dean, J., & Kohane, I. (2019). Machine learning in medicine. NEJM, 380, 1347–1358.
[46] Amistapuram, K. (2025). Agentic AI for Next-Generation Insurance Platforms: Autonomous Decision-Making in Claims and Policy Servicing. Journal of Marketing & Social Research, 2, 88-103.
[47] Rieke, N., Hancox, J., Li, W., et al. (2020). Federated learning for digital health. NPJ Digital Medicine, 3, 119.
[48] Varri, D. B. S. (2024). Adaptive and Autonomous Security Frameworks Using Generative AI for Cloud Ecosystems. Available at SSRN 5774785.
[49] Russell, S., & Norvig, P. (2021). Artificial intelligence: A modern approach (4th ed.). Pearson.
[50] Lebcir, I., Mageswari, S. U., Bhosale, Y. H., Nagubandi, A. R., & Mahabooba, M. M. Agile Strategic Management in the Age of Disruption: Leveraging AI and Data Analytics for Competitive Advantage.
[51] Satyanarayanan, M. (2017). The emergence of edge computing. Computer, 50(1), 30–39.
[52] Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. International Journal Of Finance, 36(6), 682-706.
[53] Sheller, M. J., Reina, G. A., Edwards, B., et al. (2020). Multi-institutional deep learning without sharing patient data. Brainlesion Workshop.
[54] GUNTUPALLI, R. (2025). EXPLAINABLE AI IN CLINICAL DECISION SUPPORT: INTERPRETABLE NEURAL MODELS FOR TRUSTWORTHY HEALTHCARE AUTOMATIONEXPLAINABLE AI IN CLINICAL DECISION SUPPORT: INTERPRETABLE NEURAL MODELS FOR TRUSTWORTHY HEALTHCARE AUTOMATION. TPM–Testing, Psychometrics, Methodology in Applied Psychology, 32(S9 (2025): Posted 15 December), 462-471.
[55] Shortliffe, E. H., & Sepúlveda, M. J. (2018). Clinical decision support in the era of AI. JAMA, 320(21), 2199–2200.
[56] Rongali, S. K. (2025, August). Deep Learning for Cybersecurity in Healthcare: A Mulesoft-Enabled Approach. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-6). IEEE.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

[57] Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning (2nd ed.). MIT Press.
[58] Siva Hemanth Kolla. (2023). Deep Learning–Driven Retrieval-Augmented Generation for Enterprise ITSM Automation: A Governance-Aligned Large Language Model Architecture . Journal of Computational Analysis and Applications (JoCAAA), 31(4), 2489–2502. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/4774
[59] Tsamados, A., Aggarwal, N., Cowls, J., et al. (2022). The ethics of algorithms. AI & Society, 37, 215–230.
[60] Davuluri, P. S. L. N. . (2024). AI-Driven Data Governance Frameworks for Automated Regulatory Reporting and Audit Readiness. Metallurgical and Materials Engineering, 30(4), 996–1010. Retrieved from https://metall-mater-eng.com/index.php/home/article/view/1936
[61] Wooldridge, M. (2009). An introduction to multiagent systems (2nd ed.). Wiley.
[62] Bandi, V. D. V. K. (2023). Production-Grade Machine Learning Pipelines For Healthcare Predictive Analytics. South Eastern European Journal of Public Health, 189–205. Retrieved from https://www.seejph.com/index.php/seejph/article/view/7057
[63] Zhang, A., Xing, L., Zou, J., & Wu, J. C. (2022). Shifting ML for healthcare to deployment. Nature Biomedical Engineering, 6, 1330–1345.
[64] Kolla, S. K. (2021). Architectural Frameworks for Large-Scale Electronic Health Record Data Platforms. Current Research in Public Health, 1(1), 1–19. Retrieved from https://www.scipublications.com/journal/index.php/crph/article/view/1372
[65] Benford, S., et al. (2009). Emergent multi-agent architectures. Autonomous Agents and Multi-Agent Systems, 18, 15–45.
[66] Inala, R. (2025). A Unified Framework for Agentic AI and Data Products: Enhancing Cloud, Big Data, and Machine Learning in Supply Chain, Insurance, Retail, and Manufacturing. EKSPLORIUM-BULETIN PUSAT TEKNOLOGI BAHAN GALIAN NUKLIR, 46(1), 1614-1628.
[67] Ferber, J. (1999). Multi-agent systems: An introduction. Addison-Wesley.
[68] Garapati, R. S., & Daram, D. S. B. (2025). AI-Enabled Predictive Maintenance Framework For Connected Vehicles Using Cloud-Based Web Interfaces. Available at SSRN 5524261.
[69] Kephart, J. O., & Chess, D. M. (2003). The vision of autonomic computing. Computer, 36(1), 41–50.
[70] Aitha, A. R., & Jyothi Babu, D. A. (2025). Agentic AI-Powered Claims Intelligence: A Deep Learning Framework for Automating Workers Compensation Claim Processing Using Generative AI. Available at SSRN 5505223.
[71] Huhns, M. N., & Singh, M. P. (1998). Readings in agents. Morgan Kaufmann.
[72] Nagabhyru, K. C., & Babu, A. J. Human In The Loop Generative AI: Redefining Collaborative Data Engineering For High Stakes Industries.
[73] Erl, T. (2016). Microservices design patterns. Prentice Hall.
[74] Gottimukkala, V. R. R. (2025). Generative AI for Exceptions and Investigations: Streamlining Resolution Across Global Payment Systems. Journal of International Commercial Law and Technology, 6(1), 969-972.
[75] Fowler, M. (2018). Refactoring (2nd ed.). Addison-Wesley.
[76] Segireddy, A. R. (2025). GENERATIVE AI FOR SECURE RELEASE ENGINEERING IN GLOBAL PAYMENT NETWORK. Lex Localis: Journal of Local Self-Government, 23.
[77] Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1994). Design patterns. Addison-Wesley.
[78] Amistapuram, K. (2025). GENERATIVE AI FOR CLAIMS EXCEPTIONS AND INVESTIGATIONS: ENHANCING RESOLUTION EFFICIENCY IN COMPLEX INSURANCE PROCESSES. Available at SSRN 5785482.
[79] Zaharia, M., et al. (2010). Spark: Cluster computing with working sets. HotCloud.
[80] Rongali, S. K., & Varri, D. B. S. (2025). AI in health care threat detection. World Journal of Advanced Research and Reviews, 25(3), 1784-1789.
[81] Lakshman, A., & Malik, P. (2010). Cassandra. ACM SIGOPS Operating Systems Review, 44(2), 35–40.
[82] Nagubandi, A. R. (2025). PIONEERING SELF-ADAPTIVE AI ORCHESTRATION ENGINES FOR REAL-TIME END-TO-END MULTI-COUNTERPARTY DERIVATIVES, COLLATERAL, AND ACCOUNTING AUTOMATION: INTELLIGENCE-DRIVEN WORKFLOW COORDINATION AT ENTERPRISE SCALE. Lex Localis, 23(S6), 8598-8610.
[83] Stonebraker, M., & Çetintemel, U. (2005). One size fits all? ICDE Proceedings, 2–11.
[84] Yandamuri, U. S. (2022). Big Data Pipelines for Cross-Domain Decision Support: A Cloud-Centric Approach. International Journal of Scientific Research and Modern Technology, 227.
[85] Moreira, M. W. L., et al. (2018). IoT-based smart healthcare systems. Sensors, 18(4), 1155.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 8, Issue 6, November-December 2025||

DOI:10.15662/IJRPETM.2025.0806035

[86] Guntupalli, R. (2025). Multi-Cloud vs. Hybrid Cloud Security: Key Challenges and Best Practices. Hybrid Cloud Security: Key Challenges and Best Practices (November 21, 2025).
[87] Mell, P., & Grance, T. (2011). The NIST definition of cloud computing. NIST.
[88] Rongali, S. K. (2025, August). AI-Powered Threat Detection in Healthcare Data. In 2025 International Conference on Artificial Intelligence and Machine Vision (AIMV) (pp. 1-7). IEEE.
[89] World Health Organization. (2021). Ethics and governance of artificial intelligence for health. WHO Press.
[90] Kolla, S. H. (2024). RETRIEVAL-AUGMENTED GENERATION WITH SMALL LLMS FOR KNOWLEDGE-DRIVEN DECISION AUTOMATION IN ENTERPRISE SERVICE PLATFORMS. Turkish Journal of Computer and Mathematics Education (TURCOMAT), 15(3), 476–486. https://doi.org/10.61841/turcomat.v15i3.15497
[91] Moreau, L., et al. (2015). The W3C PROV family of specifications. Future Generation Computer Systems, 29(7), 161–165.
[92] Davuluri, P. N. Integrating Artificial Intelligence into Event-Driven Financial Crime Compliance Platforms.
[93] Van Roy, P. (2009). Self-management in distributed systems. IEEE Computer, 42(12), 40–47.
[94] Vardhan Kumar Bandi, V. D. (2024). Automated Feature Engineering Systems in Large-Scale Healthcare Data Environments. Journal of Neonatal Surgery, 13(1), 2127–2141. Retrieved from https://www.jneonatalsurg.com/index.php/jns/article/view/10004
[95] Sutton, R. S. (2019). The bitter lesson. Incomplete Ideas Blog.
[96] Kolla, S. K. (2021). Designing Scalable Healthcare Data Pipelines for Multi-Hospital Networks. World Journal of Clinical Medicine Research, 1(1), 1–14. Retrieved from https://www.scipublications.com/journal/index.php/wjcmr/article/view/1376