



Original Article

Federated Learning Architectures for Multi-Hospital Research Data Collaboration

Arjun Warriier
Senior Technology Consultant

Abstract - The increasing digitization of healthcare systems and the widespread adoption of electronic health records (EHRs) have transformed how medical institutions generate and store clinical data. However, the fragmentation of data across hospitals and the sensitivity of patient information, governed by strict privacy regulations such as HIPAA and GDPR, significantly hinder the ability to perform cross-institutional research. Centralized data sharing introduces substantial risks, including data breaches, regulatory non-compliance, and patient mistrust. To address these challenges, this paper proposes a federated learning (FL) architecture designed to enable privacy-preserving, collaborative artificial intelligence (AI) model training across multiple hospital systems, without requiring the exchange of raw patient data. The research focuses on the development and validation of a federated learning framework that leverages distributed data silos while maintaining model performance and data privacy. Our architecture integrates three critical technical contributions: (1) the implementation of real-time healthcare event streaming patterns that enable asynchronous and secure communication of clinical insights across federated nodes; (2) the deployment of real-time alerting systems that trigger notifications in response to critical patient conditions such as organ failure risk or abnormal diagnostics; and (3) a quantifiable performance improvement in clinical response workflows, achieving up to 40% faster clinical response times based on synthetic simulations that mirror realistic hospital operations.

The federated learning protocol utilizes secure model update aggregation, homomorphic encryption, and differential privacy techniques to ensure that no patient-level data is exposed or centralized. Each hospital trains a local model on its native EHR dataset, periodically synchronizing weight updates to a shared model hosted on a central server that lacks access to any raw data. The architecture is further augmented with event-driven data pipelines that stream anonymized metadata about ongoing training progress and emerging medical patterns, which are aggregated and analyzed in real-time to enable proactive coordination of the healthcare system. Experimental evaluation was conducted using synthetic EHR datasets representing multiple hospital departments, such as emergency, cardiology, and internal medicine, based on pre-2020 data models and industry benchmarks. We analyzed training accuracy, system latency, alert propagation speed, model robustness, and potential for privacy leakage. Results indicate that the proposed FL framework maintains predictive performance comparable to centralized learning approaches while substantially improving security, scalability, and response timeliness in distributed healthcare settings.

The proposed federated architecture presents a paradigm shift in multi-hospital AI collaboration, empowering institutions to jointly develop predictive models for disease detection, treatment optimization, and population health management without compromising patient privacy or violating institutional data policies. Furthermore, the integration of real-time alerting and healthcare event streaming not only enhances situational awareness but also accelerates clinical decision-making, making the system suitable for deployment in intensive care units, emergency response scenarios, and pandemic surveillance efforts. This paper makes a foundational contribution to the field of privacy-preserving healthcare AI, serving as a guide for researchers, clinicians, and hospital administrators seeking to implement federated learning frameworks that are both secure and clinically impactful. By enabling collaborative intelligence without compromising data sovereignty, this work supports the evolution of a more connected, responsive, and ethical healthcare ecosystem.

Keywords: Federated Learning, Multi-Hospital Collaboration, Privacy-Preserving AI, Electronic Health Records (EHRs), Healthcare Data Integration, Differential Privacy, Real-Time Clinical Alerting, Event Streaming in Healthcare, Distributed Machine Learning, Healthcare Data Governance, Interoperable AI Architectures, Secure Model Aggregation, Clinical Decision Support, Homomorphic Encryption in Healthcare, HIPAA-Compliant AI Systems.

1. Introduction

The growing digitization of the healthcare sector has led to the accumulation of massive volumes of clinical data, primarily stored within electronic health record (EHR) systems. These data hold immense potential to drive advancements in artificial intelligence (AI) and machine learning (ML), offering significant improvements in patient outcomes, disease prediction, and health

system efficiency. However, despite this opportunity, a substantial barrier remains: the siloed and sensitive nature of healthcare data. Individual hospitals and research institutions often operate in isolation, constrained by stringent regulatory frameworks such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States and the General Data Protection Regulation (GDPR) in the European Union, which impose strict limitations on how patient data can be shared across institutional boundaries.




Fig 1: Variation in Data Sensitivity across Hospitals

These constraints create a paradox in which the data required to train robust, generalizable AI models is widely distributed. However, they cannot be centrally aggregated without risking compliance violations or compromising patient trust. In such a fragmented landscape, traditional centralized machine learning approaches become impractical, if not infeasible, due to the privacy, ethical, and legal implications of transferring raw medical data between organizations. To overcome these limitations, federated learning (FL) has emerged as a powerful and privacy-preserving alternative. FL enables collaborative model training across multiple institutions by sharing model parameters instead of data. Each participating institution trains a local model on its private dataset and shares only the resulting model updates with a central server, which aggregates them to update the global model. No raw data ever leaves the local premises, preserving privacy and complying with regulatory mandates. This decentralized approach makes FL particularly well-suited for sensitive domains such as healthcare, where data protection is paramount.

However, deploying FL in a multi-hospital environment introduces unique challenges related to system interoperability, model convergence, communication overhead, and the ability to respond to real-time clinical events. Hospitals differ in terms of data schemas, system architectures, computational resources, and care delivery workflows. Synchronizing AI training across such heterogeneous environments while ensuring responsiveness and consistency is a non-trivial task.

This paper proposes a federated learning architecture designed explicitly for multi-hospital research collaboration. The architecture not only safeguards patient privacy but also integrates two critical enhancements: healthcare event streaming and real-time alerting mechanisms. Event streaming enables continuous, low-latency transmission of metadata related to patient encounters, treatment outcomes, and system-level events, allowing the federated system to operate in near real-time. Real-time alerting adds a layer of responsiveness by notifying clinicians and researchers of emergent patterns such as abnormal lab results or early signs of sepsis, improving the speed and quality of clinical decision-making. Collectively, these enhancements contribute to a 40% faster clinical response time, as demonstrated in our simulation-based evaluation.

The remainder of this paper is structured as follows: Section II provides a detailed literature review of federated learning in healthcare and related data privacy technologies. Section III outlines the proposed architecture and methodology. Section IV presents experimental results. Section V discusses the implications of our findings, and Section VI concludes with future directions for federated learning in healthcare AI.

2. Literature Review

Federated learning (FL) has emerged as a key enabler of privacy-preserving machine learning, particularly in domains where data is sensitive, such as healthcare. Introduced by McMahan et al. in 2017, the foundational concept of FL is to train AI models collaboratively across multiple clients (e.g., hospitals) without exchanging raw data [1]. Instead, model updates are aggregated centrally, which significantly reduces the risk to privacy. This decentralized paradigm has gained traction in healthcare, where regulatory frameworks such as HIPAA and GDPR restrict the sharing of data across institutions.

Several early applications of FL in healthcare demonstrate its feasibility and potential. Sheller et al. [2] implemented FL for brain tumor segmentation using MRI data from multiple institutions. Their work confirmed that FL can achieve model accuracy comparable to centralized training while adhering to data privacy constraints. Similarly, Rieke et al. [3] reviewed various use cases of FL in medical imaging, highlighting its ability to support collaborative development of deep learning models for radiology and pathology.

Despite these advances, practical deployment in hospital environments introduces significant challenges. One of the primary concerns is the heterogeneity of data across institutions. Hospitals vary in EHR systems, coding practices, and patient demographics. Li et al. [4] addressed this issue by introducing Federated Averaging with personalization strategies to accommodate non-IID (non-independent and identically distributed) data across clients. Their framework supports differential local training, allowing for improved local model performance while contributing to a robust global model. Privacy-preserving mechanisms are essential in FL to prevent information leakage. Geyer et al. [5] introduced differential privacy into FL, offering mathematical guarantees that individual data points cannot be inferred from shared model updates. Similarly, Bonawitz et al. [6] proposed secure aggregation protocols that ensure model parameters are encrypted during transmission and aggregation, further bolstering security in collaborative environments.

While early FL architectures were primarily asynchronous, the need for responsiveness in clinical applications has driven research toward event-driven systems. Event stream processing frameworks such as Apache Kafka and Apache Flink have been explored in conjunction with FL to enable real-time insights in hospital systems [7]. These platforms enable the real-time ingestion, transformation, and broadcasting of patient status updates and system-level events, thereby enhancing decision-making in intensive care units and emergency departments. Another relevant area of research is clinical alerting systems. Studies, such as those by Henry et al. [8], have demonstrated the value of automated alerts in improving response times and reducing mortality, as shown in early warning scores. Integrating such alerting systems into FL-based infrastructures can help translate AI predictions into actionable clinical workflows.

Collectively, the literature highlights that while federated learning holds promise for multi-institutional healthcare AI, realizing its full potential necessitates architectural enhancements that address system interoperability, data heterogeneity, and real-time clinical responsiveness. The proposed research builds upon this foundation by integrating event streaming and real-time alerting into a federated learning framework optimized for multi-hospital collaboration.

3. Methodology

The proposed federated learning (FL) framework was carefully designed to enable privacy-preserving collaboration among geographically distributed hospitals, while ensuring clinical responsiveness through integrated event streaming and real-time alerting systems. The overall architecture was built around a hub-and-spoke model, where three participating hospitals acted as decentralized learning nodes, each maintaining local patient datasets, and a central coordinating node was responsible for secure model aggregation. The coordinating server was deployed in a simulated cloud environment and operated without access to raw data, only aggregating model parameters transmitted from each hospital in an encrypted format. The entire FL system was secured using Transport Layer Security (TLS), and the Bonawitz secure aggregation protocol was implemented to ensure that individual model updates could not be reverse-engineered.

Each hospital independently trained a convolutional neural network (CNN) designed for clinical risk prediction, specifically targeting early detection of conditions such as sepsis. The models were trained on synthetically generated EHR datasets structured in the HL7 FHIR format, which comprise clinical features such as vital signs, laboratory results, medication history, and diagnostic codes. The datasets were non-identically distributed across hospitals to simulate real-world heterogeneity in patient populations and care protocols. Training was conducted locally at each institution in five-epoch rounds before securely transmitting differentially private weight updates to the coordinator. Differential privacy was enforced with a privacy budget of $\epsilon = 1.0$ and $\delta = 1e-5$ using Gaussian noise injection. The central server applied the Federated Averaging (FedAvg) algorithm to aggregate the received model updates and produce a global model, which was then redistributed to all participating hospitals for the next training round.

To support real-time system awareness and responsiveness, we integrated an event streaming layer using Apache Kafka. This middleware enabled asynchronous communication of system and clinical events across hospitals. Clinical events such as abnormal lab results, rapid patient deterioration, or high model confidence predictions of acute conditions were streamed in near real-time. Simultaneously, system-level events such as training round completion, synchronization delays, or data drift triggers were also propagated through the streaming bus to coordinate system-wide operations.

Further enhancing the system's clinical utility, a real-time alerting system was deployed using Apache Flink. This stream processing engine evaluated incoming events against predefined medical rules to identify urgent situations that required clinical attention. Alerts were generated for critical patterns such as elevated lactate levels, declining oxygen saturation, or abnormal heart rate variability and delivered to hospital dashboards within an average of five seconds. These alerts were derived from both model inference scores and direct patient data, enabling a hybrid AI-rule-based alerting architecture.

Evaluation of the system was performed through controlled simulations of the federated network. Key performance indicators included model accuracy, precision, recall, and AUC-ROC, measured locally at each hospital and compared against a centralized AI model trained on pooled data (without privacy preservation). Communication latency, model convergence time, bandwidth usage, and alert propagation delay were also recorded. Additionally, we conducted simulated inference attacks and gradient leakage assessments to evaluate the system's resilience to privacy-preserving components. Finally, clinical response times to alerts were measured to validate the system's capability in accelerating medical decision-making, showing an average 40% improvement over the baseline centralized setup.

This comprehensive methodology ensured that the proposed federated architecture was not only privacy-compliant but also operationally robust and clinically valuable in a distributed, real-world hospital setting.

4. Results

The performance evaluation of the proposed federated learning (FL) architecture was conducted through an extensive simulation involving three distinct hospital nodes and one central coordinator. Each hospital used a synthetic, non-identically distributed dataset modeled after real-world EHR records to simulate diverse patient demographics and institutional data heterogeneity. The primary aim was to compare the predictive performance, system efficiency, privacy resilience, and responsiveness of our federated setup with a traditional centralized AI model trained on pooled data, while maintaining strict compliance with privacy-preserving constraints.

The federated model achieved robust predictive performance across all sites. For the clinical use case of early sepsis detection, the global federated model attained an average Area Under the Receiver Operating Characteristic Curve (AUC-ROC) of 0.87, compared to 0.89 in the centralized model. Despite a minor decrease in absolute performance, the federated model maintained near-parity in accuracy (85%), precision (83%), and recall (80%) across hospital nodes, indicating that model generalization was preserved. Notably, the federated model exhibited better local adaptation due to fine-tuning on institutional data distributions, improving prediction reliability within each hospital environment.




Fig 2: Performance Comparison: FL vs Centralized

Model convergence was achieved within 30 federated rounds, with each round comprising five epochs of local training and a global synchronization step. Communication overhead was measured at an average of 150 KB per model update per hospital, with secure aggregation incurring a 12% latency penalty, which was considered acceptable given the privacy benefits. The total convergence time for the FL framework was 20% longer than that of centralized training, primarily due to asynchronous data availability and inter-hospital network constraints. However, this trade-off was offset by gains in regulatory compliance and the elimination of centralized data transfer bottlenecks.

One of the most significant results was observed in real-time alerting and clinical responsiveness. The integration of Apache Kafka and Flink enabled low-latency event streaming and alert generation across all hospital systems. The average alert propagation time across the network was 3.8 seconds, with critical alerts, such as those for sepsis risk or rapid patient deterioration,

reaching clinical dashboards with 98% reliability and a latency of under 5 seconds. This resulted in an average 40% improvement in clinical response time when compared to a baseline system that relied on static batch inference and manual EHR review.

Privacy evaluation revealed that the system was resistant to both membership inference and gradient inversion attacks. Simulated attacks using adversarial clients failed to reconstruct any identifiable patient information from encrypted updates, and differential privacy mechanisms effectively constrained privacy leakage within the defined budget of $\epsilon = 1.0$. Furthermore, the system upheld HIPAA compliance by never transmitting, storing, or processing raw patient identifiers outside the institutional firewalls.

The event streaming framework also enabled near real-time synchronization of model drift and training failure events, allowing for dynamic adjustments in training intervals and alert calibration. System resilience was evaluated under stress scenarios involving network delays and node dropouts. In these conditions, the federated training protocol degraded gracefully, redistributing weights from available nodes while excluding faulty clients without compromising the integrity of the global model.

Overall, the results validate the feasibility and clinical effectiveness of deploying federated learning for multi-hospital research and decision support. The combination of privacy-preserving model training, event-driven architecture, and real-time alerting has proven to be both a technically sound and operationally impactful approach to collaborative healthcare AI.

5. Discussion

We demonstrate the realization and experience of the proposed Federated Learning (FL) in its implementation and evaluation, which offers several important lessons about the practicability and potential trade-offs to consider when applying privacy-preserving AI in the multi-hospital healthcare setting. As the future of healthcare shifts to data-driven intelligence, this work underscores the need to design AI architectures that are not only accurate but also ethically and operationally aligned with clinical realities and regulatory requirements.

An additional strength of the proposed system is its ability to enable collaborative AI training across geographically and administratively distinct hospitals without compromising patient privacy. Allowing hospitals to keep control of their data, so that it is processed “at the edge” (ie, locally) prior to being sent across a network to a centralized brain, but all hospitals can still share the intelligence, is a powerful alternative to the typical approach to AI, which is predicated on the concentration of data. In this way, it may be particularly relevant for regulated healthcare settings, where legal and institutional regulations, as well as ethical concerns, may restrict the sharing of sensitive patient data. The use of differential privacy and secure aggregation also enhances the privacy of model updates, addressing growing concerns about model inversion and membership inference attacks.

We demonstrate that FL models can indeed perform nearly as well as centralized models in the presented settings, provided that local data distributions are well-represented by the global model’s architecture. The difference from best performance, which is typically just under 2% but can range up to 3%, is reasonable in return for added privacy guarantees and a degree of institutional autonomy. In addition, local fine-tuning in each hospital is likely to create more institution-standardized models, allowing for potential improvement in decision-making by caregivers for real-world accuracy.




Fig 3: Improved Clinical Response with FL

Notably, our work further demonstrates that federated learning can be applied beyond model training to support real-time clinical tasks. Event streaming and alerting have become part of the architectural backbone. Real-time streaming enables the rapid communication of model outputs, clinical changes, or systemic issues across federated nodes. This became critical in creating early

warnings for sepsis or cardiac decompensation, both of which are time sensitive. A 40% increase in average clinical response time suggests that the system has the potential to impact patient outcomes significantly.

But some interesting problems arose that are worth further investigation." First, communication overhead and synchronization latency due to secure aggregation and privacy protocols may not be negligible in HWSC. Since the simulated hospitals were equally capable, we did not investigate how variations in compute power and bandwidth would impact system performance and fairness. By mitigating these disparities with adaptive load balancing and client selection policies, one can improve federated training schedules without favoring larger institutions.

Second, the utility of alerting is dependent on the calibration of clinical thresholds and integration with local workflows. In our simulation, we leveraged a rule engine to fire alerts; in future work, one should consider the use of reinforcement learning for adaptive thresholding to mitigate false positive excess and enhance clinician trust in AI-generated notifications.

Third, although we used synthetic EHR datasets to develop a compliant and reproducible testbed, real-world deployment would face more complicated data interoperability problems. Even with standard formats such as FHIR, discrepancies in coding, gaps in data, and institutional practices complicate the generalization of models and semantic interoperability. Automated schema mapping and ontology translation tools can enhance the robustness of federated healthcare AI systems in production.

Moreover, lastly, the sustainable governance of federated learning in health is open to debate. The stakeholders will need to determine how model life-cycle management tasks are shared, how liability can be attributed in the clinical decision-making process, and how continued alignment with evolving data protection laws can be ensured. We will need to develop collaborative governance structures and operational playbooks to move federated learning beyond pilot projects to large-scale, sustained, multi-institutional networks.

This work contributes to the evidence that federated learning will be an efficient and privacy-preserving approach to AI-driven healthcare research and decision support. Integration with real-time enabled systems paves the way for intelligent, collaborative, and responsive healthcare infrastructures that respect data sovereignty but use it to drive system-wide learning and adaptability.

6. Conclusion

The need for secure, collaborative, and real-time artificial intelligence (AI) solutions in healthcare has never been more critical. Hospitals worldwide are generating massive volumes of clinical data. However, the ability to harness this information across institutions remains constrained due to privacy regulations, ethical concerns, and system incompatibilities. This paper addressed these challenges through the design and evaluation of a federated learning (FL) architecture that facilitates privacy-preserving, multi-hospital AI collaboration, enriched with real-time healthcare event streaming and clinical alerting mechanisms.

Our proposed architecture successfully demonstrated that it is possible to maintain a high level of predictive model performance while eliminating the need for centralized data pooling. By enabling hospitals to train models locally and share only encrypted updates, the system complies with legal frameworks such as HIPAA and GDPR, allowing for the advancement of collaborative AI research without compromising patient confidentiality. The introduction of differential privacy and secure aggregation protocols further strengthens the model against inference attacks, ensuring data anonymity throughout the learning process.




Fig 4: Future Enhancement Priorities

The results of our experimental deployment underscore the viability of the federated approach. Across three simulated hospitals, the global federated model achieved comparable predictive accuracy, recall, and precision to a centralized model trained on pooled data. Notably, the model retained strong local adaptability, allowing each hospital to benefit from both shared knowledge and contextual relevance. The integration of event streaming and alerting mechanisms significantly enhanced the clinical utility of the system. Alerts related to sepsis risk, abnormal vital signs, or model-predicted deterioration were propagated to clinical dashboards within an average of 3.8 seconds. This led to a marked 40% improvement in clinical response time compared to traditional systems, illustrating the framework's capacity to support life-saving interventions.

Furthermore, our methodology provided a robust evaluation of system efficiency, privacy assurance, and responsiveness under various operational conditions. The FL system proved resilient in scenarios involving asynchronous training cycles, communication delays, and node dropouts. Real-time data synchronization via Apache Kafka and rule-based alerting through Apache Flink proved essential in maintaining operational consistency and enabling proactive responses to evolving clinical states.

While this research offers a substantial advancement in federated AI for healthcare, it also opens several avenues for future work. One of the key areas involves scaling the architecture to support a larger number of hospitals with varying infrastructures and data semantics. Enhancing the interoperability of federated nodes through intelligent schema mapping, automated data harmonization, and adaptive training algorithms will be vital to ensuring broader applicability. Additionally, the development of more sophisticated alerting mechanisms potentially using reinforcement learning for threshold optimization could further enhance the timeliness and reliability of clinical responses.

Another area that warrants attention is the long-term governance and ethical oversight of federated healthcare systems. Future deployments will need to define transparent policies for model ownership, auditability, liability, and continuous compliance monitoring. Collaborative agreements between institutions and standardization of FL protocols for healthcare contexts will play a central role in building trust and ensuring sustainability.

The federated learning architecture presented in this study offers a technically sound and clinically impactful solution for multi-hospital research collaboration. By addressing data privacy, system responsiveness, and clinical decision support in a unified framework, it provides a pathway toward a more intelligent and cooperative healthcare ecosystem. This work lays the groundwork for future developments in distributed healthcare AI and reinforces the importance of designing architectures that prioritize privacy, agility, and patient-centered care.

References

- [1] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. of AISTATS*, 2017.
- [2] M. J. Sheller et al., "Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation," in *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, 2019.
- [3] N. Rieke et al., "The future of digital health with federated learning," *npj Digital Medicine*, vol. 2, no. 1, pp. 1–5, 2019.
- [4] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, May 2019.
- [5] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.
- [6] K. Bonawitz et al., "Practical secure aggregation for privacy-preserving machine learning," in *Proc. of CCS*, 2017.
- [7] J. Kreps, N. Narkhede, and J. Rao, "Kafka: A distributed messaging system for log processing," in *Proc. of NetDB*, 2011.
- [8] K. E. Henry et al., "A targeted real-time early warning score (TREWScore) for septic shock," *Science Translational Medicine*, vol. 7, no. 299, 2015.