

ORIGINAL RESEARCH

Open access

A Conceptual Framework for Federated Learning in Acute Kidney Injury Prediction

Maria Hernandez^{1*}, Carlos Vega¹

Abstract

Acute kidney injury (AKI) is a common and serious condition in critical care, making early prediction essential for timely intervention, reduced mortality, and lower healthcare costs. Machine learning methods using electronic health records have shown promise in identifying at-risk patients, but their performance is often limited by reliance on single-institution datasets and poor generalizability across populations. Privacy regulations such as HIPAA and GDPR further restrict cross-hospital data sharing, hindering the development of more robust models.

To address these challenges, this study proposes a federated learning–based framework for AKI prediction, enabling multiple hospitals to collaboratively train models without exchanging raw patient data. Each institution acts as a local client that trains on its own data and shares only model updates, which are aggregated into a global model. The framework incorporates standardized feature processing, secure aggregation, and communication-efficient strategies to ensure scalability across heterogeneous healthcare environments.

This privacy-preserving approach improves model generalization by leveraging diverse multi-institutional data while maintaining regulatory compliance. Although it introduces challenges such as communication overhead and convergence complexity, these are mitigated through optimized aggregation methods. Overall, the proposed framework enhances predictive performance, supports clinical decision-making, and offers a scalable foundation for future privacy-aware healthcare AI systems in AKI management.

Keywords Federated learning, Conceptual framework, Acute kidney injury, Privacy-preserving AI, Healthcare prediction, Multi-institutional collaboration

*Correspondence:

Maria Hernandez
maria.hernandez@gmail.com

¹ Department of Artificial Intelligence in Medicine, University of Valencia, Valencia, Spain

Introduction

Acute kidney injury (AKI) is a sudden deterioration in kidney function that is diagnosed using standardized criteria such as the KDIGO guidelines [1]. This condition occurs frequently in hospitalized patients, especially those in intensive care or undergoing major procedures, leading to significant morbidity [2]. The associated mortality rates can exceed 50 percent in severe cases, highlighting the urgency of early detection [3]. Epidemiological data reveal that AKI contributes substantially to prolonged hospital stays and increased healthcare expenditures [4]. Effective

prediction strategies are therefore critical for mitigating these impacts in clinical practice.

Existing machine learning models for AKI prediction have been developed using techniques ranging from logistic regression to ensemble methods in single-center cohorts [5]. These models incorporate preoperative, intraoperative, and laboratory data to forecast AKI risk with varying degrees of success [6]. Nevertheless, their performance often degrades when applied to new populations due to differences in case mix and clinical protocols [7]. Limitations arise primarily from the inability to access large,

diverse datasets under current data governance rules [8]. Consequently, innovative decentralized frameworks are needed to enhance model robustness [9].

Privacy constraints governed by regulations such as HIPAA and GDPR prohibit the centralized pooling of sensitive patient data across healthcare institutions [10]. These legal frameworks prioritize patient confidentiality and data minimization, creating obstacles for traditional machine learning development [11]. Institutional policies additionally enforce data silos to prevent unauthorized access and breaches [12]. Such barriers result in fragmented research efforts that limit the scale and diversity of training data for AKI models [13]. Federated learning circumvents these issues by design, allowing collaboration without data transfer [14].

This manuscript introduces a conceptual framework for federated learning in the context of acute kidney injury prediction to enable privacy-compliant multi-institutional collaboration [15]. The framework addresses key gaps in existing approaches by combining decentralized training with clinical relevance [16]. It provides a structured methodology for implementing federated systems in healthcare settings [9]. Subsequent sections explore the background, framework architecture, components, training processes, and related considerations [17]. This organization offers a complete conceptual guide for advancing AI applications in AKI management [18]

Background

AKI definition and prediction

The KDIGO criteria provide a widely accepted definition for AKI based on increases in serum creatinine or decreases in urine output over specified time periods [1]. Common risk factors encompass sepsis, cardiovascular surgery, and exposure to nephrotoxic agents in critical care environments [2]. Machine learning models have been applied to predict AKI development using these clinical indicators in various patient subgroups [3]. These predictive tools aim to support early intervention strategies that can alter disease trajectories [4]. Integration of such models into electronic health systems represents a promising avenue for improved care [8].

Prediction models for AKI have demonstrated utility in settings such as pediatric critical care and post-cardiac surgery cohorts [10]. Ensemble learning techniques have

been explored for sepsis-associated AKI to enhance accuracy [7]. Despite these advances, single-institution models frequently lack the robustness required for widespread clinical adoption [5]. The conceptual framework proposes federated learning to aggregate insights from multiple sources [18]. This strategy promises to elevate the standard of AKI risk assessment [17].

Privacy constraints in healthcare

Privacy constraints in healthcare are enforced through comprehensive regulations including HIPAA and GDPR that restrict data sharing between organizations [10]. These policies are designed to safeguard personal health information from misuse while enabling ethical research [12]. Centralization of datasets is often not permissible due to the risk of re-identification and regulatory violations [11]. As a result, healthcare AI development is frequently confined to siloed environments [13]. Federated learning offers a pathway to collaborative modeling that fully complies with these mandates [19].

Institutional barriers further exacerbate the challenges of data aggregation for AKI prediction studies [20]. Policies at the hospital level prioritize data security and patient consent, limiting external collaborations [21]. The inability to share raw records hampers the creation of comprehensive predictive models [10]. Privacy-aware methodologies are therefore essential for progressing multi-center initiatives [11]. The framework incorporates these considerations to ensure regulatory alignment [15].

Federated learning fundamentals

Federated learning fundamentals center on performing model training locally at each client institution before aggregating updates centrally [14]. This paradigm ensures that raw patient data remains confined to its original location throughout the process [15]. Model parameters are shared and combined using algorithms such as FedAvg to form a global representation [16]. The approach has been conceptualized for numerous healthcare scenarios involving sensitive information [9]. It fundamentally transforms how collaborative AI is conducted in regulated domains [10].

No data sharing occurs in federated learning, as only model updates or gradients are exchanged during communication rounds [22]. This property aligns perfectly with healthcare privacy requirements for conditions like AKI [12].

Extensions of the basic framework address additional challenges in medical applications [23]. The fundamentals support the development of generalizable models from distributed sources [17]. Such principles form the basis for the proposed AKI prediction system [18].

Framework Overview

High-level architecture

The high-level architecture of the framework designates hospitals as distributed client nodes connected to a central coordinating server [24]. Clients execute local training on their proprietary datasets and forward only the resulting model updates [25]. The server then performs aggregation to refine the global model before redistributing it for the next round [26]. This flow maintains strict separation between data and model collaboration [27]. The architecture is optimized for secure and scalable operations in healthcare networks [22].

Figure 1 illustrates the hierarchical federated learning architecture enabling privacy-preserving multi-institutional collaboration for AKI prediction.

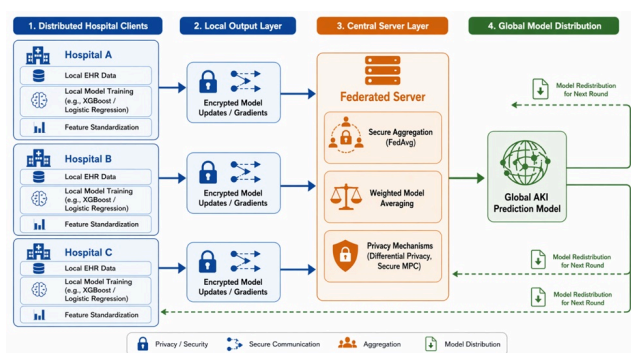


Figure 1. Hierarchical Federated Learning Architecture for Privacy-Preserving AKI Prediction

Communication flows are managed to accommodate the operational constraints of clinical environments [14]. Each hospital contributes to the collective model while retaining full control over its data [16]. The design supports varying levels of participation without disrupting ongoing care activities [9]. Overall, the architecture provides a practical foundation for federated AKI prediction [15]. It ensures efficient knowledge transfer across institutional boundaries [23].

Core assumptions

Core assumptions of the framework include the availability of compatible electronic health record systems across participating hospitals [15]. Sufficient local data volumes at each site are necessary to enable effective model training and generalization [18]. Network reliability is required to support consistent communication during training phases [24]. These assumptions ensure the operational viability of the federated setup [17]. Additional considerations involve alignment with clinical workflows for AKI monitoring [1].

EHR compatibility facilitates standardized feature extraction without the need for data centralization [16]. Adequate local datasets allow sites to capture unique population characteristics [23]. Reliable networks minimize interruptions and support timely model synchronization [26]. The framework is built upon these foundational elements to maximize effectiveness [14]. Meeting the assumptions enables the realization of privacy-preserving predictive capabilities [10].

Design principles

Design principles emphasize a privacy-first approach that prevents any exposure of raw clinical data [10]. Communication efficiency is prioritized to reduce bandwidth demands and accelerate training cycles [22]. Clinical validity is upheld by incorporating AKI-specific risk factors and outcomes into the model design [2]. These principles guide every aspect of the framework development [12]. The result is a system that is both secure and relevant to healthcare providers [19].

Privacy-first principles are non-negotiable to build trust among collaborating institutions [11]. Communication-efficient strategies such as selective updates help manage resource limitations [24]. The framework ensures clinical validity through alignment with KDIGO criteria and established prediction practices [3]. Such principles distinguish the approach from less regulated AI methods [9]. They collectively promote sustainable and ethical implementation [13].

Framework Components

Local model architecture

Local model architecture within the framework supports a range of machine learning types including ensemble methods like XGBoost and logistic regression for tabular AKI data [5]. These architectures are selected for their

interpretability and suitability to electronic health record features [6]. Each participating hospital configures the model to align with its computational infrastructure and data characteristics [7]. Flexibility in local designs accommodates institutional differences while ensuring aggregation compatibility [18]. The choice of architecture prioritizes both predictive performance and clinical usability [4].

Tree-based models offer robustness to missing data commonly encountered in clinical records for AKI prediction [3]. Logistic regression provides transparency that aids clinician acceptance and regulatory approval [8]. The framework allows for hybrid or specialized local architectures as needed [15]. Local training leverages these models to learn site-specific patterns effectively [16]. This component enhances the framework's adaptability to diverse healthcare settings [17].

Feature standardization

Feature standardization is performed independently at each hospital using local statistical properties to harmonize variables across sites without data sharing [16]. This process involves normalizing laboratory results, vital signs, and demographic features relevant to AKI risk [18]. Federated protocols guide the alignment of feature spaces to ensure consistency in the global model [24]. Such local operations prevent the disclosure of sensitive distributional information [25]. The standardization step is vital for reducing discrepancies that arise from heterogeneous EHR systems [26].

Hospitals apply transformation rules derived from their own data distributions to achieve feature compatibility [27]. Guidelines for common AKI-related variables are distributed centrally to promote uniformity [1]. This approach maintains privacy while enabling effective cross-site learning [10]. Harmonized features contribute to improved model convergence and generalizability [15]. The component addresses a critical challenge in multi-institutional federated systems [9].

Secure aggregation

Secure aggregation in the framework utilizes FedAvg and its variants to combine local model updates into a unified global model [14]. Only processed updates are transmitted, with protections against reconstruction of original data [12]. The server handles aggregation in a manner that obscures

individual client contributions [19]. This mechanism safeguards against privacy leaks during the collaboration process [11]. Weighted aggregation accounts for differences in local dataset sizes [16].

Variants of secure aggregation incorporate additional safeguards for multi-institutional scenarios [22]. The process ensures that the global model reflects collective learning without favoring any single site unduly [24]. Handling of updates includes validation steps to maintain integrity [25]. Secure aggregation is central to the trustworthiness of the federated approach [13]. It enables the framework to deliver reliable AKI predictions collaboratively [23].

Communication protocol

The communication protocol establishes synchronization mechanisms for exchanging model updates between clients and the server at regular intervals [26]. Round timing is calibrated to fit within hospital operational schedules and network capabilities [27]. Dropout handling is integrated to allow partial participation when some institutions experience temporary disconnections [25]. These features ensure uninterrupted progress toward model improvement [24]. The protocol minimizes overhead while maximizing collaboration efficiency [22].

Synchronization can be configured as synchronous or asynchronous depending on reliability requirements [14]. Timing considerations prevent interference with clinical duties at participating sites [16]. Strategies for managing dropouts include fallback aggregation from available clients [15]. The protocol supports robust operation in variable healthcare network conditions [9]. Overall, it contributes to the practical feasibility of the federated training process [10].

Training Process

Initialization and client selection

Initialization of the global model occurs at the central server prior to its distribution to selected client hospitals [9]. Client selection prioritizes institutions with adequate AKI case volumes and compatible data infrastructures [15]. The process aims to achieve balanced representation from varied clinical settings [16]. Participating hospitals are vetted for their ability to contribute meaningfully to the

federation [23]. This initial step establishes the collaborative foundation for subsequent training [17].

The global model is initialized using either a neutral starting point or preliminary knowledge from related domains [14]. Selection of clients ensures diversity in patient demographics and care practices relevant to AKI [18]. Compatibility checks are performed to confirm readiness for federated participation [1]. The approach fosters inclusive collaboration across the healthcare ecosystem [10]. Initialization and selection are critical for achieving effective knowledge synthesis [2].

Local training rounds

Local training rounds involve multiple epochs of optimization on each hospital's private dataset using suitable batch sizes and loss functions [15]. The loss function is configured to address the binary or multi-class nature of AKI prediction tasks [3]. Batch sizes are adapted to local hardware constraints to optimize computational efficiency [18]. Each round allows the local model to refine its parameters based on site-specific data [22]. This decentralized phase captures unique patterns without external data access [24].

Epoch counts per round are determined to balance local learning depth with global synchronization needs [25]. Loss functions may incorporate class weighting to handle the imbalance typical in AKI occurrence [2]. Local training is conducted in a manner consistent with clinical data quality and volume [4]. The rounds enable progressive improvement of site models prior to aggregation [16]. This process forms the iterative core of the federated learning cycle [14].

Model averaging and convergence

Model averaging is conducted via weighted techniques at the central server following the submission of local updates [14]. Weights reflect the relative contribution of each client's dataset size to the overall federation [16]. The averaged model is then redistributed to clients for the next training cycle [9]. Convergence is evaluated based on stability criteria applied to successive global models [26]. This averaging step synthesizes distributed knowledge into a cohesive predictive tool [15].

Convergence criteria include thresholds on parameter changes or proxy validation metrics across rounds [24]. The weighted averaging process helps alleviate biases

introduced by statistical heterogeneity [18]. Training continues until the global model demonstrates sufficient stability for clinical consideration [22]. The mechanism ensures efficient progression toward a high-quality federated model [23]. Model averaging and convergence complete the training process in a privacy-preserving manner [17].

Privacy Guarantees

What federated learning protects

Federated learning fundamentally protects raw patient data by ensuring that no electronic health records ever leave the originating hospital servers during the entire training lifecycle [10]. Only encrypted model updates or gradients are transmitted to the central server for aggregation, preserving the confidentiality of individual AKI cases [12]. This design directly addresses the core privacy requirements of healthcare systems handling sensitive kidney function data [11]. Hospitals maintain complete sovereignty over their local datasets while still contributing to a shared predictive capability [14]. Such protection mechanisms enable ethical collaboration across institutional boundaries without compromising regulatory standards [15].

The protection extends to preventing any reconstruction of original records from the shared updates through careful protocol design [19]. In the context of AKI prediction, this means that demographic details, laboratory trends, and clinical notes remain isolated at each site [13]. The framework leverages these safeguards to build trust among participating entities [9]. Privacy is maintained even under potential adversarial scrutiny of the communication channels [16]. Overall, the approach redefines collaborative modeling as a data-secure process aligned with healthcare ethics [20].

Additional privacy layers

Additional privacy layers in the framework incorporate differential privacy techniques to add controlled noise to model updates before transmission [12]. Secure multi-party computation can further obscure individual contributions during the aggregation phase at the central server [11]. These enhancements provide quantifiable privacy guarantees while balancing the utility of the global AKI prediction model [21]. Trade-offs arise in the form of slight reductions in convergence speed due to the added

computational overhead [22]. The layers are selectively applied based on the sensitivity of the participating institutions' data environments [19].

Integration of these layers ensures compliance with evolving privacy regulations without altering the core federated workflow [13]. For instance, differential privacy parameters can be tuned locally to match institutional risk tolerances [10]. The framework conceptualizes these additions as modular extensions that enhance baseline protections [14]. Trade-offs are managed through adaptive strategies that minimize impact on model effectiveness [23]. This comprehensive layering strengthens the overall privacy posture for multi-hospital AKI initiatives [9].

Comparison with Centralized Training

Expected performance gap

Centralized training typically aggregates all data in one location to achieve seamless model optimization, yet it incurs unavoidable privacy risks that federated approaches avoid [28]. The expected performance gap in federated settings stems from the privacy-utility trade-off inherent in decentralized aggregation methods [22]. Non-identical data distributions across hospitals may introduce minor divergences compared to a fully pooled dataset [24]. Nevertheless, the framework mitigates this through targeted design principles that prioritize clinical relevance [15]. Such gaps are conceptualized as acceptable costs for enabling collaboration under strict regulatory constraints [14].

Prior centralized models for continuous AKI prediction highlight the limitations of data pooling requirements that federated systems eliminate [29]. The privacy-utility balance favors federated learning when data sharing is prohibited [16]. Potential accuracy variations are addressed conceptually through robust averaging techniques [9]. The comparison underscores that the gap does not undermine the framework's viability for real-world deployment [17]. Instead, it positions federated training as a compliant alternative with comparable conceptual strengths [10].

Advantages of federated approach

The federated approach offers distinct advantages in privacy compliance by eliminating the need for any raw data centralization across institutions [12]. It provides

access to diverse patient populations from multiple hospitals, enriching the model's exposure to varied AKI risk profiles without ethical violations [23]. Scalability emerges as a key benefit, allowing additional sites to join the federation dynamically as resources permit [17]. These advantages align directly with the operational realities of modern healthcare networks [19]. The framework thus promotes inclusive and sustainable AI development for predictive tasks [18].

Compared to centralized methods, the federated paradigm enhances scalability by distributing computational loads across clients [24]. Privacy compliance is inherently built-in, facilitating partnerships that would otherwise be infeasible [11]. Access to broader populations supports more representative AKI prediction models [1]. The advantages collectively outweigh isolated limitations in heterogeneous settings [28]. This makes the federated strategy preferable for advancing collaborative healthcare intelligence [13].

Table 1 analytically contrasts federated and centralized learning paradigms, emphasizing the trade-offs and strategic advantages relevant to AKI prediction in regulated healthcare environments.

Table 1. Analytical Comparison of Federated and Centralized Learning Paradigms in AKI Prediction Context

Dimension	Federated Learning Framework	Centralized Learning Approach	Conceptual Implications for AKI Prediction
Data Governance	Data remains local at hospitals	Data pooled into central repository	Federated enables compliance with HIPAA/GDPR
Privacy Risk	Minimal (no raw data sharing)	High (risk of breaches and re-identification)	Federated aligns with clinical ethics
Model Generalizability	High (multi-institutional diversity)	Moderate (depends on dataset scope)	Federated improves cross-population AKI prediction

Communication Cost	Higher (iterative updates required)	Lower (single training pipeline)	Trade-off for decentralized collaborative
Scalability	High (new hospitals can join dynamically)	Limited (requires data integration pipelines)	Federated supports expanding healthcare networks
Handling Heterogeneity	Challenging (non-IID data across sites)	Easier (data homogenized centrally)	Key research area in federated AI modeling
Regulatory Compliance	Built-in by design	Difficult to maintain across jurisdictions	Federated enables global collaborative
Clinical Deployment	Flexible, institution-level integration	Centralized system dependency	Federated fits real-world hospital workflows

Challenges and Open Problems

Statistical heterogeneity

Statistical heterogeneity arises in federated systems when hospitals exhibit non-IID data distributions due to differing patient demographics and care protocols [18]. This challenge manifests as variations in AKI incidence and risk factor prevalence across sites, complicating global model convergence [17]. Different populations introduce shifts in feature distributions that must be accommodated without data exchange [16]. The framework conceptualizes mitigation through adaptive aggregation that accounts for such imbalances [24]. Addressing heterogeneity remains an open problem requiring ongoing refinement in privacy-preserving designs [9].

Even advanced centralized prediction approaches for future AKI events encounter scalability limits that heterogeneity exacerbates in federated contexts [29]. Non-IID characteristics from diverse hospital cohorts demand specialized handling to maintain predictive consistency [14]. The problem is particularly pronounced in multi-institutional AKI scenarios where case mixes vary widely

[2]. Open questions persist regarding optimal strategies for drift detection and correction [15]. The framework highlights these as priority areas for conceptual advancement [23].

System heterogeneity

System heterogeneity stems from differences in hospital hardware capabilities, electronic health record platforms, and local computational resources [26]. These variations affect the efficiency of local training for AKI-related features across the federation [25]. Incompatible EHR systems necessitate careful standardization protocols to ensure seamless integration [27]. The framework anticipates such disparities by incorporating flexible client-side adaptations [22]. System-level mismatches represent a persistent open problem in scaling federated healthcare applications [24].

Network reliability further compounds system heterogeneity when connectivity fluctuates among participating institutions [16]. Hardware differences may influence the depth of local optimization possible for prediction tasks [10]. The challenge calls for designs that accommodate a wide spectrum of infrastructure maturity [14]. Open problems include developing universal compatibility layers without sacrificing privacy [9]. The framework positions these issues as solvable through principled engineering [15].

Communication constraints

Communication constraints involve limited bandwidth and variable latency that can slow the exchange of model updates in federated rounds [22]. Dropout of participating hospitals during training phases disrupts synchronization and requires robust handling mechanisms [25]. Asynchronous versus synchronous protocols present trade-offs in convergence speed and consistency [26]. The framework conceptualizes protocols that tolerate these constraints while preserving training integrity [24]. Bandwidth limitations remain a core open problem for large-scale healthcare federations [27].

Timing of communication rounds must align with clinical operations to avoid interference [14]. Strategies for managing partial participation help maintain progress toward the global model [16]. The constraints highlight the need for efficient update compression techniques [22]. Open challenges focus on balancing communication costs with model quality [9]. The framework advances conceptual solutions tailored to these realities in AKI prediction [10].

Evaluation Strategy

Metrics for federated models

Metrics for federated models emphasize global performance across the entire federation alongside local accuracy at individual hospital sites [24]. Fairness across sites is assessed through measures that ensure equitable predictive quality regardless of local data characteristics [18]. Global metrics capture the model's overall utility for AKI risk stratification in diverse settings [17]. Local evaluations validate that each client benefits from the collaborative process [25]. This dual focus supports comprehensive assessment of the framework's effectiveness [23].

Fairness considerations prevent any single institution from experiencing degraded performance due to federation dynamics [16]. Metrics are designed to reflect clinical priorities such as early detection sensitivity [1]. The strategy avoids reliance on centralized benchmarks to maintain privacy alignment [14]. Overall, these metrics provide a conceptual lens for validating decentralized systems [15]. They ensure the framework delivers balanced value across all participants [9].

Validation protocols

Validation protocols utilize held-out hospitals as independent test sites to simulate real-world generalization without data leakage [26]. Simulated federation environments allow controlled experimentation with varying numbers of clients and heterogeneity levels [27]. Cross-silo validation divides institutions into logical groups to assess robustness under partial participation [24]. These protocols maintain strict separation of training and evaluation data flows [25]. The approach enables thorough conceptual testing of the framework's stability [22].

Held-out configurations mirror prospective deployment scenarios for AKI prediction tools [2]. Simulated setups facilitate exploration of edge cases in communication and heterogeneity [14]. Cross-silo methods promote confidence in multi-institutional scalability [16]. Validation remains privacy-preserving by design throughout the protocols [10]. The strategy establishes a rigorous foundation for future framework refinements [9].

Conclusion

The proposed conceptual framework synthesizes federated learning principles into a cohesive architecture for acute kidney injury prediction across distributed healthcare environments. It integrates local training, secure aggregation, and privacy mechanisms to enable collaborative model development. The design addresses the full spectrum of technical and regulatory considerations inherent to multi-hospital settings. This synthesis provides a blueprint for implementing decentralized AI solutions in clinical prediction tasks.

Key advantages over centralized training include superior privacy compliance and the ability to harness diverse populations without data transfer. The framework achieves scalability and ethical alignment that traditional approaches cannot match. It facilitates broader access to multi-institutional insights while upholding data sovereignty. These strengths position federated learning as a transformative paradigm for healthcare AI.

Limitations and open challenges encompass statistical and system heterogeneity that require continued conceptual innovation. Communication constraints and privacy-utility trade-offs demand ongoing refinement in practical deployments. The framework acknowledges these as areas for future methodological advancement. Addressing them will further strengthen the viability of federated systems.

Widespread implementation and validation of this framework are encouraged to realize its potential in improving AKI outcomes through privacy-preserving collaboration. Stakeholders in healthcare AI should pursue pilot integrations guided by the outlined principles. Such efforts will accelerate the adoption of federated methodologies in digital health. The conceptual foundation laid here invites interdisciplinary contributions to advance equitable and secure predictive modeling.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 17 Apr 2021 Revised: 20 Jun 2021 Accepted: 22 Jul 2021
Published online: 20 January 2022

Rights and permissions

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Dong J, Feng T, Thapa-Chhetry B, Cho BG, Shum T, Inwald DP, et al. Machine learning model for early prediction of acute kidney injury (AKI) in pediatric critical care. *Crit Care*. 2021;25(1):288.
<https://doi.org/10.1186/s13054-021-03690-4>.
- Yue S, Li S, Huang X, Liu J, Hou X, Zhao Y, et al. Machine learning for the prediction of acute kidney injury in patients with sepsis. *J Transl Med*. 2022;20(1):215.
<https://doi.org/10.1186/s12967-022-03405-0>.
- Kate RJ, Pearce N, Mazumdar D, Nilakantan V. A continual prediction model for inpatient acute kidney injury. *Comput Biol Med*. 2020;116:103580.
<https://doi.org/10.1016/j.combiomed.2019.103580>.
- Tseng PY, Chen YT, Wang CH, Chiu KM, Peng YS, Hsu SP, et al. Prediction of the development of acute kidney injury following cardiac surgery by machine learning. *Crit Care*. 2020;24(1):478.
<https://doi.org/10.1186/s13054-020-03179-9>.
- Xue B, Li D, Lu C, King CR, Wildes T, Avidan MS, et al. Use of machine learning to develop and evaluate models using preoperative and intraoperative data to identify risks of postoperative complications. *JAMA Netw Open*. 2021;4(3):e212240.
<https://doi.org/10.1001/jamanetworkopen.2021.2240>.
- Kuno T, Mikami T, Sahashi Y, Numasawa Y, Suzuki M, Noma S, et al. Machine learning prediction model of acute kidney injury after percutaneous coronary intervention. *Sci Rep*. 2022;12(1):749.
<https://doi.org/10.1038/s41598-021-04657-5>.
- Zhang L, Wang Z, Zhou Z, Li S, Huang T, Yin H, et al. Developing an ensemble machine learning model for early prediction of sepsis-associated acute kidney injury. *iScience*. 2022;25(9):105095.
<https://doi.org/10.1016/j.isci.2022.105095>.
- Song X, Liu X, Liu F, Wang C. Comparison of machine learning and logistic regression models in predicting acute kidney injury: a systematic review and meta-analysis. *Int J Med Inform*. 2021;151:104484.
<https://doi.org/10.1016/j.ijmedinf.2021.104484>.
- Xu J, Glicksberg BS, Su C, Walker P, Bian J, Wang F. Federated learning for healthcare informatics. *J Healthc Inform Res*. 2021;5(1):1-19.
<https://doi.org/10.1007/s41666-020-00082-4>.
- Sadilek A, Liu L, Nguyen D, Kamruzzaman M, Serghiou S, Rader B, et al. Privacy-first health research with federated learning. *npj Digit Med*. 2021;4(1):132.
<https://doi.org/10.1038/s41746-021-00489-2>.
- Topaloglu MY, Morrell EM, Rajendran S, Topaloglu U. In the pursuit of privacy: the promises and predicaments of federated learning in healthcare. *Front Artif Intell*. 2021;4:746497.
<https://doi.org/10.3389/frai.2021.746497>.
- Adnan M, Kalra S, Cresswell JC, Taylor GW, Tizhoosh HR. Federated learning and differential privacy for medical image analysis. *Sci Rep*. 2022;12(1):1953.
<https://doi.org/10.1038/s41598-022-05858-4>.
- Späth J, Matschinske J, Kamanu FK, Murphy SA, Zolotareva O, Bakhtiari M, et al. Privacy-aware multi-institutional time-to-event studies. *PLOS Digit Health*. 2022;1(9):e0000101.
<https://doi.org/10.1371/journal.pdig.0000101>.
- Rieke N, Hancox J, Li W, Milletari F, Roth HR, Albarqouni S, et al. The future of digital health with federated learning. *npj Digit Med*. 2020;3(1):119.
<https://doi.org/10.1038/s41746-020-00323-1>.

Zhang M, Qu L, Singh P, Kalpathy-Cramer J, Rubin DL. SplitAVG: a heterogeneity-aware federated deep learning method for medical imaging. *IEEE J Biomed Health Inform.* 2022;26(9):4635-44.
<https://doi.org/10.1109/JBHI.2022.3187566>.

Sarma KV, Harmon S, Sanford T, Roth HR, Xu Z, Tetreault J, et al. Federated learning improves site performance in multicenter deep learning without data sharing. *J Am Med Inform Assoc.* 2021;28(6):1259-64.

Dayan I, Roth HR, Zhong A, Harouni A, Gentili A, Abidin AZ, et al. Federated learning for predicting clinical outcomes in patients with COVID-19. *Nat Med.* 2021;27(10):1735-43.
<https://doi.org/10.1038/s41591-021-01506-3>.

Qu L, Balachandar N, Zhang M, Rubin D. Handling data heterogeneity with generative replay in collaborative learning for medical imaging. *Med Image Anal.* 2022;78:102424.
<https://doi.org/10.1016/j.media.2022.102424>.

Darzidehkalani E, Ghasemi-Rad M, Van Ooijen PM. Federated learning in medical imaging: part II: methods, challenges, and considerations. *J Am Coll Radiol.* 2022;19(8):975-82.
<https://doi.org/10.1016/j.jacr.2022.03.019>.

Hanif A, Lu C, Chang K, Singh P, Coyner AS, Brown JM, et al. Federated learning for multicenter collaboration in ophthalmology: implications for clinical diagnosis and disease epidemiology. *Ophthalmol Retina.* 2022;6(8):650-6.
<https://doi.org/10.1016/j.oret.2022.02.010>.

Ghazal TM, Hasan MK, Alshurideh MT, Alzoubi HM, Ahmad M, Akbar SS, et al. IoT for smart cities: machine learning approaches in smart healthcare—a review. *Future Internet.* 2021;13(8):218.
<https://doi.org/10.3390/fi13080218>.

Wu C, Wu F, Lyu L, Huang Y, Xie X. Communication-efficient federated learning via knowledge distillation. *Nat Commun.* 2022;13(1):2032.
<https://doi.org/10.1038/s41467-022-29763-x>.

Pati S, Baid U, Edwards B, Sheller M, Wang SH, Reina GA, et al. Federated learning enables big data for rare cancer boundary detection. *Nat Commun.* 2022;13(1):7346.
<https://doi.org/10.1038/s41467-022-33407-5>.

Zhou S, Landman BA, Huo Y, Gokhale A. Communication-efficient federated learning for multi-institutional medical image classification. In: *Med Imaging 2022: Imaging Informatics for Healthcare, Research, and Applications*. Bellingham (WA): SPIE; 2022. p. 6-12.
<https://doi.org/10.1117/12.2612045>.

Luo J, Wu S. FedSLD: federated learning with shared label distribution for medical image classification. In: *2022 IEEE 19th Int Symp Biomed Imaging (ISBI)*. Piscataway (NJ): IEEE; 2022. p. 1-5.
<https://doi.org/10.1109/ISBI52829.2022.9761523>.

Ma J, Zhang Q, Lou J, Xiong L, Bhavani S, Ho JC. Communication efficient tensor factorization for decentralized healthcare networks. In: *2021 IEEE Int Conf Data Mining (ICDM)*. Piscataway (NJ): IEEE; 2021. p. 1216-1221.
<https://doi.org/10.1109/ICDM51629.2021.00150>.

Ma J, Zhang Q, Lou J, Xiong L, Ho JC. Communication efficient federated generalized tensor factorization for collaborative health data analytics. In: *Proc Web Conf 2021*. New York (NY): ACM; 2021. p. 171-82.
<https://doi.org/10.1145/3442381.3449900>.

Zhang W, Zhou T, Lu Q, Wang X, Zhu C, Sun H, et al. Dynamic-fusion-based federated learning for COVID-19 detection. *IEEE Internet Things J.* 2021;8(21):15884-91.
<https://doi.org/10.1109/JIOT.2021.3056185>.

Tomašev N, Glorot X, Rae JW, Zielinski M, Askham H, Saraiva A, et al. A clinically applicable approach to continuous prediction of future acute kidney injury. *Nature.* 2019;572(7767):116-9.
<https://doi.org/10.1038/s41586-019-1390-1>.