

ORIGINAL RESEARCH

Open access

Artificial Intelligence-Enabled Hospital Command Center for Predicting Patient Transfer Bottlenecks Using Admission Requests, Unit Occupancy, Bed Cleaning Duration, Isolation Requirements, and Staffing Constraints

James Walker^{1*}, Olivia Harris¹

Abstract

Patient transfer delays from the emergency department, post-anaesthesia care unit, and outside facilities are a major source of hospital congestion. These delays can convert local unit constraints into system-wide capacity failure when demand and transfer readiness are not coordinated in real time. Hospital command centres increasingly monitor bed status, queue length, and operational pressure, but many remain reactive. They often identify congestion only after transfer queues have formed, particularly when demand, cleaning delays, isolation needs, and staffing shortages converge. This article proposes an AI-enabled hospital command centre framework for predicting patient transfer bottlenecks. The system would fuse real-time admission requests, unit occupancy, bed cleaning duration, isolation requirements, and staffing constraints into a dynamic bottleneck risk score. The framework includes an admission request projection module, a unit occupancy forecasting engine, a bed-cleaning-time estimation model, an isolation-delay calculator, and a staffing-aware transfer-capacity reasoner. Together, these components would estimate whether each receiving unit can absorb expected transfer demand. The system would provide command centre staff with a rolling risk map of potential transfer bottlenecks up to four hours ahead. This would support earlier load-balancing, cleaning prioritisation, staffing escalation, and transfer-routing decisions. A predictive command centre could shift hospital flow management from reactive queue monitoring to proactive bottleneck prevention. Such a framework should be evaluated prospectively before operational deployment.

Keywords Artificial intelligence, Patient flow, Hospital command center, Transfer bottlenecks, Bed management, Occupancy forecasting

*Correspondence:

James Walker
james.walker@gmail.com

¹ Department of Clinical Informatics and Healthcare Systems, Faculty of Medicine and Health Sciences, University of Edinburgh, Edinburgh, United Kingdom

Introduction

Patient transfer bottlenecks are a central mechanism through which local hospital constraints become system-wide congestion, particularly when emergency department admissions, intensive care step-downs, surgical recovery transfers, and external referrals compete for the same staffed beds. Studies of emergency department crowding and hospital flow have shown that delayed admission decisions, high occupancy, and boarding pressures are closely linked to downstream care disruption, safety risks, and reduced operational resilience [1-4]. High bed occupancy has also been associated with

deterioration in quality indicators and readmission-related strain, demonstrating that transfer delay is not simply a logistical nuisance but a system-level performance issue [5, 6]. In this context, predicting where the next transfer bottleneck could emerge is an essential capability for hospitals seeking to maintain access, safety, and continuity of care.

Centralised command centres were introduced to improve situational awareness by combining operational data streams into real-time dashboards for bed managers, nursing leaders, and hospital administrators. Early evaluations and benchmarking work suggest that these centres can strengthen visibility across admissions, discharges, transfers, and operational escalation, but their dominant function remains monitoring rather than anticipatory prediction [7-9]. Protocols and evaluations of AI-oriented command centres further indicate that the safety and patient-flow implications of these systems depend on how predictive outputs are integrated into decision-making routines [10, 11]. Therefore, the key design challenge is not only to display current queues, but to estimate whether specific units are likely to become transfer bottlenecks before the queue becomes clinically harmful.

AI and machine learning have been applied to hospital admission prediction, emergency department disposition, bed occupancy forecasting, and resource requirement estimation, creating a methodological foundation for predictive patient-flow systems [12-16]. These models show that data available near triage, admission request creation, and real-time census updates can support forward-looking operational estimates, although many applications remain focused on individual admission likelihood or aggregate occupancy rather than transfer-capacity constraints [2, 3, 17]. COVID-19 bed-demand models further illustrate the feasibility of forecasting ward and intensive care unit occupancy under changing demand conditions [18-21]. However, the transfer decision itself sits at the intersection of demand, bed readiness, infection-control suitability, and staff availability, which requires a more integrated AI systems framework.

This article proposes an AI-enabled hospital command centre framework that could predict patient transfer bottlenecks in real time by synthesising admission request volumes, unit occupancy, bed cleaning duration, isolation requirements, and staffing constraints. The framework builds on evidence that centralised operations centres can coordinate patient flow, that machine learning can forecast admission and occupancy pressures, and that operational bottlenecks emerge when physical beds, environmental services, infection-control rules, and staffed capacity are not aligned. Rather than presenting experimental results, the article specifies a conceptual architecture, data logic, prediction workflow, and evaluation strategy for a future deployable system. The aim is to define how a command centre could move from current-state visibility toward anticipatory transfer-capacity reasoning.

Background

The hospital transfer ecosystem

The hospital transfer ecosystem includes emergency department-to-floor admissions, intensive care unit step-downs, post-anaesthesia care unit transfers, inter-unit transfers, and direct admissions from outside facilities. Delays occur when receiving units lack clean, clinically appropriate, and staffed beds at the moment the sending location needs to move a patient, creating boarding in upstream areas and congestion across the hospital [1, 3, 4]. Emergency department admission prediction studies show that transfer demand can often be anticipated from early clinical and operational signals, but the receiving-unit consequences of that demand depend on bed supply and unit-level readiness [2, 13, 14]. A command centre framework for bottleneck prediction must therefore treat each transfer not as an isolated transaction, but as part of a dynamic system of competing requests and constrained destinations.

Table 1 analytically decomposes how distinct operational domains interact to generate transfer bottlenecks beyond simple bed availability constraints.

Table 1. Analytical Decomposition of Transfer Bottleneck Formation Across Operational Domains

Operational Domain	Primary Constraint Type	Temporal Behaviour	Interaction with Other Domains	Bottleneck Amplification Mechanism	Implication for Prediction Models
Admission Demand	Stochastic inflow pressure	Rapid, bursty, short-term variability	Drives occupancy and staffing strain	Demand exceeds short-term absorption capacity	Requires probabilistic forecasting and queue-based modelling
Unit Occupancy	Capacity saturation	Gradual accumulation with discharge uncertainty	Constrains available beds for incoming demand	High baseline occupancy reduces buffering capacity	Needs dynamic census forecasting with uncertainty distributions
Bed Cleaning	Throughput delay	Time-lagged, process-dependent	Delays conversion of discharge into usable capacity	Cleaning backlog creates hidden capacity deficit	Must model as time-to-availability, not binary state
Isolation Requirements	Compatibility constraint	Episodic but restrictive	Reduces effective bed pool and increases cleaning complexity	Limits substitutability of beds across patients	Requires constraint-based matching logic in prediction
Staffing Capacity	Functional capacity limitation	Shift-based fluctuation with predictable troughs	Governs whether physical beds are operationally usable	Understaffing converts available beds into unusable capacity	Must integrate staffing-aware capacity thresholds

Command centre operations and current capabilities

Hospital command centres and capacity coordination hubs are designed to centralise situational awareness by combining admissions, discharge planning, transfer queues, census status, and escalation workflows in one operational environment. Evaluations of command centre implementation suggest that such systems can improve data quality, visibility, and coordination, while benchmarking studies show variation in how hospitals operationalise these centres and connect them to patient-flow decision-making [7-9]. AI command centre evaluation protocols also indicate that predictive tools must be assessed not only for technical validity but for their effects on patient safety, workflow, and managerial action [10, 11]. The remaining gap is that many command centres can show where pressure exists now, but do not systematically estimate where the next transfer bottleneck will form.

The role of bed cleaning and environmental services

Bed turnover is a critical throughput process because a discharged bed does not become transfer-ready until it is cleaned, released, and confirmed as usable for the next patient. Just-in-time bed assignment research highlights that delays in the bed assignment process can be reduced when operational timing and readiness are actively managed, while patient-bed assignment models show that physical availability must be linked to decision rules for matching patients to beds [22, 23]. Infection-control modelling further demonstrates that isolation and organism-specific precautions can alter patient placement options and bed-use flexibility, especially when cleaning protocols and cohorting constraints are involved [24]. Therefore, a predictive transfer-bottleneck system should represent bed cleaning not as a binary status, but as a time-dependent process that modifies expected bed availability.

Staffing constraints and unit-level transfer capacity

A physical bed can be clean and technically available while still being unusable for transfer if nurse staffing, skill mix, or shift coverage prevents safe acceptance of another patient. Capacity and occupancy studies show that hospital flow cannot be understood solely through bed counts, because operational performance also depends on the ability of units to absorb demand without degrading safety or quality [4-6]. Patient-flow strategies aimed at reducing crowding often require coordination across departments, escalation pathways, and staffing-sensitive decisions rather than simple queue movement [25]. An AI-enabled command centre should therefore estimate staffed transfer capacity, not just physical bed supply.

Machine learning for hospital flow and bed management

Machine learning models have been used to predict hospital admission from emergency department triage, disposition decisions, and aggregate admission flows, creating useful building blocks for proactive hospital operations [2, 12-16]. Bed occupancy forecasting studies, including those developed during COVID-19, demonstrate that ward and intensive care demand can be projected under uncertainty and updated as new information becomes available [18-21, 26, 27]. Operational models for patient-bed assignment and admission estimation further show that forecasting and optimisation can be combined to support bed management decisions [22, 28, 29]. Yet these strands have not been fully integrated into a command centre system that predicts transfer bottlenecks by jointly modelling admission demand, occupancy, cleaning duration, isolation suitability, and staffing-constrained capacity.

Framework Overview

High-level architecture

The proposed framework is a multi-layer AI system embedded in the hospital command centre, where real-time admission-discharge-transfer data, bed management feeds, environmental services tracking, isolation flags, and staffing rosters are ingested into a common operational model. The architecture draws on command centre evidence showing the value of centralised flow visibility while extending it with predictive components informed by admission, occupancy, and bed assignment modelling [7, 9, 12, 22]. Each receiving unit would receive a continuously updated bottleneck risk score that reflects whether projected demand is likely to exceed clinically appropriate and staffed capacity. The dashboard would then translate this score into operational signals for bed managers, nursing supervisors, and hospital administrators.

Figure 1 illustrates the hierarchical architecture of the proposed AI-enabled command centre, showing how multi-source operational inputs are transformed into predictive bottleneck risk and actionable decision support.




Figure 1. Hierarchical AI Framework for Predicting Hospital Transfer Bottlenecks

Core inputs and outputs

The core inputs include the current admission request queue, projected emergency and elective demand, real-time unit occupancy, expected discharge timing, bed cleaning status, isolation requirements, and staffed-bed counts. Prior admission prediction and occupancy forecasting studies support the feasibility of using early patient-flow signals and time-series demand information to estimate future hospital load [2, 12, 18, 27]. The key outputs would be a per-unit bottleneck probability, an expected time-to-bottleneck, and a ranked explanation of contributing factors such as pending admissions, delayed cleaning, restricted isolation compatibility, or reduced staffed capacity. These outputs would be conceptual decision-support estimates rather than deterministic instructions, because transfer decisions remain clinically governed.

Design principles

The framework should be real-time, probabilistic, transparent, workflow-aligned, and capable of being evaluated before active deployment. Real-time operation is essential because command centre value depends on up-to-date patient-flow data, while probabilistic output is necessary because discharge timing, cleaning completion, and admission demand are uncertain [10, 11, 19, 21]. Transparency is equally important because command centre staff must understand why a unit is labelled high risk, especially when the model recommends cleaning prioritisation, staffing escalation, or transfer redistribution [8, 9, 25]. The system should therefore expose the operational drivers behind each prediction rather than presenting bottleneck risk as an unexplained algorithmic score.

Admission Request and Unit Occupancy Integration

Admission request projection

The admission request projection module would model incoming demand from the emergency department, operating rooms, procedural areas, and outside facilities using continuously updated time-series and queue-based signals. Research on emergency admission prediction shows that triage, operational, and early clinical variables can support

forward-looking estimates of hospitalisation and disposition, while aggregated admission forecasting demonstrates how real-time demand can be projected at the system level [2, 3, 12-14]. In the proposed framework, each newly registered admission request would update the demand distribution for relevant receiving units rather than merely adding to a visible queue. The output would be a probabilistic estimate of transfer demand over the next several hours, structured by unit type, acuity, service line, and clinical placement rules.

Real-time unit occupancy and discharge forecasting

The unit occupancy component would estimate future bed availability by combining current census, expected discharges, internal transfer movements, and occupancy trends within each unit. Bed occupancy forecasting studies show that ward and intensive care census can be estimated prospectively using statistical, simulation, and machine-learning approaches, especially when forecasts are updated as new data arrive [18-21, 26, 27]. In the proposed command centre, discharge timing would be represented as an uncertain distribution rather than a fixed timestamp, because late discharges and documentation delays can materially change transfer readiness. This would allow the system to estimate not only how many beds might open, but when they could become operationally useful.

Interaction of demand and supply

The core transfer-bottleneck logic would compare projected admission demand with expected unit-level bed supply and identify mismatches where cumulative demand exceeds expected availability within the prediction window. Prior models of hospital admissions, bed demand, and patient-bed assignment demonstrate that demand estimation and allocation logic need to be combined when hospitals seek to manage scarce receiving capacity [22, 23, 28, 29]. The system would also account for transfer policies, clinical appropriateness, and neighbouring-unit options, so that surplus demand is not treated as interchangeable across all beds. A bottleneck would therefore be conceptualised as a local failure of suitable, clean, and staffed capacity to absorb clinically appropriate demand at the required time.

Bed Cleaning Duration and Isolation Requirement Modelling

Bed cleaning duration estimation

The bed cleaning duration module would estimate the interval between discharge, environmental services assignment, cleaning start, cleaning completion, and bed-ready confirmation. Just-in-time bed assignment research indicates that reducing wait time for inpatients requires attention to operational timing and coordination around bed readiness, while patient-bed assignment models show that bed availability must be linked to assignment decisions rather than treated as static inventory [22, 23]. The proposed system would learn typical cleaning-to-ready patterns by unit, time of day, cleaning type, and current environmental services workload. These estimates would then modify the predicted availability of beds that are technically vacated but not yet ready to receive a transfer.

Isolation-driven delay amplification

Isolation requirements introduce both cleaning-duration effects and placement constraints because some patients require specific bed types, infection-control precautions, or compatibility rules that limit the usable bed pool. Simulation work incorporating infection-control policies for organisms such as methicillin-resistant *Staphylococcus aureus* and vancomycin-resistant *Enterococcus* illustrates how isolation and cohorting rules can affect patient flow and bed assignment options [24]. In the proposed framework, isolation status would act as both a time modifier for terminal cleaning and a constraint on which future admissions can use a given bed. This would prevent the command centre from overestimating transfer capacity when beds are physically available but clinically unsuitable for the next patient.

Dynamic bed-ready timestamp refinement

The framework would continuously refine bed-ready timestamps as cleaning status changes, new discharge information appears, or environmental services delays are recorded. Command centre studies emphasise the importance of real-time operational data quality, while AI command centre evaluation work suggests that predictive systems must be assessed in relation to the timeliness and reliability of the data streams on which they depend [7, 10, 11]. If a bed cleaning task is delayed, reprioritised, or completed earlier than expected, the bottleneck risk score for the affected unit would be recomputed immediately. This design would allow cleaning logistics to become an active component of transfer prediction rather than a downstream manual update.

Staffing Constraint and Transfer-Capacity Analysis

Staffed-bed model

The staffed-bed model would distinguish between physical bed availability and operational transfer capacity, because a bed should be considered transfer-ready only when it is clean, clinically appropriate, and supported by adequate staff coverage. Studies of hospital occupancy and patient flow show that high census pressure can degrade quality and intensify downstream congestion, while broader analyses of crowding-reduction strategies indicate that staffing-sensitive escalation is often needed to sustain flow [4-6, 25]. In the proposed framework, real-time staffing rosters, nurse-to-patient ratios, unit skill mix, and temporary closures would be linked to bed-ready forecasts so that the command centre does not overstate capacity. This would allow the AI system to identify a unit as high risk even when beds appear open, if those beds cannot safely receive transfers under current staffing conditions.

Floating and reallocation logic

The framework would incorporate floating and reallocation logic to evaluate whether a predicted bottleneck could be mitigated by temporarily shifting staff, redirecting patients, or opening contingency capacity. Patient-bed assignment and just-in-time bed allocation research suggests that operational benefit depends on matching patients to suitable beds while preserving system-wide capacity, rather than solving each unit's queue independently [22, 23]. Command centre benchmarking also shows that centralised coordination can support cross-unit action when local constraints threaten hospital-wide flow [8, 9]. In this framework, the system would present feasible reallocation options as decision-support recommendations, while leaving final action to operational leaders who can account for clinical appropriateness, staff fatigue, and local context.

Shift-change and break impact

Shift changes, handovers, meal breaks, and temporary staffing troughs can reduce a unit's ability to accept transfers even when bed and cleaning indicators appear favourable. Hospital flow studies emphasise that crowding is shaped by the interaction of demand, occupancy, staffing, and operational timing, while command centre evaluations highlight the need for real-time data that reflects current capacity rather than nominal capacity [4, 7, 11, 25]. The proposed system would therefore model predictable staffing troughs as time-varying reductions in transfer capacity and would increase bottleneck risk when projected demand overlaps with reduced staff availability. Such modelling would be expected to support earlier escalation, such as requesting float coverage or delaying non-urgent transfers until safe receiving capacity is restored.

Real-Time Bottleneck Prediction and Alerting Engine

Bottleneck risk score construction

The bottleneck prediction engine would compute a composite risk score for each receiving unit by combining projected admission demand, expected bed availability, cleaning-modified readiness, isolation compatibility, and staffing-constrained capacity. Prior work on admission prediction, occupancy forecasting, and patient-bed assignment provides the methodological basis for combining demand forecasts with allocation and capacity logic, although this framework

applies those methods specifically to transfer bottleneck prediction [19, 18, 22, 27]. A unit would be flagged as high risk when the estimated probability of demand exceeding suitable staffed capacity crosses a configurable threshold defined by hospital policy. The score would remain interpretable by displaying the relative contribution of admission volume, discharge uncertainty, cleaning delay, isolation restriction, and staffing limitation.

Tiered alerting and escalation

The alerting engine would translate bottleneck risk into tiered operational signals that match command centre escalation protocols. A moderate-risk signal could appear as a visual dashboard flag, a higher-risk state could notify the bed manager, and persistent or critical constraints could prompt escalation to hospital operations leadership [2, 4, 5]. This tiered design reflects evidence that command centre value depends on integration with real workflows and that predictive systems must be evaluated for their effects on safety, coordination, and staff behaviour [2, 3]. The alert would not dictate a single action but would focus attention on the unit, predicted timing, and operational drivers most likely to create transfer delay.

Command Centre Dashboard and Decision Support Integration

Predictive bottleneck display

The dashboard would present a map-style display in which each receiving unit is colour-coded by predicted bottleneck risk over the next operational window. Command centre studies show that centralised visibility can improve coordination, but predictive usefulness depends on whether the display helps staff understand both current state and emerging risk [7-9]. When command centre users inspect a unit, the interface would show contributing factors such as pending admission requests, uncertain discharge timing, delayed cleaning, isolation-limited beds, and insufficient staffed capacity. This design would make the system’s reasoning visible, supporting trust and reducing the risk that staff treat the risk score as an unexplained algorithmic label.

Table 2 presents a conceptual mapping between predictive signals, model components, and the operational interventions they enable within the command centre.

Table 2. Conceptual Mapping Between Predictive Signals, Model Components, and Operational Interventions

Predictive Signal	Model Component	Type of Inference	Bottleneck Insight Generated	Operational Intervention Lever	Decision-Making Value
Rising admission requests	Admission Projection Module	Probabilistic demand forecasting	Imminent surge in transfer demand	Pre-emptive bed allocation and load balancing	Enables early redistribution of demand
Delayed discharges	Occupancy Forecasting Engine	Time-to-availability estimation	Reduced near-term bed release	Discharge acceleration review	Improves timing alignment between supply and demand
Cleaning backlog	Cleaning-Time Estimation Model	Process duration prediction	Hidden delay in bed readiness	Cleaning prioritisation and staffing adjustment	Converts latent capacity into usable capacity faster
Isolation mismatch	Isolation Constraint	Compatibility filtering	Reduced usable bed pool	Cohorting and reassignment	Prevents overestimation of

	Processor			planning	effective capacity
Staffing shortfall	Staffing Capacity Model	Constraint-based capacity inference	Beds unavailable despite physical readiness	Staff reallocation or escalation	Aligns operational capacity with safety requirements
Multi-factor convergence	Transfer-Capacity Reasoning Engine	Integrated probabilistic reasoning	High-risk bottleneck formation	Coordinated multi-domain intervention	Supports system-level rather than local optimisation

Proactive decision support

The decision-support layer would suggest operational mitigations, such as prioritising a cleaning task, reviewing a pending discharge, floating staff, redirecting an appropriate transfer, or activating escalation capacity. Evidence from hospital crowding strategies and bed assignment research indicates that patient-flow improvement often requires coordinated cross-departmental action rather than isolated unit-level optimisation [23, 25]. The proposed system would estimate which mitigation is most relevant to the predicted bottleneck driver, while avoiding unsupported claims about guaranteed operational improvement. Its role would be to help command centre staff act earlier and more consistently, not to replace clinical judgement or local managerial authority.

Evaluation Strategy

Predictive performance metrics

The framework should be evaluated using unit-hour predictions that compare predicted bottleneck risk with observed transfer delay states, while avoiding premature claims of effectiveness before prospective validation. Admission prediction and resource forecasting studies commonly assess discrimination, calibration, and decision relevance, and similar principles would apply to transfer bottleneck prediction [2, 14-17]. Candidate metrics could include precision, recall, calibration, and probabilistic error across intensive care, medical-surgical, procedural, and specialty units. The evaluation should also examine whether the system remains reliable under high occupancy, infection-control restrictions, and staffing variation, because these are the conditions in which command centre decisions are most consequential.

Temporal and prospective validation

Temporal validation should use retrospective data ordered by time, so that the system is assessed on future operational periods rather than randomly mixed historical examples. COVID-19 bed occupancy and hospital demand forecasting studies illustrate the importance of evaluating predictions under changing demand conditions, while AI command centre protocols emphasise the need to study safety and workflow impacts before active intervention [10, 18-21, 26]. A prospective silent-mode deployment would be especially appropriate, because it would allow predicted bottlenecks to be compared with actual transfer delays without influencing staff behaviour during the evaluation period. This would help determine whether the framework is ready for operational trials in which alerts are visible to command centre users.

Operational impact assessment

Operational impact assessment should examine whether activating the system changes transfer-request-to-bed timing, emergency department boarding pressure, surgical flow disruption, staff workload, and user trust, while recognising that causal evidence would require careful study design. Prior command centre evaluations and crowding-reduction research suggest that operational technologies must be assessed in relation to patient safety, workflow integration, and coordination behaviour, not only predictive accuracy [7, 8, 11, 25]. Usability assessment should capture whether bed managers understand the risk score, whether explanations are actionable, and whether alerts align with escalation

protocols. The evaluation should also monitor for unintended consequences, such as excessive alerts, inappropriate transfer redistribution, or perceived loss of unit autonomy.

Limitations

Data latency and integration challenges

The framework depends on timely and accurate data from admission-discharge-transfer systems, bed management platforms, environmental services tracking, isolation documentation, and staffing rosters. Command centre studies indicate that data quality and integration are central to operational usefulness, while AI command centre evaluation work underscores the need to understand how technical systems affect safety and workflow in real settings [7, 10, 11]. Staffing data may lag actual availability, cleaning logs may not perfectly reflect bedside activity, and isolation status may change before it is updated in structured systems. These integration challenges mean the framework should include data-quality checks, uncertainty representation, and fallback workflows rather than assuming that all feeds are complete and current.

Behavioural and organisational barriers

Predictive bottleneck alerts may fail to improve flow if they are ignored, misunderstood, or perceived as conflicting with unit-level autonomy. Command centre benchmarking and patient-flow studies suggest that centralised coordination requires organisational trust, clear escalation authority, and alignment between data-driven recommendations and local clinical realities [8, 9, 25]. There is also a risk that staff may over-rely on predictions or discount alerts after false alarms, particularly during periods of extreme demand or staffing shortage. For this reason, implementation should include governance, user training, alert review, and continuous monitoring of how predictions influence real transfer decisions.

Conclusion

An AI-enabled hospital command centre for predicting patient transfer bottlenecks would provide a conceptual pathway from reactive capacity monitoring to proactive operational coordination. By estimating where transfer demand is likely to exceed suitable receiving capacity, the system could help command centre staff intervene before congestion becomes visible as a queue.

The main strength of the proposed framework is its integration of admission demand, unit occupancy, bed cleaning duration, isolation requirements, and staffing constraints into a unified bottleneck risk score. Its real-time updating and explanatory dashboard design would allow staff to see not only that a unit is at risk, but why the risk is emerging and which operational lever may be most relevant.

Important challenges remain before such a system could be used in routine hospital operations. These include integrating data across disparate vendor systems, ensuring that staffing and cleaning information are current, building trust among unit leaders, and evaluating whether predictive alerts improve coordination without creating unintended workflow burdens.

Future work should focus on pilot implementations in high-volume hospitals, silent-mode validation, prospective evaluation, and the development of open benchmark datasets for transfer bottleneck prediction. A carefully governed approach would allow hospitals to test whether predictive command centre systems can improve patient flow while preserving clinical judgement and operational accountability.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 12 Jan 2023

Revised: 19 Feb 2023

Accepted: 11 Mar 2023

Published online: 20 July 2023

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

- Lee SY, Chinnam RB, Dalkiran E, Krupp S, Nauss M. Prediction of emergency department patient disposition decision for proactive resource allocation for admission. *Health Care Manag Sci.* 2020;23(3):339-59.
<https://doi.org/10.1007/s10729-019-09485-2>.
- Hong WS, Haimovich AD, Taylor RA. Predicting hospital admission at emergency department triage using machine learning. *PLoS One.* 2018;13(7):e0201016.
<https://doi.org/10.1371/journal.pone.0201016>.
- 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.
<https://doi.org/10.1136/emered-2016-205566>.
- Long EF, Mathews KS. The boarding patient: effects of ICU and hospital occupancy surges on patient flow. *Prod Oper Manag.* 2018;27(12):2122-43.
<https://doi.org/10.1111/poms.12890>.
- Friebel R, Fisher R, Deeny SR, Gardner T, Molloy A, Steventon A. The implications of high bed occupancy rates on readmission rates in England: a longitudinal study. *Health Policy.* 2019;123(8):765-72.
<https://doi.org/10.1016/j.healthpol.2019.06.005>.
- Bosque-Mercader L, Siciliani L. The association between bed occupancy rates and hospital quality in the English National Health Service. *Eur J Health Econ.* 2023;24(2):209-36.
<https://doi.org/10.1007/s10198-022-01467-9>.

Mebrahtu TF, McInerney CD, Benn J, McCrorie C, Granger J, Lawton T, et al. The impact of hospital command centre on patient flow and data quality: findings from the UK National Health Service. *Int J Qual Health Care*. 2023;35(4):mzad072.

Franklin BJ, Mueller SK, Bates DW, Gandhi TK, Morris CA, Goralnick E. Use of hospital capacity command centers to improve patient flow and safety: a scoping review. *J Patient Saf*. 2022;18(6):e912-e921.
<https://doi.org/10.1097/PTS.0000000000000974>.

Franklin BJ, Yenduri R, Parekh VI, Fogerty RL, Scheulen JJ, Weiner JP, et al. Hospital capacity command centers: a benchmarking survey on an emerging mechanism to manage patient flow. *Jt Comm J Qual Patient Saf*. 2023;49(4):189-98.
<https://doi.org/10.1016/j.jcjq.2022.12.006>.

McInerney C, McCrorie C, Benn J, Habli I, Lawton T, Mebrahtu TF, et al. Evaluating the safety and patient impacts of an artificial intelligence command centre in acute hospital care: a mixed-methods protocol. *BMJ Open*. 2022;12(3):e054090.
<https://doi.org/10.1136/bmjopen-2021-054090>.

Mebrahtu TF, McInerney CD, Benn J, McCrorie C, Granger J, Lawton T, et al. Effect of a hospital command centre on patient safety: an interrupted time series study. *BMJ Health Care Inform*. 2023;30(1):e100653.
<https://doi.org/10.1136/bmjhci-2022-100653>.

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.
<https://doi.org/10.1038/s41746-022-00649-4>.

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.
<https://doi.org/10.1016/j.ajem.2018.10.060>.

Araz OM, Olson D, Ramirez-Nafarrate A. Predictive analytics for hospital admissions from the emergency department using triage information. *Int J Prod Econ*. 2019;208:199-207.
<https://doi.org/10.1016/j.ijpe.2018.11.018>.

Patel D, Cheetirala SN, Raut G, Tamegue J, Kia A, Rai A, et al. Predicting adult hospital admission from emergency department using machine learning: an inclusive gradient boosting model. *J Clin Med*. 2022;11(23):6888.
<https://doi.org/10.3390/jcm11236888>.

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.
<https://doi.org/10.1109/ACCESS.2018.2808843>.

Sterling NW, Brann F, Patzer RE, Di M, Koebbe M, Watkins C, et al. Prediction of emergency department resource requirements during triage: an application of current natural language processing techniques. *J Am Coll Emerg Physicians Open*. 2020;1(6):1676-83.
<https://doi.org/10.1002/emp2.12234>.

Baas S, Dijkstra S, Braaksma A, van Rooij P, Snijders FJ, Tiemessen N, et al. Real-time forecasting of COVID-19 bed occupancy in wards and intensive care units. *Health Care Manag Sci*. 2021;24(2):402-19.
<https://doi.org/10.1007/s10729-021-09549-6>.

Heins J, Schoenfelder J, Heider S, Heller AR, Brunner JO. A scalable forecasting framework to predict COVID-19 hospital bed occupancy. *INFORMS J Appl Anal*. 2022;52(6):508-23.
<https://doi.org/10.1287/inte.2022.1147>.

Bekker R, uit het Broek M, Koole G. Modeling COVID-19 hospital admissions and occupancy in the Netherlands. *Eur J Oper Res.* 2023;304(1):207-18.
<https://doi.org/10.1016/j.ejor.2022.02.051>.

Deschepper M, Eeckloo K, Malfait S, Benoit D, Callens S, Vansteelandt S. Prediction of hospital bed capacity during the COVID-19 pandemic. *BMC Health Serv Res.* 2021;21(1):468.
<https://doi.org/10.1186/s12913-021-06408-8>.

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):785-804.
<https://doi.org/10.1007/s10729-023-09641-2>.

Braaksma A, Copenhaver MS, Zenteno AC, Ugargh E, Levi R, Wiler JL, et al. Evaluation and implementation of a Just-In-Time bed-assignment strategy to reduce wait times for surgical inpatients. *Health Care Manag Sci.* 2023;26(3):501-15.
<https://doi.org/10.1007/s10729-022-09624-y>.

Shenoy ES, Lee H, Ryan EE, Hou T, Walensky RP, Hooper DC, et al. A discrete event simulation model of patient flow in a general hospital incorporating infection control policy for Methicillin-Resistant *Staphylococcus aureus* (MRSA) and Vancomycin-Resistant *Enterococcus* (VRE). *Med Decis Making.* 2018;38(2):246-61.
<https://doi.org/10.1177/0272989X17753572>.

Chang AM, Cohen DJ, Lin A, Augustine J, Handel DA, Howell E, et al. Hospital strategies for reducing emergency department crowding: a mixed-methods study. *Ann Emerg Med.* 2018;71(4):497-505.
<https://doi.org/10.1016/j.annemergmed.2017.07.022>.

Kociurzynski R, D'Ambrosio A, Papathanassopoulos A, Bürkin F, Hertweck S, Benning L, et al. Forecasting local hospital bed demand for COVID-19 using on-request simulations. *Sci Rep.* 2023;13(1):21321.
<https://doi.org/10.1038/s41598-023-48283-7>.

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.
<https://doi.org/10.1016/j.health.2023.100245>.

Teow KL, Tan KB, Phua HP, Zhu Z. Applying gravity model to predict demand of public hospital beds. *Oper Res Health Care.* 2018;17:65-70.
<https://doi.org/10.1016/j.orhc.2018.02.002>.

Latruwe T, Van der Wee M, Vanleenhove P, Michielsen K, Verbrugge S, Colle D. Improving inpatient and daycare admission estimates with gravity models. *Health Serv Outcomes Res Methodol.* 2023;23(4):452-67.
<https://doi.org/10.1007/s10742-023-00292-8>.