

ORIGINAL RESEARCH

Open access

Predictive Model for Forecasting Diagnostic Service Demand Using Ambulatory Referral Volume, Seasonal Disease Trends, Physician Ordering Patterns, Equipment Availability, and Historical Appointment Backlogs

Ahmed Youssef^{1*}, Khaled Hassan¹, Mahmoud Elamin²

Abstract

Diagnostic services are central to clinical decision-making because imaging, laboratory testing, and cardiology diagnostics often determine the next step in diagnosis, treatment, or referral. Bottlenecks in these services can delay care pathways and increase wait times when demand rises faster than available capacity. Current forecasting approaches in diagnostic departments are often reactive and based on historical averages, recent appointment counts, or manual manager judgment. Such approaches may miss upstream signals such as referral surges, seasonal disease activity, and physician ordering behavior. This article proposes a predictive model for forecasting diagnostic service demand by integrating ambulatory referral volume, seasonal disease trends, physician ordering patterns, equipment availability, and historical appointment backlogs. The model is intended to support short- and medium-term capacity planning across diagnostic services. The proposed approach uses a supervised time-series forecasting framework, such as gradient boosting with temporal features or a recurrent neural architecture, trained on historical diagnostic order and scheduling data. Inputs would be engineered from referral streams, diagnostic ordering records, seasonal indicators, equipment schedules, and backlog measures. Conceptually, the model would generate daily or weekly demand forecasts for each diagnostic modality and service line. Forecasts would include uncertainty bounds and operational alerts when projected demand is expected to exceed available appointment capacity. The proposed predictive model could enable proactive capacity management in diagnostic departments. By anticipating demand before backlogs become severe, the model could support improved scheduling, better equipment utilization, and reduced patient waiting times.

Keywords Predictive modeling, Demand forecasting, Radiology operations, Diagnostic services, Laboratory capacity, Ambulatory referrals

*Correspondence:

Ahmed Youssef

ahmed.youssef@gmail.com

¹ Department of Digital Healthcare Systems, Faculty of Medicine, University of Khartoum, Khartoum, Sudan

² Department of Intelligent Clinical Analytics, Faculty of Engineering, Sudan University of Science and Technology, Khartoum, Sudan

Introduction

Diagnostic services occupy a critical position in modern care pathways because they translate clinical suspicion into actionable evidence for diagnosis, staging, monitoring, and treatment planning. Radiology, laboratory medicine, cardiology testing, and other diagnostic services are therefore not isolated technical departments but shared operational

resources that influence downstream clinical throughput. When demand exceeds appointment capacity, delays may propagate across ambulatory clinics, emergency departments, inpatient units, and specialty services, creating wider system inefficiency [1, 2]. Forecasting diagnostic demand is therefore an operational necessity rather than a purely statistical exercise, particularly when service volumes fluctuate across days, weeks, and seasons [3, 4].

Many diagnostic departments still rely on reactive capacity management, including recent averages, informal manager experience, or manual adjustment of appointment templates after congestion has already emerged. Time-series forecasting studies in healthcare have shown that demand can be modeled more systematically, but univariate approaches may overlook the upstream drivers that explain why diagnostic demand rises or falls [5, 6]. Radiology and laboratory services face especially complex planning problems because appointment volumes reflect both clinical need and operational supply, including available scanners, staff, protocols, specimen workflows, and service hours [7, 8]. A conceptual predictive model for diagnostic demand should therefore move beyond simple historical extrapolation toward a richer representation of care-seeking, ordering, capacity, and queue dynamics [9, 10].

Upstream signals are particularly important because diagnostic demand often begins before the appointment request appears in the scheduling system. Ambulatory referral volume, specialty mix, physician ordering tendencies, and seasonal disease activity can all act as leading indicators of future diagnostic procedures [11, 12]. For example, respiratory virus activity may increase chest imaging, inflammatory markers, microbiology testing, and cardiopulmonary evaluations, while allergy seasons and chronic disease follow-up patterns may shift laboratory or diagnostic testing loads in predictable ways [13, 14]. Physician ordering patterns also contribute to variation because clinicians and specialties may differ in threshold for diagnostic testing, follow-up intervals, and reliance on imaging or laboratory confirmation [15, 16].

The thesis of this article is that a predictive model combining demand-side and supply-side signals could generate more actionable diagnostic service forecasts than models based only on historical appointment counts. The proposed model would integrate ambulatory referral volume, seasonal disease trends, physician ordering patterns, equipment availability, and historical appointment backlog to produce modality-specific forecasts over short- and medium-term planning horizons. Such a model would be expected to support scheduling, staff allocation, equipment planning, and early warning for emerging backlogs, while remaining interpretable for operational managers. The article presents a conceptual predictive model design rather than an experimental evaluation, emphasizing model structure, feature logic, workflow integration, and evaluation strategy.

Background

Diagnostic service demand and capacity planning

Diagnostic service demand includes requests for imaging, laboratory testing, cardiology procedures, and related diagnostic appointments that must be matched to finite capacity. Operational performance in these settings is commonly framed around wait time, appointment availability, throughput, utilization, backlog, and the ability to accommodate urgent demand without displacing routine care. Radiology demand forecasting has been described as a capacity planning problem in which demand projections must align with appointment slots, service times, and operational constraints [1, 7]. Similar principles apply to laboratory and outpatient diagnostic services, where test volume forecasts can inform staffing, analyzer allocation, specimen logistics, and service continuity [2, 10].

Ambulatory referrals and physician ordering as demand drivers

Ambulatory referrals and physician ordering behavior are central drivers of diagnostic demand because most nonemergency diagnostic procedures originate from outpatient encounters, specialty consultations, chronic disease monitoring, or follow-up care. Demand is shaped not only by patient volume but also by clinician-level variation in ordering tendency, specialty-specific practice norms, and the timing of referral decisions. Forecasting approaches that incorporate

outpatient visit patterns and broader healthcare demand indicators suggest that upstream clinical activity can provide useful signals before diagnostic appointment pressure becomes visible in the schedule [4, 17]. A diagnostic demand model should therefore treat referral volume and ordering rates as structured predictors rather than background noise [5, 18].

Seasonal disease trends and demand cyclicity

Diagnostic demand is cyclic because clinical presentations vary across seasons, public health conditions, and environmental exposure patterns. Respiratory illness periods, allergy peaks, cardiovascular stressors, and infection surges may increase demand for imaging, laboratory panels, microbiology tests, and cardiopulmonary diagnostics. Forecasting studies that incorporate temporal structure and external predictors demonstrate the relevance of seasonality, environmental exposure, and time-varying health demand in predicting future service use [9, 13]. For diagnostic departments, seasonal disease trend features would be expected to improve planning by separating predictable cyclical demand from unexpected operational disruption [14, 19].

Equipment availability and queue dynamics

Diagnostic demand cannot be interpreted independently of equipment availability because observed appointment volume may be constrained by scanner slots, analyzer uptime, staffing levels, room availability, or maintenance schedules. A fully booked MRI schedule may reflect limited capacity rather than stable demand, while a growing backlog can represent unmet need that spills into future periods. Studies of radiology workflow, MRI slot prediction, bed capacity planning, and queue-sensitive healthcare forecasting indicate that operational capacity and demand are tightly coupled in service systems [8, 11, 20]. A diagnostic demand model should therefore include backlog depth, wait time, available slots, and equipment status to distinguish true demand changes from supply-driven bottlenecks [15, 21].

Forecasting methods in healthcare operations

Healthcare operations forecasting has used statistical time-series models, machine learning approaches, hybrid methods, and deep learning architectures to anticipate patient flows, outpatient visits, emergency demand, and resource use. Traditional methods such as autoregressive models and smoothing approaches are useful for structured temporal patterns, while machine learning models can represent nonlinear interactions among referrals, seasonal indicators, ordering behavior, capacity, and backlog. Recent work on outpatient forecasting, emergency flow prediction, and hospital capacity planning supports the broader premise that operational demand is forecastable when temporal and contextual covariates are properly represented [3, 6, 12]. In diagnostic services, the methodological challenge is to adapt these approaches to modality-specific workflows, constrained equipment, and clinically heterogeneous ordering streams [16, 22].

Model Development Overview

High-level forecasting pipeline

The proposed forecasting pipeline would begin with a scheduled daily or weekly extraction from ambulatory referral systems, electronic order tables, diagnostic scheduling systems, equipment availability records, and backlog reports. These data streams would be harmonized into modality-specific time series, where each forecast origin summarizes recent referrals, orders, appointment capacity, seasonal indicators, and queue conditions. The model would then output demand forecasts for imaging, laboratory, cardiology, or other diagnostic services across defined planning horizons, allowing managers to anticipate pressure before appointment queues expand [1, 2]. This design follows the broader direction of healthcare forecasting studies that connect operational data streams to planning decisions rather than treating prediction as a standalone analytic product [13, 14].

Figure 1 presents the proposed prediction-to-action workflow for transforming diagnostic referrals, ordering behavior, seasonal disease signals, equipment availability, and backlog data into manager-reviewed diagnostic service capacity actions.




Figure 1. Prediction-to-Action Workflow for Forecasting Diagnostic Service Demand Across Hospital Diagnostic Departments

Core input features

Core model inputs would include outpatient referral volume by clinic and specialty, seasonal epidemiological indicators, physician ordering rates, available diagnostic equipment slots, staffing-related capacity indicators, and current backlog length. Referral features would capture the volume and source mix of likely downstream diagnostic demand, while physician ordering features would represent variation in clinician and specialty test-ordering tendency. Equipment and backlog features would indicate whether observed appointment volume reflects true demand, constrained supply, or accumulated unmet need [20, 21]. Seasonal and temporal features would further help the model represent cyclical diagnostic pressure associated with predictable illness patterns and calendar effects [9, 19].

Design principles

The model should be proactive, modality-specific, adaptive, and interpretable for operational users. A proactive design would focus on short- to medium-term forecasting horizons that are long enough to adjust staffing, extend sessions, open slots, or reprioritize nonurgent appointments, while still close enough to reflect current operational conditions. Modality-specific modeling is important because MRI, CT, ultrasound, laboratory testing, and cardiology diagnostics have different lead times, equipment constraints, referral sources, and scheduling rules [7, 8]. Interpretability is also essential because managers must understand whether a forecasted surge is driven by referrals, seasonality, physician ordering behavior, backlog spillover, or reduced capacity [16, 23].

Data Sources and Feature Engineering

Extraction of referral and ordering data

Referral and ordering data would be extracted from electronic health record order tables, ambulatory visit records, referral management systems, and diagnostic scheduling interfaces. Imaging and laboratory referrals could be identified using procedure codes, order names, ordering departments, requested modality, priority category, and intended appointment date. Physician ordering patterns would be engineered as clinician-level and specialty-level rates over recent encounter windows, while preserving the conceptual distinction between clinical need and ordering tendency [5, 17]. Similar

upstream demand and outpatient forecasting work supports the use of ambulatory activity, visit patterns, and service demand indicators as predictors of future operational volume [4, 18].

Seasonal disease trend features

Seasonal disease trend features would represent time-varying clinical demand pressure from respiratory infections, influenza-like illness, allergy-related conditions, infectious disease surges, and other predictable epidemiological cycles. These features could be derived from syndromic surveillance reports, institutional test positivity trends, clinic diagnosis patterns, or public health indicators that are aligned temporally with diagnostic ordering. The purpose is not to diagnose disease but to encode external demand drivers that may precede increases in imaging, laboratory, or cardiopulmonary diagnostic requests [9, 13]. Temporal healthcare forecasting studies suggest that external conditions, cyclical patterns, and changing patient flow dynamics should be incorporated when service demand is sensitive to seasonality and public health context [14, 19].

Equipment availability and backlog metrics

Equipment availability features would be extracted from scheduling systems, equipment maintenance logs, scanner calendars, laboratory analyzer status reports, and staffing-linked capacity plans. These features would encode available appointment slots, blocked sessions, downtime, room closure, staffing constraints, and modality-specific service capacity at the forecast origin. Backlog metrics would summarize current queue depth, pending referrals, unscheduled orders, delayed appointments, and wait-time pressure, allowing the model to represent unmet demand that may spill into future periods [11, 20]. Capacity planning and radiology operations studies indicate that service time, equipment availability, and queue conditions should be treated as structural features in operational forecasting rather than as secondary administrative details [15, 24].

Table 1 summarizes the major predictor domains used by the diagnostic service demand forecasting model and explains how each domain supports practical capacity planning decisions.

Table 1. Predictor Domains and Operational Decision Relevance for Forecasting Diagnostic Service Demand

Predictor domain	Example operational variables	Forecasting role	Diagnostic service relevance	Practical decision supported
Ambulatory referral volume	Referrals by clinic, specialty, priority level, requested modality, and intended appointment timing	Acts as an upstream lead indicator of future diagnostic appointment demand	Captures demand before it appears as completed diagnostic appointments	Advance planning for MRI, CT, ultrasound, laboratory, and cardiology appointment slots
Diagnostic order stream	Electronic orders, order type, modality, priority, ordering department, and requested completion window	Converts clinical intent into structured demand signals	Links patient care pathways to service-line workload	Early identification of rising test-order pressure by diagnostic modality
Physician ordering patterns	Clinician-level ordering rates, specialty ordering intensity, clinic-specific testing tendency, and	Captures predictable variation in test-ordering behavior	Helps distinguish demand caused by patient volume from demand caused by	Targeted communication with high-volume referring clinics and proactive schedule adjustment

	recent changes in ordering behavior	across referral sources	ordering practice variation	
Seasonal disease trends	Respiratory illness indicators, influenza-like illness activity, allergy-related cycles, institutional test positivity trends, and public health surveillance signals	Encodes cyclical and surge-related demand pressure	Anticipates imaging, laboratory, and cardiopulmonary diagnostic increases during seasonal peaks	Temporary capacity expansion before predictable seasonal diagnostic surges
Equipment availability	Scanner slots, analyzer uptime, room availability, maintenance blocks, session closures, and modality-specific capacity	Represents supply-side limits that shape observable diagnostic throughput	Prevents the model from mistaking restricted capacity for reduced clinical demand	Rescheduling maintenance, opening additional sessions, or reallocating equipment time
Staffing and operational coverage	Technologist availability, laboratory staffing, radiologist or cardiologist coverage, scheduler availability, and extended-hour capacity	Modulates deliverable capacity and expected throughput	Explains why the same referral volume may produce different backlog patterns	Staffing adjustments, overtime planning, and cross-coverage decisions
Historical appointment backlog	Pending referrals, unscheduled orders, wait-time pressure, delayed appointments, and carryover demand	Captures unmet demand that may spill into future scheduling periods	Separates new demand from accumulated queue pressure	Backlog reduction plans, prioritization rules, and appointment-template redesign
Calendar and temporal structure	Day of week, holidays, month, seasonal cycle, prior-week demand, and rolling demand summaries	Helps the model learn recurring temporal patterns and short-term demand momentum	Supports daily or weekly operational forecasts	Timing of additional appointment sessions and manager alerts
Modality-specific service characteristics	Lead time, average scheduling complexity, urgency mix, protocol requirements, and service-line workflow constraints	Allows separate forecasting logic for services with different operational behavior	Prevents one-size-fits-all modeling across MRI, CT, laboratory, and cardiology diagnostics	Modality-specific capacity plans rather than generic department-wide adjustments

Predictive Model Architecture

Model choice and rationale

The proposed architecture could use either a gradient-boosted tree model or a recurrent neural time-series model, with separate modality-specific configurations for MRI, CT, ultrasound, laboratory testing, cardiology diagnostics, and other service lines. Gradient boosting would be appropriate when the goal is to capture nonlinear interactions among referral volume, ordering tendency, seasonality, backlog, and available capacity while maintaining operational interpretability. Recurrent or sequence-learning models would be appropriate when temporal dependence, lagged demand propagation,

and evolving backlog patterns are central to the forecasting task [12, 16]. The model choice should be justified by the diagnostic department's workflow complexity, data completeness, interpretability needs, and the planning horizon required by managers [22, 25].

Feature engineering for temporal prediction

Temporal feature engineering would include lagged referral volumes, rolling demand summaries, day-of-week indicators, holiday effects, seasonal disease covariates, recent ordering intensity, equipment status at forecast origin, and backlog carryover. These features would allow the model to distinguish recurring cycles from sudden demand shifts and to represent how unmet diagnostic demand accumulates over time. Hybrid and machine learning forecasting approaches in healthcare show that temporal covariates, recent flow patterns, and context-sensitive features can improve the conceptual richness of operational demand models [3, 6]. For diagnostic services, these engineered predictors would be expected to make forecasts more actionable because they connect projected demand to specific operational levers such as appointment templates, staffing, scanner allocation, and backlog management [26, 27].

Output: demand forecast with prediction intervals

The model would output projected daily or weekly appointment demand for each diagnostic modality, accompanied by prediction intervals and threshold-based alerts when expected demand exceeds available service capacity. These outputs would be designed for operational use, enabling managers to see whether a future pressure point is likely to arise from referral growth, seasonal disease activity, physician ordering behavior, backlog carryover, or reduced equipment availability. Forecast uncertainty is important because diagnostic managers need to plan capacity without treating point estimates as fixed operational truth [14, 22]. The forecast should therefore be presented as decision support rather than an automated scheduling command, preserving managerial review and local clinical judgment [23, 24].

Handling Temporal and Structural Factors

Backlog and unmet demand as a leading indicator

Historical appointment backlog should be treated as a leading indicator rather than merely an outcome of poor capacity alignment. When unscheduled referrals accumulate, future booking pressure may remain elevated even if new referral volume stabilizes, because deferred diagnostic need spills forward into later scheduling periods. The proposed model would encode backlog depth, delayed orders, pending appointment requests, and recent wait-time pressure as temporal predictors that could help distinguish ordinary cyclical demand from accumulated unmet need [11, 20]. Queue-sensitive healthcare forecasting and radiology operations studies support this logic because observed service volume may reflect both patient need and the capacity limits that prevent all demand from being scheduled immediately [15, 24].

Encoding physician ordering patterns

Physician ordering patterns would be represented through static and dynamic features that capture specialty, clinic, recent ordering tendency, and changes in ordering behavior over time. These features would not be used to judge individual clinicians but to anticipate downstream diagnostic volume from referral sources whose test-ordering patterns are predictably different. For example, a specialty clinic with a rising visit volume and a historically high imaging-ordering rate could generate greater future diagnostic demand than another clinic with similar visit volume but lower test-ordering intensity [5, 17]. Incorporating clinician and specialty ordering features would therefore help the model connect ambulatory care activity to modality-specific diagnostic demand in a way that simple appointment-count forecasting may miss [18, 23].

Adapting to sudden equipment downtime or reduced staffing

The model should be able to revise forecasts when scanner availability, laboratory analyzer capacity, room access, or staffing levels change unexpectedly. In this setting, the forecast should separate projected clinical demand from deliverable appointment capacity, because a temporary loss of MRI slots or laboratory processing capacity may reduce completed appointments while increasing unmet demand. Encoding equipment status at the forecast origin would allow the model to identify whether a future pressure point is likely to reflect true referral growth, backlog spillover, or a supply-side restriction [8, 20]. This distinction is operationally important because the appropriate response may differ across added staffing, rescheduled maintenance, temporary slot reallocation, or prioritization of urgent diagnostic requests [21, 24].

Model Interpretability for Operational Planning

Explaining surge predictions to managers

Interpretability is essential because diagnostic managers need to understand why the model expects demand to rise before they change appointment templates, staffing plans, or equipment allocation. SHAP-style explanation could show that a predicted CT or laboratory surge is primarily driven by rising respiratory illness indicators, increased outpatient referrals from specific clinics, and a backlog of unscheduled diagnostic orders. Such explanations would help managers distinguish a seasonal surge from a physician-ordering shift or a capacity-driven backlog effect, making the forecast more credible and actionable [16, 22]. Explainable forecasting is especially important when models are embedded in operational settings, because managers must be able to translate prediction drivers into practical capacity responses rather than simply accepting a black-box output [23, 28].

From forecast to actionable decision support

The forecast should be paired with decision-support logic that links projected demand pressure to feasible operational responses. When predicted demand exceeds available diagnostic capacity, the system could recommend actions such as opening additional appointment sessions, reallocating slots across modalities, prioritizing urgent orders, adjusting staffing coverage, or reviewing equipment downtime schedules. These recommendations should remain advisory, because final decisions require managerial judgment, clinical priority review, and awareness of local constraints [13, 26]. The model would therefore function as an operational planning aid rather than an autonomous scheduling system, consistent with healthcare forecasting approaches that emphasize decision support and workflow integration [27, 28].

Integration Into Diagnostic Service Workflow

Embedding forecasts in scheduling and command center dashboards

The proposed forecast would be embedded into diagnostic scheduling systems, radiology or laboratory operations dashboards, and hospital command center views. Daily or weekly updates could present projected demand, available capacity, backlog pressure, and uncertainty ranges by modality, referral source, and planning horizon. Embedding forecasts into routine operational displays would help managers compare expected demand with available service capacity before bottlenecks become severe [7, 21]. This design aligns with medical machine learning operations principles, which emphasize deployment pathways, workflow alignment, monitoring, and operational usability when predictive models are introduced into clinical environments [23].

Triggering proactive capacity adjustments

Forecasted demand above a predefined operational threshold could trigger alerts to diagnostic department managers, schedulers, and capacity planning teams. These alerts could recommend targeted responses such as extending hours for a specific modality, temporarily redistributing appointment slots, reserving capacity for urgent referrals, or coordinating with referring clinics to smooth demand. The alert logic should incorporate uncertainty so that managers can distinguish a

mild projected increase from a high-confidence capacity risk [14, 29]. Prior forecasting work in emergency and outpatient settings supports the use of predictive information to support proactive staffing, resource allocation, and planning decisions before crowding or backlog becomes operationally disruptive [13, 28].

Evaluation Strategy

Forecast accuracy metrics

The model should be evaluated using standard forecast accuracy metrics such as Mean Absolute Error, Mean Absolute Percentage Error, horizon-specific error, calibration of prediction intervals, and modality-specific forecast reliability. Evaluation should compare predicted daily or weekly demand against observed diagnostic appointment requests, completed appointments, and backlog-adjusted demand definitions, while recognizing that observed volume may be constrained by capacity. Accuracy should be assessed separately for modalities such as MRI, CT, ultrasound, laboratory testing, and cardiology diagnostics because each service has different temporal behavior and operational constraints [1, 2]. Comparative evaluation against simpler time-series baselines would help determine whether adding referrals, seasonality, ordering patterns, equipment availability, and backlog features provides meaningful conceptual value [3, 10].

Temporal validation and prospective simulation

Temporal validation should use walk-forward evaluation so that the model is trained on past periods and assessed on later periods without leaking future information. A prospective silent run could then generate forecasts in parallel with existing scheduling practice, allowing managers to observe how the model would have anticipated demand surges without yet changing live operations. This approach would be suitable for evaluating whether the model remains stable under seasonal changes, concept drift, equipment disruptions, and changing referral patterns [14, 16]. Studies comparing statistical, machine learning, and operational forecasting methods support the need for time-aware validation rather than random splitting when the target is future healthcare demand [25, 29].

Operational impact assessment

Operational impact assessment should examine whether model-guided planning could reduce patient waiting time, improve equipment utilization, decrease avoidable backlog growth, and reduce last-minute staffing corrections. Because this article proposes a predictive model rather than reporting experimental outcomes, these effects should be framed as intended evaluation endpoints rather than demonstrated results. The assessment should compare historical practice with model-informed planning through prospective simulation, silent deployment, or pilot implementation in high-volume diagnostic departments [12, 26]. Operational outcomes should be interpreted alongside forecast accuracy because a statistically accurate model may still fail to improve workflow if its outputs are not timely, interpretable, or actionable for managers [27, 28].

Table 2 outlines the validation and deployment readiness checks required before the proposed forecasting model can be safely embedded into diagnostic service operations.

Table 2. Validation and Deployment Readiness Framework for a Diagnostic Service Demand Forecasting Model

Readiness domain	What should be evaluated	Why it matters	Practical acceptance criterion	Implementation action if weak
Temporal validation	Whether the model is tested using future-facing walk-forward validation	Diagnostic demand is time-dependent and vulnerable to leakage	Forecasts remain conceptually stable across different time	Rebuild validation using strictly chronological forecast windows

	rather than random splitting	if future information influences training	periods and seasonal conditions	
Baseline comparison	Whether the model is compared with simple historical averages, seasonal averages, or standard time-series baselines	Complex models should add operational value beyond simpler planning methods	The model provides clearer, more actionable forecasts than basic demand extrapolation	Simplify the model or improve feature engineering before deployment
Modality-specific reliability	Whether forecast behavior is assessed separately for MRI, CT, ultrasound, laboratory, cardiology, and other diagnostic services	Each diagnostic modality has different referral patterns, capacity limits, and scheduling rules	Forecasts are interpretable and operationally plausible for each service line	Build separate modality models or use modality-specific feature sets
Prediction interval usefulness	Whether uncertainty ranges are calibrated and understandable to managers	Capacity decisions require awareness of uncertainty, not only point forecasts	Managers can distinguish routine variation from high-risk capacity pressure	Redesign forecast displays to emphasize uncertainty and confidence bands
Backlog sensitivity	Whether the model captures the effect of existing queues and unmet demand on future booking pressure	Backlog can create future demand even when new referrals stabilize	Forecasts reflect carryover demand and delayed appointment pressure	Add queue depth, wait-time, and pending-order features
Equipment and staffing sensitivity	Whether the model responds appropriately to scanner downtime, analyzer limits, closed sessions, or staffing shortages	Completed appointments may fall because capacity is unavailable, not because demand is lower	Forecast outputs separate clinical demand from deliverable appointment capacity	Improve equipment, staffing, and schedule availability feeds
Interpretability for managers	Whether forecast drivers are presented in practical terms such as referrals, seasonality, ordering, backlog, and capacity	Managers need explanations to translate forecasts into operational action	Forecast summaries identify major drivers of predicted demand pressure	Add SHAP-style explanations or structured driver summaries
Workflow integration	Whether forecasts appear in scheduling systems, operations dashboards, or command center views at the right time	A useful model must reach decision-makers before capacity decisions are finalized	Forecasts are available during daily or weekly planning routines	Embed outputs into existing scheduling and management workflows
Human oversight and governance	Whether operational decisions remain manager-reviewed and clinically appropriate	Forecasts should support planning, not automate clinical prioritization	Alerts are advisory and require human review before schedule changes	Define governance rules, escalation pathways, and review responsibilities
Prospective silent run	Whether the model is tested in live operations	Silent deployment reveals feasibility, data	Managers find forecasts timely,	Refine alert thresholds, display

	without changing decisions initially	latency, and alert burden before full use	understandable, and operationally relevant	format, and data refresh processes
Operational impact assessment	Whether the model is later assessed for wait-time management, utilization, backlog control, and staffing coordination	Forecast accuracy alone does not guarantee workflow improvement	Model-informed planning is linked to practical service improvement goals	Redesign implementation pathway if forecasts do not change decisions
Monitoring after deployment	Whether model drift, data quality, seasonal changes, and service-line changes are tracked over time	Diagnostic demand patterns can shift as referral behavior, disease patterns, and capacity change	Ongoing monitoring identifies when recalibration or retraining is needed	Establish periodic review, recalibration, and governance reporting

Limitations

Data latency and quality

The proposed model would depend on timely and accurate extraction of referral, ordering, scheduling, equipment, and backlog data. Referral entries may be delayed, order metadata may be inconsistent, seasonal indicators may be imperfect proxies for local diagnostic demand, and equipment status may not always be logged in a structured or real-time format. These data limitations could reduce forecast reliability or make model outputs less actionable during periods of rapid operational change [6, 19]. Data governance, monitoring, and local validation would therefore be necessary before using the model to guide staffing, appointment templates, or equipment allocation [22, 23].

Generalizability across diagnostic modalities

A model trained for MRI demand may not transfer directly to laboratory testing, cardiology diagnostics, CT, or ultrasound because each modality has different ordering thresholds, turnaround expectations, capacity constraints, and backlog dynamics. Some diagnostic services are scheduled days or weeks ahead, whereas others respond rapidly to urgent clinical demand, making a single universal forecasting architecture potentially inappropriate. Modality-specific models, hybrid architectures, or multi-task designs may be needed to balance shared temporal signals with local workflow differences [7, 20]. Generalizability should therefore be evaluated across diagnostic departments, referral sources, and operational contexts before the model is adopted as a broad capacity planning tool [21, 24].

Conclusion

The proposed predictive model is designed to forecast diagnostic service demand by integrating ambulatory referral volume, seasonal disease trends, physician ordering patterns, equipment availability, and historical appointment backlogs. It frames diagnostic demand as the result of interacting upstream clinical activity and downstream capacity constraints rather than as a simple continuation of previous appointment counts. By producing modality-specific short- and medium-term forecasts, the model could support earlier recognition of demand pressure across radiology, laboratory, cardiology, and related diagnostic services.

A major strength of the proposed framework is its fusion of upstream referral signals, seasonal epidemiological indicators, clinician ordering behavior, equipment capacity, and backlog dynamics into a single operational forecasting structure. This integrated view could help managers understand not only how much diagnostic demand may occur, but also why demand is expected to rise or fall. The model is therefore intended to support actionable planning rather than passive reporting.

Important challenges remain in data integration, feature reliability, modality-specific calibration, and prospective validation in live diagnostic environments. Forecasts must be carefully interpreted when observed appointment volume is constrained by equipment availability or staffing rather than by actual clinical need. Successful implementation would also require clear governance, transparent interpretation, and alignment with scheduling workflows.

Pilot deployments in high-volume radiology and laboratory departments would be a logical next step for evaluating feasibility and operational usefulness. Such pilots should assess whether forecast-informed planning can improve wait-time management, equipment utilization, staffing decisions, and backlog control. The broader objective is to shift diagnostic capacity planning from reactive adjustment to anticipatory, data-informed management.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 23 Dec 2024

Revised: 06 Feb 2025

Accepted: 10 Mar 2025

Published online: 20 July 2025

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

- Côté MJ, Smith MA. Forecasting the demand for radiology services. *Health Syst.* 2018;7(2):79-88.
- 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.
- Luo L, Luo L, Zhang X, He X. Hospital daily outpatient visits forecasting using a combinatorial model based on ARIMA and SES models. *BMC Health Serv Res.* 2017;17(1):469.

- Huang D, Wu Z. Forecasting outpatient visits using empirical mode decomposition coupled with back-propagation artificial neural networks optimized by particle swarm optimization. *PLoS One*. 2017;12(2):e0172539.
- Klute B, Homb A, Chen W, Stelpflug A. Predicting outpatient appointment demand using machine learning and traditional methods. *J Med Syst*. 2019;43(9):288.
- Asheim A, Bache-Wiig Bjørnsen LP, Næss-Pleyrn LE, Uleberg O, Dale J, Nilsen SM. Real-time forecasting of emergency department arrivals using prehospital data. *BMC Emerg Med*. 2019;19(1):42.
- Zhang Y, Luo L, Zhang F, Kong R, Yang J, Feng Y, et al. Emergency patient flow forecasting in the radiology department. *Health Informatics J*. 2020;26(4):2362-74.
- 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.
- Qiu H, Luo L, Su Z, Zhou L, Wang L, Chen Y, et al. Machine learning approaches to predict peak demand days of cardiovascular admissions considering environmental exposure. *BMC Med Inform Decis Mak*. 2020;20(1):83.
- Huang Y, Xu C, Ji M, Xiang W, He D. Medical service demand forecasting using a hybrid model based on ARIMA and self-adaptive filtering method. *BMC Med Inform Decis Mak*. 2020;20(1):237.
- Tello M, Reich ES, Puckey J, Maff R, Garcia-Arce A, Bhattacharya BS, et al. Machine learning based forecast for the prediction of inpatient bed demand. *BMC Med Inform Decis Mak*. 2022;22(1):55.
- Liu X, Gu F, Bai Z, Huang Q, Ma G. Forecasting of daily outpatient visits based on genetic programming. *Iran J Public Health*. 2022;51(6):1313-21.
- King Z, Farrington J, Utley M, Kung E, Elkhodair S, Harris S, et al. Machine learning for real-time aggregated prediction of hospital admission for emergency patients. *NPJ Digit Med*. 2022;5(1):104.
- Susnjak T, Maddigan P. Forecasting patient flows with pandemic induced concept drift using explainable machine learning. *EPJ Data Sci*. 2023;12(1):11.
- 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.
- Peláez-Rodríguez C, Torres-López R, Pérez-Aracil J, López-Laguna N, Sánchez-Rodríguez S, Salcedo-Sanz S. An explainable machine learning approach for hospital emergency department visits forecasting using continuous training and multi-model regression. *Comput Methods Programs Biomed*. 2024;245:108033.
- Zhou L, Zhu Q, Chen Q, Wang P, Huang H. Predicting hospital outpatient volume using XGBoost: a machine learning approach. *Sci Rep*. 2025;15(1):17028.
- Orhan F, Kurutkan MN. Predicting total healthcare demand using machine learning: separate and combined analysis of predisposing, enabling, and need factors. *BMC Health Serv Res*. 2025;25(1):366.
- Chen CW, Hsieh LL, Chu BX. Structural time series modelling for weekly forecasting of enterovirus outpatient, inpatient, and emergency department visits. *PLoS One*. 2025;20(5):e0323070.
- Wang X, Aski SN, Uhlemann F, Gupta V, Amthor T. Predicting slot lengths of MRI exams to decrease observed discrepancies between planning and execution. *Curr Probl Diagn Radiol*. 2024;53(3):359-68.
- Zander A, Reuter-Oppermann M. Data in ambulatory care logistics: What modelers need and what practice can offer. *Health Care Manag Sci*. 2025;28(3):505-26.

Blanco J, Ferreras M, Cosido O. Predictive modeling of hospital emergency department demand using artificial intelligence: A systematic review. *Int J Med Inform.* 2025;193:106215.

de Almeida JG, Messiou C, Withey SJ, Matos C, Koh DM, Papanikolaou N, et al. Medical machine learning operations: a framework to facilitate clinical AI development and deployment in radiology. *Eur Radiol.* 2025;35(11):6828-41.

Aloini D, Benevento E, Berdini M, Stefanini A. Predicting radiology service times for enhancing emergency department management. *Socio Econ Plann Sci.* 2025;99:102208.

Rahimi Z, Li N, Down DG, Arnold DM. Clustering-based demand forecasting with an application to immunoglobulin products. *Oper Res Health Care.* 2025;45:200469.

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.

Ricciardi C, Marino MR, Trunfio TA, Majolo M, Romano M, Amato F, et al. Evaluation of different machine learning algorithms for predicting the length of stay in the emergency departments: a single-centre study. *Front Digit Health.* 2024;5:1323849.

Vural O, Ozaydin B, Aram KY, Booth J, Lindsey BF, Ahmed A, et al. An artificial intelligence-based framework for predicting emergency department overcrowding: development and evaluation study. *JMIR Med Inform.* 2025;13:e73960.

Ackermann AE, Fani V, Bandinelli R, Sellitto MA. Short-term prediction in an emergency healthcare unit: comparison between ARIMA, ANN, and logistic map models. *Forecasting.* 2025;7(3):52.