

ORIGINAL RESEARCH

Open access

Privacy-Preserving Artificial Intelligence Platform for Cross-Hospital Learning of Patient Access Patterns Using Secure Aggregation of Scheduling Demand, Referral Completion, Appointment Lead Time, and No-Show Trends

Hassan Rahman^{1*}, Tariq Mahmood¹, Ali Raza²

Abstract

Patient access performance—scheduling efficiency, referral completion, wait times, and appointment attendance—varies widely across healthcare organizations. These organizations rarely learn from each other because operational data are sensitive, locally governed, and often competitively protected. Isolated access analytics limit the discovery of generalizable patterns and prevent hospitals from learning from peer institutions with different patient populations and workflows. No current operational platform fully enables multi-hospital learning about patient access without exposing patient-level scheduling, referral, and attendance data. This article proposes a privacy-preserving AI platform that uses federated learning and secure aggregation to support cross-hospital modeling of scheduling demand, referral completion, appointment lead time, and no-show risk. Raw data remain within each participating hospital, while only protected model updates or aggregate statistics contribute to shared learning. The platform consists of local data adapters, standardized access-feature pipelines, a federated model trainer, a secure aggregation layer, differential privacy controls, and local operational dashboards. Each hospital receives a shared model that can be adapted locally while preserving institutional data control. The framework could enable hospitals to benefit from broader operational learning while maintaining confidentiality, competitive neutrality, and governance accountability. It would be expected to support more consistent access analytics across heterogeneous health systems without requiring centralized pooling of sensitive records. Privacy-preserving AI could support a new collaborative analytics paradigm for patient access and healthcare operations. Such platforms should be evaluated through multi-institutional pilots that assess technical feasibility, governance readiness, privacy protection, and operational usefulness.

Keywords Federated learning, Healthcare operations, Privacy-preserving AI, Secure aggregation, Patient access, Appointment scheduling

*Correspondence:

Hassan Rahman

hassan.rahman@gmail.com

¹ Department of Health Information Systems and Analytics, Faculty of Medicine, University of Karachi, Karachi, Pakistan

² Department of Clinical Informatics Engineering, Faculty of Engineering, National University of Sciences and Technology, Islamabad, Pakistan

Introduction

Patient access is a core operational function because scheduling demand, appointment lead time, referral completion, and attendance behavior shape whether patients receive timely and coordinated care. Access failures may lead to

unused capacity, delayed diagnosis, fragmented referral pathways, and inequitable service availability across populations. Prior work on appointment no-shows, referral leakage, electronic consultation, and outpatient scheduling analytics shows that access processes are not merely administrative tasks but measurable determinants of system responsiveness and continuity of care [1-4].

Most hospitals still analyze patient access using local dashboards, local prediction models, and locally defined operational metrics. This approach allows each institution to tailor models to its own scheduling systems and patient population, but it also restricts learning to the patterns available inside one organization. Predictive modeling studies for outpatient no-shows and appointment attendance demonstrate the practical value of machine learning for access operations, yet they also illustrate how models are often developed within bounded institutional settings rather than across privacy-preserving networks [5-9].

Cross-hospital access learning is constrained by privacy obligations, competitive concerns, fragmented data standards, and inconsistent operational definitions. Appointment logs, referral records, lead-time histories, and attendance outcomes can reveal sensitive patient information as well as strategically important service-line performance. Federated learning, secure healthcare analytics, and interoperability work suggest that collaborative modeling can be technically feasible, but standardization and governance remain critical barriers when institutions seek to learn jointly without sharing raw data [10-16].

This article proposes an AI Systems/Frameworks platform for privacy-preserving cross-hospital learning of patient access patterns. The platform would allow hospitals to keep raw operational data inside local environments while contributing encrypted model updates, securely aggregated gradients, or differentially private statistical signals to a shared learning process. The design draws on federated healthcare AI, privacy-preserving data synthesis, common data models, and operational access analytics to frame a system that could support scheduling demand prediction, referral completion modeling, appointment lead-time analysis, and no-show risk estimation without centralized data pooling.

Background

Patient access metrics as operational drivers

Scheduling demand, referral completion, appointment lead time, cancellation behavior, and no-show risk are operational signals that describe how effectively a health system converts clinical need into completed care. No-show research shows that missed appointments can be modeled using demographic, historical, appointment-type, and service-context features, while referral studies show that incomplete referral loops can weaken care coordination and obscure unmet need. In an access-focused AI framework, these metrics would function as connected indicators of clinic capacity, care continuity, and patient engagement rather than isolated administrative outcomes [1-9].

Inter-hospital variation and the need for collaborative learning

Patient access outcomes vary across hospitals because sites differ in patient demographics, payer mix, geography, specialty availability, referral networks, scheduling rules, and care delivery norms. A model trained only on one institution's experience may miss patterns that are visible across a broader network, particularly when specific access barriers occur unevenly across populations or service lines. Collaborative learning could allow participating hospitals to benefit from shared statistical structure while still preserving the local adaptations needed for different workflows and patient communities [2, 3, 7, 17, 18].

Privacy-preserving machine learning and federated analytics

Privacy-preserving machine learning provides a technical foundation for collaborative analytics without centralized sharing of raw records. Federated learning trains models across distributed sites by sending model parameters or

updates rather than source data, while secure aggregation and secure multi-party computation can prevent the coordinator from inspecting individual institutional contributions. Differential privacy can further reduce the risk that model updates reveal patient-level information, although the privacy-utility balance must be governed carefully for operational use [10-13, 19-21].

Secure multi-institutional collaboration in healthcare

Prior healthcare federated learning work has focused heavily on clinical prediction, medical imaging, and disease-specific applications, demonstrating that multi-institutional learning can be organized without moving patient data into a central repository. These studies provide design lessons about site heterogeneity, governance, workflow integration, and distributed model validation that are relevant to healthcare operations. The gap addressed here is the translation of those privacy-preserving methods from clinical prediction settings to patient access analytics, where the protected data include scheduling demand, referral completion, lead time, and attendance histories [12, 19, 22-27].

Standardization of operational data

Cross-hospital patient access analytics require shared definitions before federated learning can produce interpretable models. Concepts such as “referral completion,” “appointment lead time,” “cancellation,” “no-show,” and “scheduling demand” may be represented differently across EHRs, scheduling systems, and referral management platforms. Interoperability standards, FHIR-based exchange concepts, and common data model approaches provide a foundation for mapping local operational data into a comparable structure while allowing each hospital to retain control of its source systems [12-15].

Platform Overview

High-level architecture

The proposed platform would use a distributed architecture in which each hospital installs a local gateway that connects to scheduling, referral, and EHR-derived operational data sources. The gateway would transform local access data into standardized features, participate in federated model training, and receive an updated shared model for local inference and dashboard display. This architecture follows the privacy-preserving logic of federated healthcare AI: computation occurs close to the data, and collaborative learning occurs through protected model updates rather than raw record exchange [10-14].

Figure 1 illustrates the proposed privacy-preserving federated architecture through which hospitals retain local custody of patient access data while contributing protected model updates to secure cross-hospital learning.




Figure 1. Privacy-Preserving Federated AI Architecture for Cross-Hospital Learning of Patient Access Patterns

Core input metrics and outputs

The platform’s conceptual inputs include historical scheduling demand, referral orders and completion statuses, appointment lead-time patterns, cancellation events, and no-show records. Its outputs would include shared predictive models for access outcomes, uncertainty-aware predictions where appropriate, and locally usable risk estimates for operational planning, referral follow-up, and scheduling support. These outputs are grounded in prior work showing that no-show behavior, referral completion, and appointment access processes can be represented as prediction and decision-support problems, while remaining careful not to claim realized performance before empirical evaluation [1-9, 17, 18, 28].

Design principles

The design principles are privacy by design, minimal data movement, model-centric collaboration, equitable participation, and robustness to heterogeneous local data distributions. Privacy by design means that identifiable scheduling and referral records remain inside each hospital, while model-centric collaboration means that the network learns from updates that are protected before aggregation. Equitable participation is important because hospitals with different patient populations and operational resources should be able to contribute to and benefit from the shared model without being required to expose proprietary access patterns [10, 11, 13, 26].

Table 1 defines the architectural layers through which the proposed platform transforms locally governed patient access data into shared, privacy-preserving operational intelligence.

Table 1. Federated Patient Access Learning Architecture: From Local Data Custody to Shared Operational Intelligence

Architectural layer	Patient access function	Local hospital responsibility	Shared network function	Privacy-preserving mechanism	Operational value added
Local access data environment	Captures scheduling demand, referral activity, appointment lead time, cancellations, and no-show outcomes	Maintains raw appointment logs, referral records, attendance histories, and local operational definitions within institutional systems	No raw data are transferred to the network	Local data custody; institutional firewall; source-system control	Preserves sensitive operational and patient-level information while enabling participation in collaborative analytics
Local data adapter	Translates heterogeneous scheduling and referral fields into a common access schema	Maps EHR, scheduling, referral-management, and access-center fields into standardized definitions	Supports comparable feature construction across hospitals	Local schema mapping; data-quality validation before participation	Reduces semantic inconsistency across sites and improves interpretability of federated model behavior

Common access-feature pipeline	Converts local records into model-ready access features	Computes features for appointment context, prior attendance, referral pathway status, lead-time patterns, clinic attributes, and workflow meta-features	Aligns feature meaning across hospitals without pooling feature-level records centrally	Local feature engineering; configurable site mappings	Enables cross-hospital learning while preserving workflow-specific relevance
Local model trainer	Learns from each hospital's standardized access features	Trains or updates the model locally using only locally governed data	Produces site-level model updates that can contribute to global learning	Federated learning; model-centric collaboration	Allows each hospital to benefit from broader learning without exporting patient-level access data
Privacy-protection layer	Protects model updates before transmission	Applies clipping, noise addition, encryption, or masking according to governance rules	Ensures that outbound signals are protected before aggregation	Differential privacy; cryptographic masking; secure transmission	Reduces risk that patient histories, referral patterns, or institutional performance signals can be inferred from updates
Secure aggregation layer	Combines protected model updates across hospitals	Submits protected updates only when participation and privacy conditions are met	Aggregates updates without exposing individual hospital contributions	Secure aggregation; secure multi-party computation; participant thresholding	Supports collaborative learning while protecting institutional confidentiality and competitive neutrality
Shared model coordination	Produces a network-informed patient access model	Receives updated shared model parameters or model packages	Learns generalizable patterns in no-show risk, referral completion, lead-time pressure, and scheduling demand	Aggregate-only model update construction; controlled model release	Creates broader operational intelligence than isolated single-site analytics alone
Local adaptation and calibration	Converts shared learning into site-specific decision support	Calibrates predictions, thresholds, and intervention logic for local workflows and patient populations	Allows global model benefits while preserving local usability	Local personalization layers; site-aware calibration	Prevents shared learning from erasing meaningful local differences in access workflows

Operational dashboard layer	Translates predictions into scheduling, referral, and access-management actions	Uses dashboards for outreach queues, referral follow-up, capacity planning, and attendance-risk mitigation	May receive anonymized or thresholded benchmark summaries	Local dashboard deployment; privacy-preserving benchmarking	Converts federated AI outputs into practical access operations without exposing raw peer data
Governance and audit layer	Oversees privacy, participation, model release, and accountable use	Maintains audit logs, validates mappings, monitors privacy-budget use, and reviews interventions	Enforces consortium rules, data-use agreements, and model-release criteria	Audit trails; role-based governance; privacy-budget tracking	Frames the platform as accountable socio-technical infrastructure rather than a purely technical model

Data Standardization and Local Feature Engineering

Common data schema for access metrics

A common access schema would define the operational meaning of scheduling demand, referral initiation, referral completion, appointment lead time, cancellation, and no-show events before model training begins. Local adapters would map EHR, scheduling, and referral-management fields into these definitions, allowing hospitals to preserve their source systems while contributing comparable features to federated learning. Common data models and interoperability standards are essential here because a privacy-preserving platform cannot correct conceptual inconsistency if participating sites label or timestamp access events in incompatible ways [14-16, 29].

Local feature extraction without data exposure

Each hospital would compute model-ready features inside its own firewall, using local appointment histories, referral pathways, patient-context variables, clinic attributes, and prior attendance patterns as permitted by governance policies. This local feature engineering approach supports privacy because sensitive patient-level records do not need to leave the site, and it supports operational relevance because features can reflect local scheduling rules and service-line practices. The approach aligns with federated learning principles and with prior appointment-access modeling work, where predictive features often derive from patient history, appointment context, and workflow characteristics [2, 5-10, 12].

Handling heterogeneity in care delivery

Hospitals differ in specialty structures, appointment types, referral workflows, reminder practices, and definitions of completion or nonattendance. The platform would therefore include site and workflow meta-features, configurable mappings, and local validation checks so that the shared model can learn common patterns while allowing local differences to remain visible. Such heterogeneity is expected in federated healthcare systems and is especially important in access analytics, where operational meaning depends on local care delivery arrangements as well as patient characteristics [3, 17-19, 23-27].

Federated Learning and Secure Aggregation Module

Federated model training process

In the federated training process, participating hospitals would initialize or receive a shared model, update it locally using their standardized access features, and transmit only protected model updates to the coordination layer. The central coordinator would not receive appointment logs, referral records, lead-time tables, or patient-level attendance histories. This process adapts the core logic of federated healthcare learning to operational access tasks, allowing cross-hospital model development while preserving local data custody [10, 12, 13, 22, 25].

Secure aggregation protocol

Secure aggregation would protect institutional contributions by ensuring that the coordinator can combine model updates without inspecting any single hospital's update in isolation. Depending on the governance and technical environment, the protocol could use secure multi-party computation, cryptographic masking, homomorphic aggregation, or related methods to produce only the aggregate update needed for model improvement. This layer is important because model updates may still encode sensitive operational or patient-level information if transmitted without additional protection, particularly in small or specialized participating networks [11, 13, 19-21].

Handling non-IID data and model heterogeneity

Patient access data are likely to be non-identically distributed across hospitals because local populations, payer structures, specialty availability, scheduling policies, and referral networks differ. The platform would therefore support methods such as proximal regularization, site-aware calibration, local personalization layers, or hybrid global-local models so that shared learning does not erase meaningful local variation. Prior federated healthcare studies and reviews emphasize that heterogeneous data distributions are a central implementation challenge, and that operational systems should be designed to preserve local usefulness rather than forcing one uniform model behavior across all sites [19, 22-27].

Privacy Mechanisms and Differential Privacy Guarantees

Differential privacy implementation

The platform would include differential privacy controls that add calibrated noise to model updates or statistical summaries before they are transmitted for aggregation. This mechanism would be intended to reduce the likelihood that an adversary could infer whether a specific patient, appointment history, or referral event contributed to the training process. Because excessive privacy noise can weaken operational usefulness, the system should be evaluated through a privacy-utility framework that considers whether access predictions remain meaningful for local scheduling, referral follow-up, and attendance-risk workflows [10, 11, 13, 20, 21].

Threat model and residual risk analysis

The platform's threat model would primarily address honest-but-curious coordinators, external attackers, and unintended disclosure through model updates or aggregate statistics. Secure aggregation would limit visibility into individual hospital contributions, while local processing would reduce exposure of patient-level scheduling and referral data. Residual risks would remain if participant pools were very small, if institutions colluded, or if governance procedures failed to constrain repeated queries and model-release practices, so the platform should combine technical protection with institutional oversight [11, 13, 19, 26].

Auditability and governance

Auditability would be implemented through logs that document local data transformations, schema mappings, privacy-preserving computation steps, aggregation events, model-release decisions, and privacy-budget use. These logs would support accountability by allowing hospitals, consortium stewards, or external reviewers to verify that the platform is operating according to agreed rules. Governance is especially important because federated learning does not remove the

need for data-use agreements, role definitions, interoperability controls, or procedures for managing disputes about model behavior and network participation [14-16, 26, 27, 29].

Shared Model for Patient Access Enhancement

Predicting no-shows, referral leakage, and lead times

The shared model would generate locally applicable predictions for access-related outcomes such as appointment nonattendance, delayed scheduling, referral noncompletion, and potential referral leakage. These predictions could help operations teams identify patients who may need additional outreach, clinics that may require scheduling adjustments, and referral pathways where follow-through may be at risk. The framework is grounded in prior access analytics showing that no-show behavior, outpatient scheduling patterns, referral completion, and referral leakage can be modeled as operational decision-support problems, although any claim of superiority over local-only models should be reserved for formal evaluation [1-9, 17, 18, 28].

Local adaptation and continuous improvement

Local adaptation would allow each hospital to adjust the shared model to its own workflows, service lines, patient populations, and scheduling practices between global learning rounds. This could be implemented through local calibration, personalization layers, or site-specific decision thresholds that preserve the benefits of shared learning while maintaining local relevance. Such adaptation is important because federated healthcare systems must handle heterogeneous data distributions, and patient access models are especially sensitive to local operational routines and referral network structures [3, 17-19, 22-27].

Operational Dashboards and Localized Insights

Site-specific access dashboards

Each hospital would receive dashboard outputs that translate model predictions into locally interpretable views for scheduling managers, referral coordinators, access-center leaders, and clinic administrators. These dashboards could surface trends in appointment demand, lead-time pressure, referral follow-through, and no-show risk while keeping patient-level records within the local environment. The design should align with clinical informatics and operations principles by presenting predictions as decision support rather than autonomous directives, particularly because access interventions must account for patient needs, equity, and workflow feasibility [1, 2, 5-9, 14].

Privacy-preserving benchmarking

The platform could support privacy-preserving benchmarking by allowing hospitals to compare aggregate access indicators or model-derived operational summaries without revealing raw appointment logs, referral records, or institution-specific patient cohorts. Benchmarks would need to be anonymized, thresholded, and governed so that hospitals cannot infer a competitor's underlying volumes, service-line weaknesses, or patient mix from shared reports. This approach extends federated analytics from model training into operational learning, where institutions can identify improvement opportunities while maintaining confidentiality and competitive neutrality [10, 13, 20, 21, 26].

Evaluation Strategy

Predictive accuracy gains over local-only models

The platform should be evaluated by comparing shared federated models with local-only models for patient access tasks such as no-show prediction, referral completion estimation, appointment lead-time forecasting, and scheduling demand

modeling. Evaluation should use appropriate predictive measures, but results should be interpreted as evidence of whether collaborative learning offers operationally meaningful improvement rather than as proof of universal superiority. Prior no-show, referral, and federated learning studies provide the conceptual basis for such comparisons while also showing why access models must be validated within real workflows before deployment [1-9, 12, 17, 18, 22, 25, 28].

Privacy-utility trade-off

A privacy-utility evaluation should examine how stronger privacy protection affects the usefulness of the shared model for local access decisions. Differential privacy, secure aggregation, and local feature restrictions may each change the information available to the learning process, so the platform should assess whether predictions remain sufficiently actionable for scheduling and referral management under different privacy settings. This evaluation would be expected to inform governance decisions about acceptable privacy budgets, minimum participant requirements, and the degree of model detail that can be safely returned to participating hospitals [10, 11, 13, 20, 21, 26].

Network scalability and participation metrics

Scalability evaluation should examine whether hospitals with different technical infrastructures, data standards, and operational maturity can participate reliably in the federated network. Relevant assessment areas include communication burden, training coordination, failure recovery, local gateway maintainability, and the ability of the common schema to support diverse scheduling and referral workflows. These considerations are central because a privacy-preserving access platform would only be useful if it can sustain multi-site participation while preserving interoperability, governance, and local control [14-16, 19, 23-27, 29].

Table 2 presents an evaluation and governance framework for determining whether privacy-preserving cross-hospital access learning is technically valid, operationally useful, equitable, and institutionally trustworthy.

Table 2. Evaluation and Governance Framework for Privacy-Preserving Cross-Hospital Patient Access AI

Evaluation domain	Core question	Suggested assessment approach	Key indicators	Governance interpretation	Risk if not evaluated
Predictive usefulness	Does federated learning improve access prediction compared with local-only modeling?	Compare federated, local-only, and hybrid global-local models across no-show prediction, referral completion, lead-time forecasting, and scheduling demand modeling	Discrimination, calibration, prediction error, decision-curve utility, subgroup performance, service-line performance	Shared learning should be retained only if it produces operationally meaningful gains without degrading local performance	Hospitals may adopt a technically sophisticated model that offers little advantage over existing local analytics
Local calibration	Are predictions usable within each hospital's own workflow and patient population?	Evaluate site-specific calibration, threshold performance, and workflow fit after local adaptation	Calibration slope, calibration-in-the-large, threshold sensitivity, alert burden, intervention yield	Local personalization should be treated as a required deployment step, not an optional refinement	A global model may appear accurate on average while producing poorly targeted interventions at individual hospitals

Privacy-utility balance	How much privacy protection can be added before operational usefulness declines?	Test differential privacy settings, update clipping, aggregation thresholds, and feature restrictions across model tasks	Privacy budget, performance degradation, uncertainty widening, subgroup stability, operational actionability	Privacy settings should be selected through explicit governance rather than technical convenience alone	Excessive privacy noise may make predictions unusable, while weak privacy settings may create unacceptable disclosure risk
Secure aggregation robustness	Can the platform protect institutional contributions during model coordination?	Simulate participant dropout, small-network conditions, malicious or honest-but-curious coordinator scenarios, and repeated aggregation queries	Minimum participant thresholds, aggregation failure rates, update exposure risk, protocol resilience	Secure aggregation should include rules for minimum participation, failure recovery, and release suppression	Individual hospital updates or sensitive access patterns may become inferable under weak participation conditions
Data standardization validity	Are access concepts defined consistently enough for cross-hospital learning?	Conduct schema-mapping audits, timestamp validation, event-definition testing, and cross-site concept review	Mapping error rates, missingness, definition concordance, no-show/referral-completion consistency, lead-time calculation agreement	Governance must treat operational definitions as model infrastructure, not clerical metadata	Federated learning may amplify inconsistent definitions rather than discovering meaningful shared patterns
Equity and access fairness	Does the model support equitable patient access across populations and sites?	Evaluate subgroup calibration, intervention allocation, referral follow-up rates, and no-show outreach patterns	Performance by demographic, social-risk, payer, geography, language, specialty, and site groups	Equity monitoring should guide threshold setting, outreach design, and model-release decisions	The platform may reproduce or conceal unequal access patterns under the appearance of technical neutrality
Operational impact	Do model outputs improve scheduling, referral management, and attendance workflows?	Pilot dashboards in access-center, referral-coordination, and clinic-management workflows using pre/post or stepped implementation designs	Wait-time change, referral-loop closure, unused appointment capacity, outreach completion, scheduling backlog, staff workload	Deployment should be justified by measurable workflow improvement, not prediction metrics alone	Accurate models may fail to improve access if outputs do not fit operational decision points
Competitive neutrality	Can hospitals learn from peer patterns without exposing	Review benchmark design, anonymization rules, thresholding,	Re-identification risk, benchmark granularity, minimum group	Benchmarking should be governed as a controlled	Shared analytics may reveal service-line weaknesses, referral leakage,

	strategically sensitive performance information?	peer-group construction, and report-release policies	size, service-line masking, report suppression events	consortium function rather than an unrestricted reporting layer	capacity constraints, or market-sensitive trends
Auditability and accountability	Can participating institutions verify how the system used data, updates, privacy budgets, and model releases?	Maintain audit logs for schema mappings, local training events, aggregation rounds, privacy-budget consumption, model release, and dashboard deployment	Audit completeness, traceability, policy exceptions, release approvals, access logs	Accountability requires auditable technical and governance processes across the full platform lifecycle	Participants may lose trust if model behavior, data transformations, or privacy decisions cannot be reconstructed
Sustainability and scalability	Can diverse hospitals participate reliably over time?	Assess gateway maintainability, communication burden, infrastructure requirements, onboarding time, and failure recovery across sites	Training-round completion, site dropout, maintenance burden, interoperability defects, implementation cost, staff support needs	The platform should scale only when technical participation requirements are realistic for heterogeneous hospitals	The network may become biased toward technically mature hospitals and exclude smaller or resource-constrained organizations

Limitations

Standardization complexity and data quality

A major limitation is that operational access data are often inconsistently defined, incompletely captured, and embedded in local scheduling or referral workflows that differ across organizations. Even with a common schema, hospitals may vary in how they record cancellations, no-shows, referral closure, appointment urgency, and lead-time start points. These differences could reduce model consistency unless the consortium invests in shared definitions, mapping validation, data-quality monitoring, and ongoing governance of operational terminology [1, 4, 14-16, 29].

Legal and competitive frameworks

Technical privacy protections do not eliminate the need for legal agreements, institutional trust, and clear rules for competitive neutrality. Hospitals may remain concerned that aggregate benchmarks, model behavior, or participation patterns could reveal strategically sensitive information about service-line capacity, referral leakage, or access bottlenecks. The platform should therefore be treated as a socio-technical infrastructure that requires data-use agreements, antitrust-aware governance, audit rights, and transparent procedures for admitting, suspending, or removing participating organizations [10, 13, 18, 26, 28].

Conclusion

The proposed platform frames patient access as a cross-hospital learning problem that can be addressed without centralizing raw operational data. By combining local data processing, standardized access features, federated learning, secure aggregation, and privacy controls, it could support collaborative modeling of scheduling demand, referral completion, appointment lead time, and no-show risk.

A key strength of the framework is that it separates learning from data pooling. Hospitals could contribute to shared operational intelligence while retaining local custody of appointment histories, referral records, and patient-level access data. This architecture would be expected to support stronger privacy protection, greater institutional trust, and more generalizable access modeling than isolated local analytics alone.

Important challenges remain before such a system could be responsibly implemented. Operational data standardization, legal governance, privacy-budget management, local workflow integration, and multi-institutional accountability would all require careful design. Real-world pilot testing should examine feasibility, usability, governance readiness, and the conditions under which shared learning provides practical value.

Future work should focus on consortium-based implementation projects in integrated delivery networks, regional health information exchanges, or trusted multi-hospital collaboratives. These projects could demonstrate how privacy-preserving AI can support patient access improvement while respecting confidentiality, competitive boundaries, and local operational autonomy. A mature version of this framework could help healthcare organizations learn together while keeping sensitive access data where they belong.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 04 Feb 2026

Revised: 01 Apr 2026

Accepted: 15 May 2026

Published online: 20 July 2026

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Dantas LF, Fleck JL, Oliveira FL, Hamacher S. No-shows in appointment scheduling: a systematic literature review. *Health Policy*. 2018;122(4):412-21.
- Mohammadi I, Wu H, Turkcan A, Toscos T, Doebbeling BN. Data analytics and modeling for appointment no-show in community health centers. *J Prim Care Community Health*. 2018;9:2150132718811692.
- Lai L, Liddy C, Keely E, Afkham A, Kurzawa J, Abdeen N, et al. The impact of electronic consultation on a Canadian tertiary care pediatric specialty referral system: a prospective single-center observational study. *PLoS One*. 2018;13(1):e0190247.
- Patel MP, Schettini P, O'Leary CP, Bosworth HB, Anderson JB, Shah KP, et al. Closing the referral loop: an analysis of primary care referrals to specialists in a large health system. *J Gen Intern Med*. 2018;33(5):715-21.
- Chua SL, Chow WL. Development of predictive scoring model for risk stratification of no-show at a public hospital specialist outpatient clinic. *Proc Singapore Healthc*. 2019;28(2):96-104.
- Chong LR, Tsai KT, Lee LL, Foo SG, Chang PC. Artificial intelligence predictive analytics in the management of outpatient MRI appointment no-shows. *AJR Am J Roentgenol*. 2020;215(5):1155-62.
- Ferro DB, Brailsford S, Bravo C, Smith H. Improving healthcare access management by predicting patient no-show behaviour. *Decis Support Syst*. 2020;138:113398.
- Fan G, Deng Z, Ye Q, Wang B. Machine learning-based prediction models for patients no-show in online outpatient appointments. *Data Sci Manag*. 2021;2:45-52.
- Hamdan AF, Bakar AA. Machine learning predictions on outpatient no-show appointments in a Malaysia major tertiary hospital. *Malays J Med Sci*. 2023;30(5):169-85.
- 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.
- Kaissis GA, Makowski MR, Rückert D, Braren RF. Secure, privacy-preserving and federated machine learning in medical imaging. *Nat Mach Intell*. 2020;2(6):305-11.
- Guo P, Wang P, Zhou J, Jiang S, Patel VM. Multi-institutional collaborations for improving deep learning-based magnetic resonance image reconstruction using federated learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021. p. 2423-32.
- 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.
- Lehne M, Sass J, Essenwanger A, Schepers J, Thun S. Why digital medicine depends on interoperability. *NPJ Digit Med*. 2019;2(1):79.
- Reinecke I, Zoch M, Reich C, Sedlmayr M, Bathelt F. The usage of OHDSI OMOP: a scoping review. In: *German Medical Data Sciences 2021: Digital Medicine: Recognize–Understand–Heal*. 2021. p. 95-103.
- Finster M, Wenzel M, Taghizadeh E. Common data models and data standards for tabular health data: a systematic review. *BMC Med Inform Decis Mak*. 2025;25(1):422.
- Nabelsi V, Lévesque-Chouinard A, Liddy C, Pilon MD. Improving the referral process, timeliness, effectiveness, and equity of access to specialist medical services through electronic consultation: pilot study. *JMIR Med Inform*. 2019;7(3):e13354.

Liddy C, de Man G, Moroz I, Afkham A, Mercer J, Keely E. Effective integration of an eConsult service into an existing referral workflow within a primary care clinic. *Telemed J E Health*. 2020;26(5):659-64.

Warnat-Herresthal S, Schultze H, Shastri KL, Manamohan S, Mukherjee S, Garg V, et al. Swarm learning for decentralized and confidential clinical machine learning. *Nature*. 2021;594(7862):265-70.

Beaulieu-Jones BK, Wu ZS, Williams C, Lee R, Bhavnani SP, Byrd JB, et al. Privacy-preserving generative deep neural networks support clinical data sharing. *Circ Cardiovasc Qual Outcomes*. 2019;12(7):e005122.

Yoon J, Drumright LN, Van Der Schaar M. Anonymization through data synthesis using generative adversarial networks (ADS-GAN). *IEEE J Biomed Health Inform*. 2020;24(8):2378-88.

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.

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.

Pati S, Kumar S, Varma A, Edwards B, Lu C, Qu L, et al. Privacy preservation for federated learning in health care. *Patterns (N Y)*. 2024;5(7):100000.

Teo ZL, Jin L, Liu N, Li S, Miao D, Zhang X, et al. Federated machine learning in healthcare: a systematic review on clinical applications and technical architecture. *Cell Rep Med*. 2024;5(2):101397.

Eden R, Chukwudi I, Bain C, Barbieri S, Callaway L, de Jersey S, et al. A scoping review of the governance of federated learning in healthcare. *NPJ Digit Med*. 2025;8(1):427.

Li M, Xu P, Hu J, Tang Z, Yang G. From challenges and pitfalls to recommendations and opportunities: implementing federated learning in healthcare. *Med Image Anal*. 2025;101:103497.

O'Connor GE, Cook LA. Reducing referral leakage: an analysis of health-care referrals in a service ecosystem. *J Serv Mark*. 2020;34(4):513-28.

Klann JG, Estiri H, Weber GM, Moal B, Avillach P, Hong C, et al. Validation of an internationally derived patient severity phenotype to support COVID-19 analytics from electronic health record data. *J Am Med Inform Assoc*. 2021;28(7):1411-20.