

ORIGINAL RESEARCH

Open access

Machine Learning Model for Predicting Delayed Patient Placement after Emergency Admission Using Bed Assignment Logs, Isolation Needs, Unit Census, Nurse Staffing Levels, and Specialty Service Availability

Luis Herrera^{1*}, Daniela Rojas¹, Andres Castro²

Abstract

Emergency admissions frequently wait for inpatient bed placement when hospital capacity, infection control needs, staffing limitations, and specialty-bed requirements collide. These delays can prolong emergency department boarding and disrupt hospital-wide patient flow. Current bed management is often reactive, relying on bed coordinators, charge nurses, manual communication, and local escalation routines. Without a prospective warning system, teams may recognize an impending placement delay only after the admission queue has already stalled. The objective is to develop a machine learning model that predicts, at the time of admission decision, whether a patient is likely to experience delayed placement beyond a defined operational threshold. The model would use bed assignment logs, isolation requirements, unit census, nurse staffing levels, and specialty service availability as core predictors. A supervised classification model based on gradient-boosted trees would be trained on historical emergency admissions and linked operational data. The model would generate a placement delay risk score that can be refreshed as bed status, staffing, and unit conditions change. Conceptually, the model would identify admissions at elevated risk for delayed placement and attribute risk to operational constraints such as limited isolation rooms, high census, low staffing, or unavailable specialty beds. These explanations would give bed managers lead time to intervene before the delay becomes entrenched. This predictive model could shift hospital bed management from a reactive queue-based process to proactive, data-driven placement coordination. It would support earlier escalation, more targeted resource allocation, and improved alignment between emergency admissions and inpatient capacity.

Keywords Machine learning, Nurse staffing, Bed management, Emergency department boarding, Patient placement, Isolation precautions

*Correspondence:

Luis Herrera

luis.herrera@gmail.com

¹ Department of Clinical Informatics and Health Analytics, Faculty of Medicine, Pontifical Catholic University of Chile, Santiago, Chile

² Department of Intelligent Healthcare Engineering, Faculty of Engineering, University of Concepcion, Concepcion, Chile

Introduction

Delayed patient placement after an emergency admission is a major operational failure point because the admission decision has been made, but the patient remains physically in the emergency department while awaiting an inpatient bed. This creates downstream crowding, reduces emergency department treatment capacity, and may expose admitted patients to care environments not designed for prolonged inpatient management [1, 2]. Prior emergency department

admission prediction research has shown that information available early in the visit can be used to anticipate downstream disposition, suggesting that placement delay could likewise be framed as a prospective prediction problem [3, 4]. A predictive model focused specifically on post-admission placement delay would address the operational interval between acceptance for admission and arrival on the appropriate inpatient unit.

Current bed management practice often depends on bed managers, charge nurses, house supervisors, and unit-level staff reconciling multiple incomplete information streams in real time. Bed status boards, admission-discharge-transfer feeds, staffing rosters, pending discharges, infection-control restrictions, and specialty-service needs may be reviewed manually or semi-manually, which limits the ability to detect complex interactions before a delay occurs [5, 6]. Studies of emergency department throughput and boarding emphasize that staffing, crowding, and inpatient capacity are intertwined rather than isolated bottlenecks [2, 7]. A model-oriented approach could augment these human workflows by converting fragmented operational signals into a forward-looking placement risk estimate.

Hospitals increasingly possess the data infrastructure needed to support such prediction because bed assignment logs, electronic health records, admission-discharge-transfer feeds, census systems, and staffing schedules capture the key events that shape placement. Machine learning models have already been applied to emergency department triage, admission prediction, intensive care need, inpatient demand forecasting, and patient-flow forecasting, demonstrating the feasibility of integrating clinical and operational variables for real-time decision support [8-10]. Bed occupancy forecasting and operational patient-bed assignment studies further suggest that dynamic bed availability and assignment history can be translated into actionable features [6, 11]. These precedents support the development of a placement-delay model that is explicitly grounded in hospital operations rather than only clinical acuity.

The proposed model would forecast, for each emergency admission, the probability that placement will be delayed beyond a locally defined threshold. Its core contribution would be to combine bed assignment log history, isolation needs, unit census, nurse staffing levels, and specialty service availability into a single risk score interpretable by bed coordinators and nursing leaders. Rather than replacing human judgment, the model would function as an early warning layer that helps teams identify which admitted patients need proactive escalation, rapid cleaning, targeted staffing adjustment, or specialty-bed coordination. This article therefore outlines a conceptual predictive modeling framework for delayed patient placement after emergency admission using routinely available hospital operations data.

Background

Patient placement delays and ED boarding

Patient placement delay can be understood as the interval after an emergency admission decision during which the patient remains in the emergency department because an appropriate inpatient bed has not yet been assigned, prepared, staffed, or made available. This delay is closely related to emergency department boarding, a phenomenon associated with crowding, impaired throughput, and operational strain across the hospital [1, 2]. Prior predictive modeling studies have largely focused on admission disposition at triage or during the emergency encounter, but the next operational challenge is whether an admitted patient can actually be placed in the right unit in a timely manner [3, 12]. A placement-delay model would therefore extend emergency department prediction beyond the admission decision into the inpatient-capacity interface.

Bed assignment logs as a predictive data source

Bed assignment logs record the operational sequence of requests, candidate units, bed assignments, bed readiness, patient transport, and arrival, making them a rich source for identifying bottlenecks. When linked with admission-discharge-transfer data, these logs can reveal recurring patterns such as delays during shift change, repeated reassignments, unavailable specialty beds, or extended waits after a bed is nominally assigned [5, 6]. Forecasting work on inpatient bed demand and patient-bed assignment indicates that operational bed data can support both prediction and

optimization when structured into time-sensitive features [5, 11]. For delayed placement prediction, the bed log would provide both historical learning examples and live signals of current system congestion.

Isolation needs and their effect on placement

Isolation requirements can substantially narrow the set of acceptable beds because patients requiring contact, droplet, or airborne precautions may need private rooms, negative-pressure rooms, or cohorting restrictions. Infection-control precautions can affect patient flow by constraining bed choice, increasing environmental cleaning requirements, and limiting cross-unit flexibility [13]. During infectious disease surges, dynamic bed allocation becomes especially important because isolation-capable capacity may be consumed unevenly across units and time [14]. A predictive model should therefore treat isolation status not as a simple patient attribute, but as an operational constraint interacting with bed type, unit census, and room turnover.

Unit census and nurse staffing levels

Unit census and nurse staffing levels determine whether a theoretically empty bed is practically usable for an incoming admission. Emergency department throughput studies show that staffing constraints can adversely affect patient movement, while national boarding analyses highlight the relevance of crowding and staffing levels to emergency department function [2, 7]. Nurse staffing workload methods further demonstrate that required staffing depends on patient volume, acuity, and operational demand rather than headcount alone [15]. A placement-delay model should therefore include real-time census, occupancy pressure, nurse-to-patient ratios, skill mix, and shift-level staffing adequacy as predictors of whether a target unit can accept the next patient.

Specialty service availability

Specialty service availability introduces additional constraints because admitted patients may require telemetry, step-down monitoring, neurology observation, orthopedic beds, cardiology coverage, or other service-specific environments. Emergency department disposition and admission models show that clinical destination is not merely admitted versus discharged, but often depends on level of care and specialty pathway [16, 17]. Prediction across healthcare settings further suggests that site-specific care pathways and resource structures influence model behavior and must be represented in the prediction task [18]. For placement-delay prediction, specialty availability features would indicate whether the clinically appropriate bed type and responsible service are available at the moment of admission decision.

Model Development Overview

High-level predictive pipeline

The proposed predictive pipeline would begin when the emergency clinician enters or confirms the admission decision, because this is the moment when bed placement becomes operationally urgent. At that point, the system would assemble a feature vector from the patient record, bed assignment logs, admission-discharge-transfer feeds, unit census, staffing rosters, infection-control status, and specialty-bed availability [10, 19]. A supervised classification model would then output a placement delay risk score for the incoming admission, reflecting whether the patient is likely to wait beyond the hospital's defined delay threshold. Similar real-time emergency department and hospital admission prediction frameworks support the feasibility of this pipeline, provided that operational feeds are sufficiently timely and interpretable [20, 21].

Figure 1 illustrates the proposed end-to-end machine learning workflow for predicting delayed patient placement after emergency admission, from real-time operational data capture to calibrated risk scoring, explanation, validation, and human-supervised bed management action.




Figure 1. Machine Learning Workflow for Predicting Delayed Inpatient Placement after Emergency Admission

Core input features

Core input features would include current candidate-bed availability, pending discharges, recent bed-request queues, historical placement intervals by unit and shift, current patient isolation requirements, real-time unit census, and shift-level nurse staffing. Additional features would describe whether specialty-specific beds, monitored beds, or step-down capacity are available for the admitting service at the time of decision [5, 17]. Emergency department triage and admission prediction studies demonstrate that models can use mixed clinical and administrative variables, while bed demand forecasting studies show that operational capacity indicators can improve anticipation of inpatient resource needs [11, 22]. The distinctive feature of this model is that it would combine these domains around the specific operational endpoint of delayed placement.

Design principles

The model should be designed for real-time operation, workflow compatibility, interpretability, and support of human decision-making. Bed coordinators and charge nurses would need a concise risk score, a plain-language explanation, and a clear indication of which constraints are driving the predicted delay, rather than a black-box alert disconnected from action [23, 24]. Prior work on machine-learning triage and disposition prediction highlights the importance of embedding predictions into clinical workflows where staff can understand and contest the model's implication [18, 25]. Therefore, the model should support bed management decisions without automating placement decisions or overriding clinical prioritization.

Data Sources and Feature Engineering

Extracting from Bed assignment logs and ADT feeds

Bed assignment logs and admission-discharge-transfer feeds would provide the temporal backbone of the model by capturing bed request time, admission decision time, provisional assignment, bed readiness, transport initiation, and arrival on the inpatient unit. These timestamps would allow construction of historical placement intervals by unit, bed type, service line, day of week, and shift, as well as live features describing queue length and pending bed turnover [5, 6]. Patient-flow and bed occupancy forecasting studies show that operational event streams can be transformed into predictive signals when aligned to the relevant decision time [11, 26]. In this framework, feature engineering would deliberately avoid future leakage by using only information available at or before the admission decision.

Constructing isolation, census, and staffing features

Isolation features would be encoded as indicators for contact, droplet, airborne, protective, or combined precautions, while also reflecting whether compatible rooms are available or pending cleaning. Census and occupancy features would describe the current load on candidate units, the proportion of beds occupied or blocked, and the presence of discharge-dependent capacity [2, 14]. Staffing features would represent scheduled and actual nurse availability, nurse-to-patient pressure, skill mix, and shift context, recognizing that placement can be delayed when a bed exists but staffing cannot safely absorb another admission [7, 15]. These features would help the model distinguish between physical bed scarcity, infection-control incompatibility, and staffing-constrained capacity.

Specialty service availability features

Specialty service availability features would identify whether the patient requires a service-specific unit, telemetry-capable bed, step-down environment, intensive care pathway, or consulting team that affects placement feasibility. Emergency department models predicting inpatient and intensive care admission indicate that level-of-care and destination categories can be anticipated from information available during the emergency encounter [16, 17]. Hospital admission location prediction work further supports the idea that destination-specific modeling can improve operational alignment between patient need and bed assignment [10]. In the proposed model, specialty availability would be updated as admissions, discharges, transfers, and unit closures change the feasible destination set.

Table 1 defines the operational data domains and feature-engineering logic required to translate bed assignment logs, isolation needs, census pressure, staffing adequacy, and specialty-bed availability into predictors of delayed patient placement.

Table 1. Operational data domains, engineered feature logic, and expected predictive contribution for delayed patient placement risk.

Operational domain	Example source systems	Manuscript-specific feature examples	Predictive logic	Practical interpretation for bed management
Admission decision context	Emergency department record; admission order timestamp; ADT feed	Admission decision time; admitting service; requested level of care; decision-to-placement interval start	Defines the prediction anchor and prevents leakage by limiting features to information available at or before the admission decision	Establishes when the placement-delay risk score should be generated
Bed assignment history	Bed management platform; bed request logs; transfer logs	Prior placement intervals by unit, shift, service line, and bed type; number of reassignments; pending bed requests	Historical bottlenecks reveal recurring operational constraints and unit-specific delay patterns	Helps identify whether similar admissions have repeatedly experienced placement delays
Current bed availability	Bed board; ADT feed; environmental services status	Open beds; blocked beds; beds pending cleaning; discharge-dependent beds; target-unit bed availability	A physically listed bed may not be usable if it is blocked, dirty, pending discharge, or clinically incompatible	Distinguishes apparent capacity from actionable placement capacity
Isolation requirements	Infection-control documentation; EHR precautions; room-type registry	Contact, droplet, airborne, protective, or combined precautions; compatible	Isolation requirements restrict feasible destination beds and may interact with	Indicates whether infection-control compatibility is the

		room availability; isolation-room turnover	cleaning delay and unit census	dominant placement constraint
Unit census and occupancy pressure	Unit census system; ADT feed; capacity command center	Unit occupancy rate; admission queue length; discharge backlog; bed turnover pressure	Placement delay becomes more likely when preferred units are near capacity or discharge-dependent	Helps bed managers distinguish unit congestion from patient-specific placement complexity
Nurse staffing and skill mix	Staffing rosters; scheduling platform; charge nurse staffing updates	Scheduled nurses; actual available nurses; nurse-to-patient pressure; shift coverage; specialty skill availability	A bed may be physically available but operationally unusable if staffing cannot safely absorb another admission	Identifies staffing-constrained capacity rather than simple bed scarcity
Specialty service availability	Service census; specialty unit bed board; admission pathway documentation	Telemetry need; step-down bed requirement; monitored-bed availability; specialty-unit capacity; consulting service constraints	Specialty-specific placement needs narrow the feasible destination set and may produce delays even when general beds are available	Supports early negotiation for telemetry, step-down, or service-specific placement
Temporal operating conditions	Calendar; shift schedules; historical throughput patterns	Hour of day; weekend or holiday status; shift change window; seasonal respiratory-virus pressure	Similar census levels may carry different placement risk depending on staffing patterns, discharge timing, and seasonal demand	Helps explain predictable operational rhythms that affect placement flow
Data quality and latency indicators	Interface logs; missingness flags; update timestamps	Missing staffing feed; stale bed status; delayed isolation documentation; incomplete cleaning status	Real-time operational prediction is vulnerable to stale or incomplete data; missingness may itself be informative	Supports safe use by warning staff when the prediction is based on uncertain or delayed information

Predictive Model Architecture

Model choice and rationale

Gradient-boosted trees would be an appropriate conceptual model class because they can handle mixed categorical, binary, and continuous features while capturing nonlinear interactions among bed availability, isolation status, unit census, staffing, and service constraints. Prior emergency department admission and disposition models have used machine learning to integrate heterogeneous data available early in the care process, supporting the suitability of flexible supervised models for operational prediction [3, 8, 21]. A tree-based architecture would be expected to capture interaction patterns such as high census being more consequential when nurse staffing is low or isolation-compatible rooms are scarce. The model should nevertheless be compared conceptually with simpler baselines to ensure that added complexity provides interpretable operational value rather than unnecessary opacity [4, 25].

Input feature vector and preprocessing

The input feature vector would include patient-level indicators, admission timing, candidate-unit characteristics, current bed status, isolation requirements, census pressure, staffing adequacy, specialty-service availability, and recent placement history. Continuous features such as occupancy pressure or recent placement interval could be normalized or binned, while binary indicators would represent isolation needs, telemetry requirements, and specific bed-type availability [20, 22]. Rare categories, such as uncommon specialty pathways or infrequent isolation configurations, should be grouped carefully to preserve operational meaning while reducing sparsity [13, 14]. Missing values should be handled in a way that distinguishes true absence from delayed documentation, because real-time operational systems often contain incomplete but informative data [27, 28].

Output: placement delay probability

The model output would be a calibrated placement delay probability that can be translated into low-, medium-, and high-risk tiers for use by bed management staff. Rather than presenting the prediction as a deterministic decision, the score would communicate the relative likelihood that placement will exceed the hospital's predefined delay threshold given the current operational state [18, 23]. High-risk predictions could trigger review of pending discharges, expedited room cleaning, early transport coordination, staffing escalation, or specialty-service negotiation before the patient remains boarded for an extended period [1, 24]. This output design aligns the predictive task with practical bed management action while preserving human oversight and clinical prioritization [19, 26].

Handling Temporal Dynamics and Data Sparsity

Time-varying features and model refresh

Placement-delay prediction is inherently time-varying because the state of the hospital can change between the admission decision and physical transfer to the inpatient unit. A model could therefore run at the admission decision and refresh when key operational features change, including new discharges, bed cleaning completion, staffing updates, isolation-room release, or specialty-bed reassignment [5, 11]. Emergency department admission prediction and hospital admission location prediction studies support the use of evolving information streams to refine operational forecasts as care progresses [10, 29]. In this framework, refresh logic should be designed to update the risk score without creating excessive alert noise or undermining the bed manager's situational awareness.

Weekend, shift, and seasonal effects

Weekend, shift, and seasonal patterns should be represented because placement delay risk may increase when discharge processing slows, staffing patterns change, or infectious disease activity increases. Time-based features could include day of week, hour of admission decision, holiday periods, shift transition windows, and seasonal respiratory-virus pressure, all of which may affect bed turnover and unit acceptance capacity [2, 14]. Staffing and workload studies further indicate that throughput cannot be separated from the temporal distribution of available nurses and operational demand [7, 15]. These temporal features would help the model learn recurring operational rhythms without assuming that the same census level has the same meaning at all times.

Handling rare isolation or specialty constraints

Rare isolation conditions and unusual specialty-bed requirements may be infrequent but operationally important because they can sharply restrict feasible placement options. A model could address sparsity by grouping clinically similar constraints, applying regularisation, preserving specialist-defined flags, or using secondary review pathways when a rare condition is detected [13, 14]. Emergency department disposition and intensive care admission prediction work suggests that destination-specific needs can be modelled, but rare pathways require careful representation to avoid unstable or

misleading predictions [16, 17]. The goal would be to retain the operational signal of rare constraints while avoiding overfitting to unusual historical episodes.

Model Interpretability and Clinical Trust

Explaining placement delay predictions

Interpretability is essential because bed managers and charge nurses need to understand why an alert is generated before acting on it. SHAP-style explanations could show that a high-risk prediction is driven by limited step-down availability, low night-shift staffing, high census on the preferred unit, or the absence of compatible isolation rooms [23, 24]. Prior work on emergency department triage, clinical disposition, and prediction across healthcare settings emphasizes that model outputs are more useful when they can be translated into local workflow reasoning rather than presented as opaque scores [18, 25]. The explanation layer should therefore identify modifiable bottlenecks and distinguish them from fixed patient-level requirements.

Dashboard integration for action

The placement-delay risk score should be embedded in the existing bed-management dashboard rather than introduced as a separate tool that staff must check manually. A dashboard could display the risk tier, principal contributing factors, and context-specific suggested actions, such as expediting cleaning of an isolation-compatible room, confirming discharge timing, or escalating staffing review [5, 19]. Operational forecasting and bed-capacity management studies suggest that predictions are most valuable when they are connected to concrete decisions about bed assignment, resource allocation, and patient flow [6, 26]. This design would make the model a coordination aid rather than an abstract analytics product.

Deployment and Bed Management Integration

Integration with existing bed-tracking systems

Deployment would require integration with bed-tracking systems, admission-discharge-transfer messages, staffing platforms, infection-control documentation, and electronic health record data. A model service could listen for admission decisions and relevant operational updates, construct the live feature vector, and return a prediction through an application programming interface to the bed board or command center platform [10, 20]. Prior studies using emergency department and hospital operational data show that predictive models can be connected to real-time clinical workflows when data streams are structured around the decision point [17, 21]. The deployment architecture should prioritize reliability, auditability, and graceful degradation when one data source becomes delayed or unavailable.

Proactive escalation and mitigation protocols

High-risk predictions should be linked to predefined mitigation protocols so that the alert produces coordinated action rather than passive awareness. The protocol could notify the bed manager, charge nurse, environmental services, transport, and admitting service when the explanation indicates a specific operational constraint, such as isolation-bed scarcity, high target-unit census, or specialty-bed unavailability [1, 24]. Studies of boarding, staffing, and emergency department throughput show that placement delay is a system-level problem, so escalation should involve the departments that control bed release, unit acceptance, and patient movement [2, 7]. The model would therefore act as an early trigger for parallel work rather than a replacement for existing command-center judgment.

Evaluation Strategy

Predictive performance metrics

The model should be evaluated with discrimination, calibration, and decision-oriented measures that reflect the operational purpose of identifying patients likely to experience delayed placement. AUROC, precision-recall analysis for the delayed-placement class, calibration assessment, and threshold-specific review could be used conceptually, while avoiding reliance on a single headline performance measure [3, 8]. External validation studies of emergency department admission prediction emphasize that a model may appear useful in one environment but require careful prospective evaluation before adoption elsewhere [27, 28]. Evaluation should therefore ask whether the score is reliable enough to support earlier action, not merely whether it separates delayed from non-delayed cases retrospectively.

Temporal validation and multi-site testing

Temporal validation should use strictly later admissions for testing so that the model is assessed under future operational conditions rather than random partitions that blur seasonal, staffing, or capacity changes. Multi-site testing would be important because tertiary hospitals, community hospitals, and specialty centers differ in bed mix, service configuration, isolation capacity, nurse staffing models, and admission pathways [18, 28]. Prediction across healthcare settings and hospital admission location studies suggest that model transportability depends on the alignment between local workflow, available data, and destination structure [10, 18]. A robust evaluation plan should therefore examine whether feature definitions remain meaningful across hospitals before assuming generalizability.

Operational impact assessment

Operational evaluation should examine whether predictions would support earlier intervention, smoother bed coordination, and reduced boarding pressure after implementation. Relevant outcomes could include placement delay duration, emergency department boarding burden, bed assignment rework, escalation frequency, and staff-perceived usefulness, while avoiding unsupported claims of improvement before prospective testing [1, 2]. Bed-demand forecasting, dynamic bed allocation, and capacity optimization studies indicate that prediction alone is insufficient unless paired with operational actions that change resource timing and assignment decisions [5, 14, 26]. A future pilot should therefore evaluate the full decision-support pathway, including prediction, explanation, escalation, and response.

Table 2 presents an evaluation and governance framework showing how predictive accuracy, calibration, interpretability, fairness, data quality, and human oversight should be assessed before routine deployment of the placement-delay model.

Table 2. Evaluation, interpretability, governance, and operational-use framework for a delayed patient placement prediction model.

Evaluation or governance dimension	Recommended manuscript-specific approach	Why it matters for this model	Operational risk if neglected	Decision-use implication
Discrimination	Assess AUROC and precision-recall performance for delayed placement classification	The model must distinguish admissions likely to exceed the delay threshold from those likely to be placed on time	A model with weak discrimination may overwhelm staff with low-value alerts	Use risk scores only if they reliably prioritize admissions needing early review

Calibration	Evaluate calibration plots, observed-versus-predicted delay rates, and risk-tier reliability	Bed managers need the predicted probability to reflect realistic operational risk	Poor calibration could exaggerate or understate delay risk, leading to inappropriate escalation	Convert probabilities into locally validated low-, medium-, and high-risk tiers
Threshold selection	Review threshold-specific sensitivity, precision, workload burden, and escalation capacity	The best threshold depends on how many alerts the bed management team can safely act on	Too low a threshold may produce alert fatigue; too high a threshold may miss preventable delays	Select thresholds based on operational capacity, not only statistical performance
Temporal validation	Test on later admissions rather than random splits alone	Placement patterns vary by season, staffing, census, and hospital demand	Random splits may overestimate performance by mixing similar operational periods	Require temporally later testing before pilot deployment
Multi-site validation	Evaluate performance across tertiary, community, rural, pediatric, or specialty hospitals when applicable	Placement constraints differ by bed mix, service structure, isolation capacity, and staffing model	A model trained at one hospital may fail when transferred to another setting	Treat generalizability as empirical and require local recalibration
Interpretability	Use SHAP-style explanations or comparable local feature-attribution methods	Staff need to know whether the risk is driven by isolation rooms, staffing, census, or specialty-bed constraints	Opaque alerts may be ignored or misunderstood	Present the top modifiable bottlenecks alongside the risk tier
Data-quality monitoring	Track stale feeds, missing staffing updates, delayed isolation documentation, and inconsistent bed status	The model depends on time-sensitive operational data	Stale or incomplete inputs can generate misleading placement-risk estimates	Display data-quality warnings and allow staff to discount uncertain predictions
Fairness and equity review	Examine whether delay predictions or escalation patterns differ by age, language, disability, insurance status, race/ethnicity where ethically and legally available, or service line	Operational models can unintentionally reproduce inequitable placement patterns	Some patient groups may receive less timely escalation or be systematically deprioritized	Use fairness review to monitor decision support, not to automate bed allocation
Human oversight	Require bed manager, charge nurse, or supervisor review before any action	The model is designed to support coordination, not replace clinical or operational judgment	Automated use could misalign placement decisions with clinical priority or infection-control policy	Keep the model as an early warning and coordination aid

Prospective operational impact	Evaluate placement delay duration, ED boarding burden, escalation frequency, bed assignment rework, and staff-perceived usefulness	Retrospective accuracy does not prove that the model improves patient flow	A statistically accurate model may fail if no workflow action follows the alert	Test the full pathway: prediction, explanation, escalation, response, and outcome
--------------------------------	--	--	---	---

Limitations

Real-time data latency and accuracy

A major limitation is that the model would depend on operational data that may be delayed, incomplete, or inconsistently documented across systems. Bed assignment logs may not immediately reflect cleaning status, staffing rosters may differ from actual available nurses, and isolation status may change after additional clinical assessment [7, 13]. Real-time prediction studies in emergency care show that model reliability depends on the timeliness and accuracy of input variables at the moment the prediction is generated [20, 27]. The model should therefore include data-quality monitoring, missingness indicators, and workflow checks to prevent stale information from producing misleading placement-risk estimates.

Generalizability across hospital types

A model trained in a large tertiary center may not transfer directly to a community hospital, rural hospital, pediatric hospital, or specialty center because placement constraints differ by institutional structure. Service availability, monitored-bed supply, isolation-room capacity, nurse staffing models, and bed-management workflows may vary substantially across sites [18, 28]. Studies of admission prediction and patient-bed assignment suggest that operational models are shaped by local practice patterns as much as by general clinical features [6, 19]. Generalizability should therefore be treated as an empirical question requiring local validation, recalibration, and stakeholder review before routine deployment.

Conclusion

A machine learning model for predicting delayed patient placement after emergency admission could provide a structured way to anticipate which admitted patients are likely to remain in the emergency department while awaiting an appropriate inpatient bed. By generating a placement delay risk score at the time of admission decision, the model would focus attention on the operational interval where bed management teams can still intervene.

The central strength of the proposed model is its integration of bed assignment logs, isolation needs, unit census, nurse staffing levels, and specialty service availability into a single decision-support output. This integrated design would reflect the reality that placement delay is rarely caused by one factor alone, but by the interaction of physical capacity, staffing capacity, infection-control compatibility, and service-specific bed requirements.

Important challenges remain before such a model could be used safely in routine practice. These include ensuring real-time data accuracy, validating the model across sites, preventing alert fatigue, and embedding predictions into bed management culture so that staff understand how to act on the risk score.

Future work should prioritize pilot implementations and prospective studies that test whether prediction-informed escalation can reduce boarding pressure and improve patient flow. The model should ultimately be evaluated not only as an algorithm, but as part of a broader operational system for safer and more proactive hospital placement coordination.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 13 Apr 2024

Revised: 14 May 2024

Accepted: 03 Jul 2024

Published online: 20 July 2024

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

- Salehi L, Phalpher P, Valani R, Meaney C, Amin Q, Ferrari K, et al. Emergency department boarding: a descriptive analysis and measurement of impact on outcomes. *Can J Emerg Med.* 2018;20(6):929-37.
- Wretborn J, Henricson J, Ekelund U, Wilhelms DB. Prevalence of crowding, boarding and staffing levels in Swedish emergency departments: a national cross-sectional study. *BMC Emerg Med.* 2020;20(1):50.
- Hong WS, Haimovich AD, Taylor RA. Predicting hospital admission at emergency department triage using machine learning. *PLoS One.* 2018;13(7):e0201016.
- Barak-Corren Y, Israelit SH, Reis BY. Progressive prediction of hospitalisation in the emergency department: uncovering hidden patterns to improve patient flow. *Emerg Med J.* 2017;34(5):308-14.
- Tello M, Reich ES, Puckey J, Maff R, Garcia-Arce A, Feijoo F, et al. Machine learning based forecast for the prediction of inpatient bed demand. *BMC Med Inform Decis Mak.* 2022;22(1):55.
- Schäfer F, Walther M, Grimm DG, Hübner A. Combining machine learning and optimization for the operational patient-bed assignment problem. *Health Care Manag Sci.* 2023;26(4):785-803.

Ramsey Z, Palter JS, Hardwick J, Moskoff J, Christian EL, Bailitz J. Decreased nursing staffing adversely affects emergency department throughput metrics. *West J Emerg Med.* 2018;19(3):496-500.

Graham B, Bond R, Quinn M, Mulvenna M. Using data mining to predict hospital admissions from the emergency department. *IEEE Access.* 2018;6:10458-69.

Goto T, Camargo CA Jr, Faridi MK, Freishtat RJ, Hasegawa K. Machine learning-based prediction of clinical outcomes for children during emergency department triage. *JAMA Netw Open.* 2019;2(1):e186937.

El-Bouri R, Eyre DW, Watkinson P, Zhu T, Clifton DA. Hospital admission location prediction via deep interpretable networks for the year-round improvement of emergency patient care. *IEEE J Biomed Health Inform.* 2021;25(1):289-300.

Kutafina E, Bechtold I, Kabino K, Jonas SM. Recursive neural networks in hospital bed occupancy forecasting. *BMC Med Inform Decis Mak.* 2019;19(1):39.

Parker CA, Liu N, Wu SX, Shen Y, Lam SSW, Ong MEH. Predicting hospital admission at the emergency department triage: a novel prediction model. *Am J Emerg Med.* 2019;37(8):1498-504.

Tran K, Bell C, Stall N, Tomlinson G, McGeer A, Morris AM, et al. The effect of hospital isolation precautions on patient outcomes and cost of care: a multi-site, retrospective, propensity score-matched cohort study. *J Gen Intern Med.* 2017;32(3):262-8.

Ma X, Zhao X, Guo P. Cope with the COVID-19 pandemic: dynamic bed allocation and patient subsidization in a public healthcare system. *Int J Prod Econ.* 2022;243:108320.

Wundavalli L, Kumar P, Dutta S. Workload Indicators of Staffing Need as a tool to determine nurse staffing for a high volume academic emergency department: an observational study. *Int Emerg Nurs.* 2019;46:100780.

Fernandes M, Mendes R, Vieira SM, Leite F, Palos C, Johnson A, et al. Predicting intensive care unit admission among patients presenting to the emergency department using machine learning and natural language processing. *PLoS One.* 2020;15(3):e0229331.

Fenn A, Davis C, Buckland DM, Kapadia N, Nichols M, Gao M, et al. Development and validation of machine learning models to predict admission from emergency department to inpatient and intensive care units. *Ann Emerg Med.* 2021;78(2):290-302.

Barak-Corren Y, Chaudhari P, Perniciaro J, Waltzman M, Fine AM, Reis BY. Prediction across healthcare settings: a case study in predicting emergency department disposition. *NPJ Digit Med.* 2021;4(1):169.

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.

Yao LH, Leung KC, Tsai CL, Huang CH, Fu LC. A novel deep learning-based system for triage in the emergency department using electronic medical records: retrospective cohort study. *J Med Internet Res.* 2021;23(12):e27008.

Patel D, Cheetirala SN, Raut G, Tamegue J, Kia A, Glicksberg BS, et al. Predicting adult hospital admission from emergency department using machine learning: an inclusive gradient boosting model. *J Clin Med.* 2022;11(23):6888.

Raita Y, Goto T, Faridi MK, Brown DFM, Camargo CA Jr, Hasegawa K. Emergency department triage prediction of clinical outcomes using machine learning models. *Crit Care.* 2019;23(1):64.

Moreno-Sánchez PA, Aalto M, van Gils M. Prediction of patient flow in the emergency department using explainable artificial intelligence. *Digit Health.* 2024;10:20552076241264194.

Chang YH, Lin YC, Huang FW, Chen DM, Chung YT, Chen WK, et al. Using machine learning and natural language processing in triage for prediction of clinical disposition in the emergency department. *BMC Emerg Med.* 2024;24(1):237.

Levin S, Toerper M, Hamrock E, Hinson JS, Barnes S, Gardner H, et al. Machine-learning-based electronic triage more accurately differentiates patients with respect to clinical outcomes compared with the Emergency Severity Index. *Ann Emerg Med*. 2018;71(5):565-74.

Kim JN, Oh IH. Optimizing hospital bed capacity and resource allocation using inflow and outflow indices for effective healthcare management. *Inquiry*. 2024;61:469580241304244.

Kishore K, Braitberg G, Holmes NE, Bellomo R. Early prediction of hospital admission of emergency department patients. *Emerg Med Australas*. 2023;35(4):572-88.

Dadabhoy FZ, Driver L, McEvoy DS, Stevens R, Rubins D, Dutta S. Prospective external validation of a commercial model predicting the likelihood of inpatient admission from the emergency department. *Ann Emerg Med*. 2023;81(6):738-48.

Lucini FR, Fogliatto FS, da Silveira GJ, Neyeloff JL, Anzanello MJ, Kuchenbecker RS, et al. Text mining approach to predict hospital admissions using early medical records from the emergency department. *Int J Med Inform*. 2017;100:1-8.