

ORIGINAL RESEARCH

Open access

Federated Multi-Task Learning Framework for Predicting Service Demand Across Emergency, Imaging, Pharmacy, Laboratory, and Inpatient Units without Sharing Patient-Level Operational Data

Gabriel Costa^{1*}, Rafael Mendes¹, Bruno Teixeira², Lucas Ribeiro¹

Abstract

Accurate prediction of demand for emergency, imaging, pharmacy, laboratory, and inpatient services is critical for hospital planning. However, forecasting models are typically built separately by department or institution, which limits their ability to learn from shared demand patterns. Hospitals generate rich operational demand streams, but patient-level data cannot usually be pooled across organizations. This creates a need for collaborative forecasting methods that preserve institutional control over sensitive operational records. This article proposes a federated multi-task learning framework for predicting service demand across multiple hospital units. The framework trains a shared predictive model across hospitals while each institution contributes only protected model updates. The framework includes local data adapters, a shared temporal learning backbone, task-specific forecasting heads, a federated aggregation layer, differential privacy mechanisms, and site-specific personalization modules. Together, these components support collaborative forecasting without transferring patient-level operational data. The framework could improve demand prediction by learning common temporal patterns across hospitals and service lines. It would also support local adaptation, reduce duplicated model development, and preserve data confidentiality. A privacy-preserving, collaborative approach to hospital demand forecasting could become a core infrastructure for multi-site operational coordination. Federated multi-task learning offers a practical conceptual foundation for such a system.

Keywords Federated learning, Healthcare operations, Demand forecasting, Privacy-preserving analytics, Multi-task learning, Hospital capacity

*Correspondence:

Gabriel Costa
gabriel.costa@gmail.com

¹ Department of Healthcare Information Engineering, Faculty of Medicine, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

² Department of Intelligent Clinical Analytics, Faculty of Engineering, University of Campinas, Campinas, Brazil

Introduction

Hospitals depend on reliable forecasts to plan staffing, beds, diagnostics, medication supply, laboratory capacity, and patient flow. Emergency department studies show that arrivals and occupancy contain structured temporal patterns that could support planning when modeled appropriately [1]. Similar forecasting needs arise in radiology, where examination volume affects workforce scheduling, scanner utilization, and downstream care coordination [2]. Laboratory and pharmacy demand also create operational pressure because test volumes and medication workload influence turnaround times, inventory planning, and clinical throughput [3, 4]. Inpatient bed demand adds another layer of complexity because

admission, discharge, and transfer dynamics connect upstream emergency demand with downstream capacity constraints [5].

Traditional hospital forecasting models are often developed for one department, one site, or one operational target. This siloed approach may overlook shared drivers, such as day-of-week patterns, seasonal variation, admission surges, or changes in local disease burden. Multi-task learning offers a way to represent related prediction problems jointly while still preserving task-specific outputs [6]. Because hospital service lines are operationally connected, a model that learns across emergency, imaging, pharmacy, laboratory, and inpatient tasks could capture demand relationships that isolated models may miss [7]. Such an approach would be especially valuable when individual hospitals have limited historical data for certain services.

A second challenge is that the most informative operational learning environment would span multiple hospitals, but centralized pooling of patient-level operational data is often infeasible. Federated learning addresses this barrier by allowing institutions to train models collaboratively while keeping data within local environments [8]. Prior healthcare applications have shown that federated learning can support multi-institutional model development without direct data sharing [9]. The same privacy-preserving logic has been extended to digital health and medical imaging, where institutional collaboration must be balanced against confidentiality and governance requirements [10]. These precedents suggest that a federated architecture could be adapted to hospital operations forecasting.

Federated multi-task learning combines distributed collaboration with shared representation learning. The federated component keeps operational records local, while the multi-task component allows related service-line forecasts to benefit from common temporal representations. Secure aggregation could prevent the coordinating server from inspecting individual hospital updates. Differential privacy could further reduce the risk that model updates reveal sensitive utilization patterns. This article therefore proposes a conceptual AI Systems/Frameworks architecture for predicting demand across emergency, imaging, pharmacy, laboratory, and inpatient units without exposing patient-level operational data.

Background

Service demand in hospitals

Hospital service demand is not a single stream but a set of connected operational time series. Emergency departments experience temporal variation in arrivals and occupancy, making short-horizon forecasting important for crowding management and resource allocation [11]. Imaging demand creates related planning problems because examination volume affects scanner scheduling, radiologist workload, and diagnostic throughput [12]. Laboratory test volumes require their own forecasting logic because specimen flows, test ordering patterns, and turnaround expectations vary over time [13]. Pharmacy demand and inpatient bed capacity add further interdependence, since medication workload and bed occupancy are shaped by admissions, care pathways, and discharge timing [4, 5].

Multi-task learning in healthcare

Multi-task learning is relevant to healthcare operations because many hospital prediction tasks are distinct but related. A shared representation can learn common temporal features, while task-specific layers can adapt outputs to emergency visits, imaging orders, pharmacy workload, laboratory volumes, or bed demand [6]. Dense prediction and broader multi-task learning research show that shared backbones can support related tasks when careful task-specific modeling is retained [7]. In the proposed framework, this means that common demand signals would be learned jointly, while each service line would maintain a specialized forecasting head. The result is a conceptual architecture that treats hospital demand as an interconnected operational system rather than a collection of isolated forecasting problems.

Federated learning and privacy-preserving analytics

Federated learning enables multiple institutions to collaborate on model training without transferring raw data. Early work on federated predictive modeling from electronic health records demonstrated the feasibility of learning across distributed clinical datasets while preserving data locality [8]. Medical imaging studies further showed that multi-institutional deep learning could be organized without centralizing patient data [14]. Broader reviews of federated healthcare and digital health emphasize that this approach is especially useful when privacy, regulation, and institutional governance limit data sharing [10, 15]. For hospital operations, the same principle applies because utilization data can reveal sensitive patient flows, capacity constraints, and institutional operating patterns.

Time-series forecasting for hospital operations

Hospital demand forecasting requires models that can represent temporal dependence, seasonality, and known future covariates. Emergency department studies have used time-series and machine learning methods to forecast arrivals and occupancy across operational planning horizons [16]. Radiology forecasting research shows that diagnostic service volumes can also be modeled as structured time series for resource planning [2]. Laboratory and pharmacy forecasting studies extend this logic to test demand and medication workload, where historical volumes and calendar effects may inform future demand [3, 17]. Temporal Fusion Transformers are relevant as a conceptual backbone because they were designed for interpretable multi-horizon time-series forecasting with static, historical, and future inputs [18].

Cross-institutional collaboration and barriers

Cross-hospital collaboration could improve forecasting because each institution observes different demand patterns, service configurations, and local operating environments. However, legal, technical, competitive, and governance barriers often prevent direct pooling of operational data. Federated learning has been identified as a promising strategy for digital health collaboration because it allows model development while data remain inside participating institutions [10]. Privacy-preserving machine learning methods, including secure aggregation and privacy-aware training, can further strengthen trust in such collaborations [19, 20]. These concerns are central to hospital operations because even aggregated-looking demand streams may encode sensitive information about patient flows and institutional capacity.

Framework Overview

High-level federation design

The proposed framework assumes that each hospital maintains its own data warehouse, local compute node, and service-line data connectors. A coordinating server would distribute model parameters, orchestrate training rounds, and receive only protected model updates rather than raw emergency, imaging, pharmacy, laboratory, or inpatient records. This design follows established federated learning principles in which clients train locally and share model updates for aggregation [21]. Healthcare applications of federated learning show that such an architecture can be adapted to multi-institutional settings where direct data sharing is restricted [9]. The framework therefore treats each hospital as both a local forecasting environment and a participant in a broader collaborative learning network.

Figure 1 presents the proposed federated multi-task learning architecture, showing how hospitals can jointly train service-demand forecasting models across emergency, imaging, pharmacy, laboratory, and inpatient units while retaining patient-level operational data locally.




Figure 1. Federated Multi-Task Learning Framework for Privacy-Preserving Hospital Service Demand Forecasting.

Core input features and task definitions

The framework would use a common feature schema that each hospital maps locally from its own operational systems. Inputs could include historical demand, calendar variables, admission and discharge context, recent occupancy, and locally available external signals. Emergency department forecasting studies support the use of historical volume and temporal structure for short-term planning [22]. Imaging, laboratory, pharmacy, and bed-capacity studies similarly indicate that service volumes can be represented through operational time-series features [2-5]. Each service demand stream would be defined as a regression task, allowing the model to produce task-specific forecasts from a shared temporal representation.

Design principles

The framework is organized around privacy-by-design, task modularity, non-IID robustness, and practical hospital onboarding. Privacy-by-design means that patient-level operational records never leave the local institution, and only protected updates are shared for aggregation [20]. Task modularity means that hospitals can participate even when they do not provide every service or cannot immediately standardize every data stream. Non-IID robustness is necessary because hospital demand varies by size, geography, case mix, referral role, and service configuration [23]. These design principles align with federated multi-task learning, where related clients and tasks are learned collaboratively without assuming that all local distributions are identical [24].

Federated Multi-Task Learning Design
Shared backbone and task-specific heads

The core model would contain a shared temporal backbone that encodes demand history, calendar context, and operational state across all participating hospitals. This backbone could be implemented conceptually using a recurrent architecture or a Transformer-based sequence model suitable for multi-horizon forecasting [18]. Task-specific heads would then generate separate predictions for emergency visits, imaging orders, pharmacy workload, laboratory test volumes, and inpatient bed demand. Multi-task learning research supports this separation between shared parameters and specialized outputs because related tasks can benefit from common representations while retaining task-level adaptation [6]. Site-specific personalization heads would further allow each hospital to calibrate the shared model to its local demand environment [7].

Table 1 maps each architectural layer of the proposed framework to its technical function, operational role, privacy-preserving contribution, and added analytical value.

Table 1. Architecture-to-Function Map of the Federated Multi-Task Hospital Demand Forecasting Framework

Framework layer	Core technical function	Hospital operations role	Privacy-preserving contribution	Analytical value added to the manuscript
Local hospital node	Maintains institutional data warehouse, local compute environment, and service-line connectors	Preserves each hospital as the primary site of operational forecasting and decision use	Prevents centralized pooling of patient-level records, timestamps, service volumes, and local capacity patterns	Clarifies that the framework is a distributed hospital operations system rather than a conventional centralized prediction model
Local service-line demand streams	Converts emergency, imaging, pharmacy, laboratory, and inpatient activity into forecastable time series	Represents hospital demand as a connected multi-service operational environment	Keeps raw demand streams within each institution	Shows why the model must address cross-service interdependence rather than isolated departmental forecasting
Local data adapters	Maps native systems into a common feature schema	Enables hospitals with different information systems to participate in the same modeling framework	Performs standardization locally without exporting source records	Identifies local adapter design as the bridge between institutional heterogeneity and shared learning
Temporal harmonization layer	Aligns timestamps, forecast horizons, calendar variables, and service-specific frequencies	Allows hourly, shift-based, daily, or multi-horizon planning needs to coexist	Avoids transferring raw timestamped patient events or operational sequences	Explains how multi-service forecasting becomes technically comparable across hospitals
Shared temporal learning backbone	Learns common temporal patterns across sites and service lines	Captures shared demand drivers such as seasonality, weekday effects, surges, and capacity pressure	Learns from distributed model updates instead of centralized raw data	Establishes the framework's main theoretical contribution: shared representation learning for hospital demand
Task-specific forecasting heads	Produces separate predictions for emergency, imaging, pharmacy, laboratory, and inpatient demand	Preserves the distinct operational meaning of each service-line output	Limits unnecessary sharing by allowing task-level specialization within local training	Prevents the architecture from collapsing diverse hospital services into a single generic demand target
Site-specific personalization layer	Calibrates model behavior to local case mix, hospital role, geography, and service configuration	Supports local relevance for trauma centers, community hospitals, academic centers, and specialty hospitals	Keeps institution-specific adaptation parameters local	Addresses non-IID variation as a core design requirement rather than an implementation detail

Federated aggregation layer	Combines protected model updates across participating hospitals	Allows multi-hospital learning without direct data exchange	Receives protected updates rather than patient-level data or raw operational time series	Defines the collaborative learning mechanism that distinguishes the framework from siloed local models
Secure aggregation protocol	Prevents the coordinating server from inspecting individual hospital updates	Reduces institutional concern that participation will reveal operational weaknesses or capacity stress	Aggregates updates so that no single hospital contribution is exposed	Functions as the technical trust layer needed for multi-institutional operations forecasting
Differential privacy mechanism	Adds calibrated noise to gradients or parameters before update transmission	Protects against inference of sensitive utilization or workload patterns	Reduces the risk that updates reveal patient-level or institution-level demand signals	Frames privacy as an evaluable model-design parameter rather than a general assurance claim
Communication-efficiency controls	Uses periodic updates, compression, sparsification, or selective parameter exchange	Makes participation feasible for hospitals with varied IT resources	Minimizes transmitted model information	Connects technical feasibility to real-world hospital network participation
Local dashboard and governance interface	Converts forecasts into staffing, capacity, inventory, queue, and bed-management views	Supports managers, charge nurses, diagnostic leaders, pharmacy planners, laboratory supervisors, and bed managers	Keeps operational dashboards local while sharing only protected learning signals	Translates the framework from model architecture into operational decision support

Federated training process

During each training round, participating hospitals would receive the current model and update it locally using their own operational data. The hospital would compute model updates without transmitting patient-level records to the federation. Secure aggregation could then allow the coordinating server to combine updates while preventing inspection of any individual hospital's contribution [19]. Federated healthcare studies show that this type of local-training and central-aggregation workflow can support multi-institutional collaboration without direct data sharing [8, 9]. Differential privacy could be added before transmission so that updates are less likely to reveal sensitive information about local utilization patterns [25].

Handling heterogeneous tasks and missing modalities

Hospitals differ in available services, data maturity, coding practices, and operational definitions. The framework would therefore mask unavailable service-line tasks during local training rather than requiring every hospital to contribute every output. Multi-task learning can support this kind of modular participation when the shared representation is trained from partially overlapping task information [6]. Federated time-series forecasting research also emphasizes that client heterogeneity must be handled directly rather than treated as a minor implementation issue [23]. In this framework, a hospital could contribute emergency and inpatient demand updates even if imaging, pharmacy, or laboratory feeds were not yet standardized.

Multi-Modal Demand Data Standardization

Local data adapters and common feature mapping

Each hospital would implement local adapters that transform native operational data into a common modeling schema. Emergency department data could be mapped from tracking systems, imaging demand from radiology information systems, laboratory volumes from laboratory information systems, pharmacy workload from medication systems, and bed demand from admission-discharge-transfer platforms. Prior work on emergency forecasting shows that local operational timestamps can be converted into forecastable demand series [26]. Radiology and laboratory forecasting studies similarly demonstrate that departmental information systems can support volume prediction when data are structured into time-indexed features [2, 3]. The adapter layer allows this standardization to occur locally, so raw records remain inside each hospital.

Temporal alignment and frequency harmonization

Demand streams may be recorded at different temporal resolutions across services and hospitals. Emergency departments may require hourly forecasts, while imaging, pharmacy, laboratory, and inpatient planning may use daily or shift-based horizons. Time-series forecasting methods for emergency occupancy, radiology volume, laboratory testing, pharmacy workload, and bed capacity all depend on consistent temporal indexing before modeling can occur [5, 11-13, 17]. The framework would therefore harmonize timestamps locally before training, allowing each hospital to align service demand to the common forecasting horizon. This step would support shared learning while avoiding transfer of raw timestamped patient events.

Data quality and completeness monitoring

Local nodes would perform automated checks for missing periods, delayed feeds, implausible values, and inconsistent service definitions before participating in model training. These checks are necessary because operational forecasting depends on coherent time-series construction and stable definitions of demand. Emergency and inpatient forecasting studies illustrate how sensitive planning models can be to occupancy definitions and temporal aggregation choices [5, 16, 22]. Laboratory, pharmacy, and radiology studies likewise depend on accurate representation of service volumes over time [2-4]. The framework would flag data quality issues to local administrators while sharing only training-readiness signals or protected metadata with the federation.

Privacy-Preserving Aggregation and Communication

Secure aggregation protocol

The framework would use secure aggregation so that the coordinating server receives only an aggregate update rather than any individual hospital's model contribution. Practical secure aggregation protocols are relevant because they are designed to protect client updates during collaborative learning while still permitting model improvement through averaged parameters [19]. In the proposed hospital setting, this would mean that updates from emergency, imaging, pharmacy, laboratory, and inpatient forecasting tasks could be combined without exposing the operational signature of any single institution. Secure aggregation would be especially important when hospitals differ in case mix or capacity, because even model updates could otherwise reveal information about local utilization patterns [20]. This protocol would therefore function as a technical trust layer between participating hospitals and the federation.

Differential privacy guarantees

Differential privacy would be incorporated by adding calibrated noise to local gradients or model weights before updates are transmitted. This mechanism would reduce the risk that an attacker could infer sensitive patient-level or institution-level operational information from a hospital's contribution to the federated model [25]. In a demand forecasting context,

privacy protection matters because service volumes may encode patterns related to emergency surges, imaging backlogs, medication demand, laboratory testing intensity, or bed pressure. The framework would treat privacy budgeting as a governance and technical parameter that should be evaluated alongside forecast utility rather than as an afterthought. This aligns with broader privacy-preserving healthcare machine learning, where model development must be balanced against confidentiality and institutional trust [27].

Communication efficiency and bandwidth

Communication efficiency is important because hospitals may have different computing resources, network constraints, and IT support capacity. Instead of requiring continuous data transfer, the framework would exchange periodic model updates, using compression, sparsification, or selective parameter sharing where appropriate. Federated learning research identifies communication cost and systems heterogeneity as persistent challenges, particularly when client environments vary widely [21]. In healthcare networks, these challenges may be amplified by security review processes, restricted infrastructure, and uneven analytics maturity across institutions [10]. The framework would therefore prioritize lightweight communication protocols that allow participation without imposing excessive technical burden on local hospital teams.

Local Personalization and Non-iid Robustness

Site-specific adaptation layers

Site-specific adaptation layers would allow each hospital to adjust the shared model to its own operational environment. A trauma center, community hospital, pediatric hospital, and academic referral center may experience different baseline demand, service mix, and temporal variation, even when they share broad forecasting tasks. Federated multi-task learning is well suited to this situation because it does not require every client to learn an identical model; instead, it can support related local models connected through shared structure [24]. Federated time-series forecasting research similarly emphasizes the need to handle heterogeneous local distributions rather than assuming that all sites produce interchangeable data [23]. In the proposed framework, personalization layers would remain local so that adaptation improves site relevance without exposing institution-specific operational patterns.

Handling concept drift and local seasonality

The framework would support continuing local fine-tuning between federation rounds so that each hospital can adapt to changing conditions. Concept drift may occur when service demand changes because of seasonal illness, local outbreaks, staffing changes, referral shifts, new imaging capacity, medication supply disruptions, or changes in admission policy. Emergency department and inpatient forecasting studies show that demand and occupancy are sensitive to temporal context and operational state, making static models insufficient for long-term deployment [1, 5, 16]. Temporal forecasting architectures such as Temporal Fusion Transformers are relevant because they are designed to incorporate known future variables and evolving historical context [18]. Local adaptation would therefore allow the framework to remain responsive while preserving the broader benefits of federated learning.

Operational Dashboard and Decision Support Integration

Forecast delivery to hospital operations teams

At each hospital, the trained local model would deliver multi-horizon forecasts through an operational dashboard designed for managers, charge nurses, staffing coordinators, diagnostic service leaders, pharmacy planners, laboratory supervisors, and bed managers. The dashboard would present expected demand direction, uncertainty, and service-line interactions rather than treating forecasts as isolated numerical outputs. Prior work on emergency, radiology, laboratory, pharmacy, and inpatient demand forecasting shows that prediction becomes operationally meaningful only when it is

connected to planning decisions such as staffing, equipment allocation, inventory preparation, and capacity coordination [2-5, 22]. The framework would therefore translate model outputs into decision-support views that remain local to each hospital. This design avoids exposing operational dashboards externally while still allowing the shared federated model to improve over time.

Privacy-preserving benchmarking across sites

The framework could also support privacy-preserving benchmarking by reporting only aggregated or anonymized model performance indicators across the participating network. Each hospital could compare its forecasting process with a network-level reference while avoiding disclosure of raw volumes, departmental backlogs, or identifiable operational stress points. Federated learning reviews emphasize that collaboration requires both technical privacy controls and institutional trust, especially when participants may be legally separate or strategically cautious [10, 27]. Multi-institutional medical imaging work has shown that collaborative model development can occur without direct data sharing, offering a useful precedent for operational benchmarking [28, 29]. In the proposed framework, benchmarking would be designed as a governance-sensitive feedback mechanism rather than as a competitive ranking system.

Evaluation Strategy

Forecast accuracy across tasks and sites

The framework should be evaluated using forecast accuracy measures appropriate for local held-out test sets at each participating hospital. Candidate measures could include mean absolute error, symmetric mean absolute percentage error, and quantile loss, but the article does not claim any experimental outcome or numerical performance level. Prior emergency department studies motivate evaluation across arrival and occupancy forecasting tasks, while radiology, laboratory, pharmacy, and inpatient capacity studies show that each service line requires task-specific interpretation of forecast quality [2-5, 11-13, 17, 26]. Evaluation should therefore examine whether the framework could support reliable predictions across departments and hospitals without assuming that one metric fully captures operational usefulness. Results should be reported conceptually or empirically only when future studies use real deployments and transparent validation protocols.

Federated vs. Local and centralized baselines

A future evaluation should compare the federated multi-task framework with local single-task models and, where legally and ethically possible, a centralized pooled-data reference model used only as an analytical upper-bound comparator. Local single-task baselines would show whether collaboration across tasks and hospitals adds value beyond isolated departmental modeling. Centralized baselines would help estimate the potential value of pooled information, while acknowledging that such pooling may be unrealistic under privacy and governance constraints. Federated healthcare studies and reviews provide the rationale for this comparison because they position federated learning as a compromise between isolated local learning and direct centralization [8, 9, 15, 21]. The comparison should avoid overstating superiority unless supported by transparent future experiments.

Privacy, communication, and participation metrics

Evaluation should also include non-accuracy criteria because the value of the framework depends on privacy protection, communication feasibility, and sustained hospital participation. Relevant measures would include privacy budget consumed, volume of model-update communication, successful completion of training rounds, failure handling, and the ability of hospitals with incomplete modalities to remain active in the federation. Secure aggregation work supports evaluating whether individual updates remain protected during aggregation, while federated learning surveys highlight communication efficiency and participation heterogeneity as central systems concerns [19, 21, 27]. Privacy-preserving healthcare machine learning also suggests that technical safeguards should be assessed in relation to governance

expectations and data sensitivity [20, 25]. These criteria would ensure that evaluation reflects the full AI systems framework rather than only forecasting accuracy.

Table 2 proposes a multidimensional evaluation and governance framework for assessing whether federated multi-task demand forecasting is accurate, privacy-preserving, operationally usable, and institutionally sustainable.

Table 2. Evaluation and Governance Framework for Federated Multi-Task Hospital Demand Forecasting

Evaluation domain	Key question for future pilots	Candidate measures or evidence sources	Why this domain matters for hospital operations	Governance implication
Forecast accuracy by service line	Does the framework produce useful predictions for each operational task?	Mean absolute error, symmetric mean absolute percentage error, quantile loss, calibration plots, prediction interval coverage	Emergency, imaging, pharmacy, laboratory, and inpatient forecasts have different planning consequences and cannot be judged by one generic metric	Evaluation reports should be stratified by service line rather than reported only as an aggregate model score
Cross-site generalizability	Does shared learning improve performance across hospitals with different demand patterns?	Site-level held-out validation, leave-one-site-out testing, performance dispersion across hospital types	Hospitals differ by size, referral role, geography, case mix, and operational configuration	Consortium governance should define acceptable variation and avoid penalizing hospitals with structurally different demand patterns
Federated versus local baselines	Does collaboration add value beyond isolated single-site models?	Comparison with local single-task models, local multi-task models, federated single-task models, and federated multi-task models	Demonstrates whether the added complexity of federation is operationally justified	Participation should be justified by transparent evidence that networked learning improves planning utility
Centralized upper-bound comparison	How close can federated learning come to pooled-data performance when centralization is infeasible?	Centralized pooled-data reference model where legally and ethically permissible	Helps estimate the performance cost or benefit of privacy-preserving collaboration	Centralized comparison should be framed as an analytical benchmark, not as the preferred implementation model
Non-IID robustness	Can the model handle hospitals with different distributions, missing services, or uneven task availability?	Performance by hospital type, task-masking success, partial-participation analysis, subgroup error profiles	Real hospital networks rarely contain identical institutions or complete service feeds	Governance should allow modular onboarding rather than requiring all hospitals to meet identical data maturity standards
Privacy protection	Are model updates protected against	Privacy budget reporting, differential	Demand patterns can reveal emergency	Privacy safeguards should be documented,

	inference of sensitive patient-level or institutional operational patterns?	privacy parameters, secure aggregation verification, attack-resistance testing	surges, backlogs, bed pressure, or institutional capacity limitations	audited, and reviewed as part of participation agreements
Communication and systems feasibility	Can hospitals participate without excessive infrastructure burden?	Update size, bandwidth use, training-round completion rate, dropped-client frequency, compute burden	Federated learning can fail operationally if communication requirements exceed hospital IT capacity	The consortium should define minimum technical requirements and support lightweight participation pathways
Data quality and semantic alignment	Are local demand definitions sufficiently comparable for shared learning?	Missingness checks, timestamp completeness, service definition audits, feature-mapping validation	Forecasting quality depends on stable definitions of visits, orders, tests, workload units, and bed states	A shared data dictionary and local validation protocol should be maintained before and during deployment
Dashboard usability and decision uptake	Do operations teams understand and use the forecasts in planning workflows?	User testing, dashboard interaction logs, decision-action tracking, qualitative feedback from managers and supervisors	Forecasts matter only if they change staffing, capacity, inventory, scheduling, or escalation decisions	Human oversight should be embedded in dashboard review rather than treated as an optional downstream step
Benchmarking fairness	Can hospitals compare performance without exposing raw volumes or creating punitive rankings?	Aggregated performance bands, anonymized network references, fairness review of benchmark displays	Benchmarking can support learning but may also create reputational or competitive concerns	Network reporting should emphasize improvement and shared learning rather than competitive institutional ranking
Model drift and maintenance	Does the framework remain useful as demand patterns, staffing, policies, or seasonal pressures change?	Drift detection, rolling-window validation, recalibration frequency, local fine-tuning performance	Hospital demand is dynamic and affected by outbreaks, staffing changes, service expansion, and policy shifts	The governance model should specify retraining schedules, drift response procedures, and accountability for model updates
Implementation readiness	Is the framework mature enough for real operational use?	Pilot feasibility, governance approval, privacy review, workflow integration, failure-mode testing	Deployment requires more than predictive performance; it requires safe integration into hospital operations	Future pilots should evaluate technical, organizational, legal, and workflow readiness together

Limitations

Non-standardized operational data

Even with local adapters, hospitals may define and record demand events differently. An emergency visit, imaging order, laboratory test, pharmacy workload unit, or inpatient bed occupancy state may vary across institutions because of workflow, coding, system configuration, or reporting practice. Service-specific forecasting studies depend on consistent operational definitions, and differences in temporal aggregation or volume construction could limit the comparability of model inputs across sites [1-5]. Multi-task and federated time-series learning can help manage heterogeneity, but they cannot fully eliminate ambiguity in the underlying operational concepts [6, 23]. The framework would therefore require ongoing data governance and semantic alignment in addition to technical model design.

Governance and incentive alignment

A federated hospital forecasting network would require more than algorithms; it would require governance agreements, participation incentives, accountability structures, and trust in the privacy technology. Competing health systems may hesitate to collaborate if they believe model updates or benchmarks could reveal capacity constraints, market position, or operational weaknesses. Federated learning literature emphasizes that collaboration is shaped by privacy, communication, incentives, and institutional context, not only by model architecture [10, 21, 27]. Secure aggregation and privacy-preserving machine learning can reduce technical risk, but they do not automatically resolve legal, organizational, or strategic concerns [19, 20, 25]. The framework should therefore be implemented through a consortium model with clear rules for data locality, update protection, performance reporting, and shared governance.

Conclusion

A federated multi-task learning framework for hospital service demand prediction offers a conceptual pathway for collaborative forecasting across emergency, imaging, pharmacy, laboratory, and inpatient units. The central idea is to learn shared temporal representations across hospitals and service lines while keeping patient-level operational data within each local institution. This approach reframes hospital forecasting as a networked learning problem rather than a collection of isolated departmental prediction tasks.

The key strength of the proposed framework is its combination of simultaneous multi-service learning, privacy-preserving collaboration, and local personalization. A shared model could capture common demand dynamics, while task-specific and site-specific components would allow adaptation to each service line and hospital environment. By exchanging protected model updates instead of raw records, the system could support collective learning without requiring centralized patient-level data storage.

Several challenges remain before such a framework could be implemented in real hospital networks. Operational data definitions differ across institutions, governance requirements can be complex, and hospital leaders must see clear value before committing resources to federated collaboration. Future pilots would need to test technical feasibility, workflow integration, privacy safeguards, and decision-support usability under real operational conditions.

A multi-hospital consortium would be a suitable setting to demonstrate the framework and refine its governance model. Such a consortium could establish common feature definitions, privacy protocols, evaluation standards, and dashboard expectations while allowing each hospital to maintain local control. Over time, this type of privacy-preserving operational forecasting network could become an important foundation for coordinated, data-driven hospital capacity management.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 23 Nov 2025

Revised: 15 Dec 2025

Accepted: 29 Jan 2026

Published online: 25 February 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

- Porto BM, Fogliatto FS. Enhanced forecasting of emergency department patient arrivals using feature engineering approach and machine learning. *BMC Med Inform Decis Mak.* 2024;24(1):377.
- Becker AS, Erinjeri JP, Chaim J, Kastango N, Elnajjar P, Hricak H, et al. Automatic forecasting of radiology examination volume trends for optimal resource planning and allocation. *J Digit Imaging.* 2022;35(1):1–8.
- Mohammed EA, Naugler C. Open-source software for demand forecasting of clinical laboratory test volumes using time-series analysis. *J Pathol Inform.* 2017;8(1):7.
- Pall R, Gauthier Y, Auer S, Mowaswes W. Predicting drug shortages using pharmacy data and machine learning. *Health Care Manag Sci.* 2023;26(3):395–411.
- Mahmoudian Y, Nemati A, Safaei AS. A forecasting approach for hospital bed capacity planning using machine learning and deep learning with application to public hospitals. *Healthc Anal.* 2023;4:100245.
- Zhang Y, Yang Q. A survey on multi-task learning. *IEEE Trans Knowl Data Eng.* 2021;34(12):5586–609.
- Vandenhende S, Georgoulis S, Van Gansbeke W, Proesmans M, Dai D, Van Gool L. Multi-task learning for dense prediction tasks: A survey. *IEEE Trans Pattern Anal Mach Intell.* 2022;44(7):3614–33.
- Brisimi TS, Chen R, Mela T, Olshevsky A, Paschalidis IC, Shi W. Federated learning of predictive models from federated electronic health records. *Int J Med Inform.* 2018;112:59–67.
- 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: *Proc IEEE/CVF Conf Comput Vis Pattern Recognit*; 2021. p. 2423–32.

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.

Harrou F, Dairi A, Kadri F, Sun Y. Forecasting emergency department overcrowding: A deep learning framework. *Chaos Solitons Fractals.* 2020;139:110247.

Pala Z, Atıcı R, Yıldız E. Forecasting future monthly patient volume using deep learning and statistical models. *Wireless Pers Commun.* 2023;130(2):1479–502.

Rojas F, Yáñez J, Cortés M. Hybrid statistical–metaheuristic inventory modeling: Integrating SARIMAX with skew-normal and zero-inflated errors in clinical laboratory demand forecasting. *Mathematics.* 2025;13(18):3001.

Chang K, Balachandar N, Lam C, Yi D, Brown J, Beers A, et al. Distributed deep learning networks among institutions for medical imaging. *J Am Med Inform Assoc.* 2018;25(8):945–54.

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.

Tuominen J, Pulkkinen E, Peltonen J, Kanninen J, Oksala N, Palomäki A, et al. Forecasting emergency department occupancy with advanced machine learning models and multivariable input. *Int J Forecast.* 2024;40(4):1410–20.

Barrow D, Kourentzes N, Utley M, Kirkland H. Drug demand forecasting for hospital pharmacies using temporal hierarchies. *J Oper Res Soc.* 2026;77(1):1–8.

Lim B, Arık SÖ, Loeff N, Pfister T. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *Int J Forecast.* 2021;37(4):1748–64.

Bonawitz K, Ivanov V, Kreuter B, Marcedone A, McMahan HB, Patel S, et al. Practical secure aggregation for privacy-preserving machine learning. In: *Proc ACM SIGSAC Conf Comput Commun Secur*; 2017. p. 1175–91.

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.

Li T, Sahu AK, Talwalkar A, Smith V. Federated learning: Challenges, methods, and future directions. *IEEE Signal Process Mag.* 2020;37(3):50–60.

Whitt W, Zhang X. Forecasting arrivals and occupancy levels in an emergency department. *Oper Res Health Care.* 2019;21:1–8.

Yuan W, Ye G, Zhao X, Nguyen QV, Cao Y, Yin H. Tackling data heterogeneity in federated time series forecasting. *arXiv.* 2024;arXiv:2411.15716.

Smith V, Chiang CK, Sanjabi M, Talwalkar AS. Federated multi-task learning. In: *Adv Neural Inf Process Syst.* 2017;30.

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.

Cheng Q, Argon NT, Evans CS, Liu Y, Platts-Mills TF, Ziya S. Forecasting emergency department hourly occupancy using time series analysis. *Am J Emerg Med.* 2021;48:177–82.

Kairouz P, McMahan HB, Avent B, Bellet A, Bennis M, Bhagoji AN, et al. Advances and open problems in federated learning. *Found Trends Mach Learn.* 2021;14(1–2):1–210.

Sheller MJ, Reina GA, Edwards B, Martin J, Bakas S. Multi-institutional deep learning modeling without sharing patient data: A feasibility study on brain tumor segmentation. In: *Int MICCAI Brainlesion Workshop*; 2018. p. 92–104.

Manthe M, Duffner S, Lartizien C. Federated brain tumor segmentation: An extensive benchmark. *Med Image Anal.* 2024;97:103270.