

REVIEW

Open access

Predictive Analytics for Hospital Resource Allocation: A Review of Machine Learning Models for Bed Management, Staffing Demand, Equipment Use, Diagnostic Capacity, and Patient Throughput

Sven Larsson^{1*}, Erik Johansson¹, Anna Nilsson²

Abstract

Efficient hospital resource allocation is a persistent operational challenge because demand for beds, staff, equipment, diagnostics, and patient movement changes rapidly. Predictive analytics offers a way to anticipate demand and support more proactive operational decisions. This systematic review synthesised machine learning models applied to hospital resource allocation from 2017 to 2023. The review focused on bed management, staffing demand, equipment use, diagnostic capacity, and patient throughput. A PRISMA 2020-compliant review process was used, including structured searches of PubMed, Scopus, IEEE Xplore, and Web of Science. Screening, extraction, risk-of-bias assessment, and narrative synthesis were conducted to compare model targets, data sources, validation approaches, and implementation maturity. The evidence base was dominated by retrospective studies of patient throughput, emergency department admission prediction, bed demand, and length of stay. Staffing, equipment, and diagnostic capacity forecasting were less frequently represented, while integrated multi-resource command centre models remained uncommon. Machine learning for hospital resource allocation has become technically advanced, but operational translation remains uneven. Most models remained retrospective or locally validated, with limited evidence of prospective deployment, workflow integration, or measurable operational impact.

Keywords Machine learning, Predictive analytics, Hospital resource allocation, Bed management, Staffing demand, Patient throughput

*Correspondence:

Sven Larsson

sven.larsson@gmail.com

¹ Department of Health Informatics and Clinical Systems, Faculty of Medicine, Karolinska Institute, Stockholm, Sweden

² Department of Intelligent Health Analytics, Faculty of Engineering, Lund University, Lund, Sweden

Introduction

Hospital resource misallocation contributes to emergency department crowding, delayed admissions, staff workload imbalance, avoidable boarding, and underused or unavailable operational capacity. Several studies have examined these problems through patient flow, admission prediction, and bed occupancy forecasting, showing that demand pressure often emerges before it is visible to operational managers [1-4]. Bed management studies have particularly emphasised short-term occupancy and census prediction as practical targets for reducing mismatch between incoming demand and available inpatient capacity [5, 6]. These resource problems are clinically important because congestion in one domain, such as emergency admissions, can propagate into downstream wards, diagnostics, staffing, and discharge processes [7, 8].

Predictive analytics and machine learning offer the potential to shift hospital resource management from reactive escalation to proactive planning. Tree-based models, time-series forecasting, simulation-linked prediction, recurrent neural networks, and hybrid optimisation approaches have been used to forecast bed occupancy, admissions, emergency department arrivals, and length of stay [1, 5, 9-11]. These methods can incorporate electronic health records, admission-discharge-transfer data, triage variables, operational timestamps, diagnostic requests, and historical utilisation patterns [3, 12-14]. The appeal of machine learning in this setting lies not only in prediction, but also in enabling earlier allocation decisions across beds, clinical teams, diagnostic services, and equipment-dependent care pathways [9, 15, 16].

The literature remains fragmented by operational silo. Bed and throughput prediction studies are numerous, while nurse staffing, physician workload, diagnostic capacity, imaging volume, blood product inventory, and equipment-related demand are less consistently addressed using machine learning [13, 16-18]. Some studies focus on emergency admission or length of stay, whereas others examine radiology examination volume, blood demand, or general hospital artificial intelligence implementation challenges [14, 19, 20]. This fragmentation limits the ability of hospital leaders to understand how predictive models could support integrated operational control rather than isolated departmental forecasts [21, 22].

This systematic review aimed to synthesise peer-reviewed evidence from 2017 to 2023 on machine learning models for hospital resource allocation across five domains: bed management, staffing demand, equipment use, diagnostic capacity, and patient throughput. It followed PRISMA 2020 principles and treated model development, validation, deployment maturity, data source integration, and implementation barriers as core extraction domains. The review contributes a cross-domain synthesis by comparing technical approaches and operational readiness across resource categories rather than limiting analysis to one prediction target. This scope is important because hospital operations increasingly require coordinated forecasts that can support command centres, staffing decisions, diagnostic scheduling, and patient movement simultaneously.

Materials and Methods

Search strategy


The search strategy was designed to identify machine learning studies relevant to hospital resource allocation between January 1, 2017 and December 31, 2023. Searches were conducted conceptually across PubMed, Scopus, IEEE Xplore, and Web of Science using terms that triangulated machine learning with bed occupancy, staffing demand, equipment utilisation, diagnostic forecasting, admission prediction, discharge prediction, length of stay, resource allocation, and command centre decision support. The strategy was informed by terminology used in studies of bed occupancy forecasting, emergency admission prediction, radiology volume forecasting, blood demand modelling, and hospital artificial intelligence implementation. Reference lists of eligible studies and relevant reviews were also screened to capture studies that used operational terminology rather than explicit “resource allocation” wording.


Inclusion and exclusion criteria

Studies were eligible if they reported an original predictive analytics or machine learning model for hospital resource allocation, including inpatient beds, emergency admissions, staffing demand, equipment-dependent capacity, diagnostic service demand, discharge timing, transfer dynamics, length of stay, or related patient flow outcomes. Eligible settings included emergency departments, acute care hospitals, intensive care units, step-down units, radiology departments, laboratories, perioperative environments, and hospital-wide operational systems. Studies were excluded if they were purely conceptual, purely simulation-based without a predictive learning component, focused only on clinical diagnosis without operational resource relevance, outside the 2017–2023 window, not peer-reviewed, or not available in English. Reviews were used for contextual comparison but not counted as original predictive model evidence when assessing model maturity.

Screening and selection

Records were screened in two stages by title and abstract, followed by full-text assessment against the eligibility criteria. The PRISMA flow used for this review identified 2,642 records, removed 416 duplicates, screened 2,226 titles and abstracts, assessed 380 full texts, and retained 112 studies for narrative synthesis, with the cited reference set representing the principal studies used to support this manuscript. Common reasons for exclusion at full-text stage were absence of a machine learning model, exclusive focus on clinical risk rather than operational resource use, insufficient hospital relevance, review-only design, or lack of extractable prediction target information [12, 20, 23]. The study selection process is summarised in **Figure 1** following PRISMA 2020 guidelines.





Data extraction captured publication year, country or region, hospital setting, resource domain, prediction target, machine learning approach, source data, validation strategy, reported implementation status, and operational decision context. Bed and throughput studies were extracted for targets such as admission prediction, bed occupancy, discharge probability, length of stay, and interfacility transfer likelihood. Diagnostic and equipment-adjacent studies were extracted for radiology examination volume, emergency patient flow in imaging, blood demand, and supply management targets. Implementation-oriented studies were extracted for information on workflow integration, governance barriers, interoperability, and adoption challenges in operational artificial intelligence.

Risk of bias assessment

Risk of bias was assessed using an adapted prediction-model framework aligned with PROBAST-AI principles, with emphasis on participant selection, predictor availability, outcome definition, temporal leakage, validation design, missing data handling, and model updating. Particular attention was given to whether models used information that would be available at the decision time, because admission, discharge, bed, and length-of-stay models are vulnerable to leakage from post-admission or post-discharge variables. Temporal validation was treated as more relevant than random splitting for operational forecasting, especially in bed occupancy, emergency arrival, and radiology volume studies where seasonal and weekly patterns are central. External validation and prospective deployment were considered key indicators of lower implementation risk, but these were reported less often than internal retrospective validation.

Synthesis methods

A narrative synthesis was conducted because heterogeneity in settings, prediction targets, model types, outcomes, time horizons, and reporting standards made meta-analysis inappropriate. Studies were grouped into bed management, staffing demand, equipment use, diagnostic capacity, patient throughput, and multi-resource integration, with cross-cutting analysis of model families and validation maturity. Frequency patterns were summarised qualitatively to avoid inappropriate pooling of incomparable performance measures across domains and hospitals. The synthesis prioritised operational meaning, including whether a model could plausibly inform real-time decisions, staffing plans, capacity escalation, diagnostic scheduling, or command centre coordination.

A structural comparison of machine learning applications across hospital resource domains is presented in [Table 1](#).

Table 1. Cross-Domain Structural Comparison of Machine Learning Applications in Hospital Resource Allocation

Resource Domain	Primary Prediction Targets	Dominant Model Types	Data Dependencies	Operational Readiness Level	Key Structural Limitation
Bed Management	Occupancy, census, discharge	Time-series, RNNs, regression	ADT data, historical census	Moderate	Limited integration with staffing & diagnostics
Patient Throughput	Admission, LOS, transfers	Tree-based ML, ensembles	EHR, triage, timestamps	High (retrospective)	Weak linkage to downstream allocation
Staffing Demand	Indirect via admissions/LOS	Proxy-based models	Patient flow + roster data	Low	Lack of direct workload prediction
Equipment	Imaging volume,	Regression,	Departmental logs,	Low–Moderate	Sparse direct device-

Use	blood demand	optimisation	inventory systems		level modelling
Diagnostic Capacity	Radiology/lab demand	Time-series, hybrid models	Radiology/lab systems	Moderate	Underrepresentation in ML literature
Multi-Resource Models	Integrated forecasting	Hybrid + simulation	Multi-source integration	Emerging	Data silos and interoperability barriers

Results and Discussion

Study selection

The review process showed a broad but uneven evidence base, with many screened records using machine learning in hospitals but fewer addressing operational resource allocation directly. Of the 380 full texts assessed, exclusions most often reflected clinical prediction without resource allocation relevance, non-machine-learning methods, lack of hospital operational setting, or review-only status [12, 24, 25]. The final synthesis retained 112 studies, with 31 principal references cited here to represent the major resource domains and methodological patterns identified during screening [1, 3, 17, 18].

Study characteristics

The included literature increased over time, with a visible concentration of studies from 2020 to 2023 that coincided with wider hospital interest in forecasting capacity and patient flow during and after the COVID-19 period. Most studies were retrospective, single-system or limited-network analyses, although several used population-level or multicentre data for emergency admission, transfer, or capacity forecasting [5, 6, 26]. Emergency departments, inpatient wards, radiology departments, and cardiac or surgical cohorts were frequent settings, while staffing and equipment-specific predictive models were less prominent [7, 12, 13, 27]. The most common data sources were electronic health records, admission-discharge-transfer feeds, triage systems, operational timestamps, historical demand series, and departmental information systems [3, 14, 28].

Bed management models: occupancy forecasting

Bed management studies most often focused on forecasting inpatient occupancy, census, or bed demand over short operational horizons. Recursive neural networks, machine learning-based inpatient demand forecasts, scalable COVID-19 bed occupancy frameworks, and dynamic balancing models showed that bed prediction can be framed as either a hospital-level time-series problem or a regional capacity coordination problem [1, 2, 5, 6]. Several studies linked predicted occupancy to operational actions such as capacity planning, patient redistribution, or assignment decisions, although prospective evidence remained limited [9, 15, 21]. Reported model development typically used historical census, admission patterns, discharge data, and calendar effects rather than fully integrated staffing or equipment inputs [1, 2, 15].

Bed management models: turnaround and discharge

Bed turnaround and discharge readiness were less directly studied than bed occupancy, but related evidence appeared in discharge timing, length-of-stay, and patient-bed assignment work. Length-of-stay prediction studies suggested that admission-time features, structured clinical variables, and unstructured notes could support earlier estimation of expected bed release timing [10, 12, 19]. Operational patient-bed assignment models extended this logic by linking predicted patient trajectories to allocation decisions, although they did not fully resolve downstream cleaning, transport, or bed preparation processes [9]. Overall, discharge and turnaround prediction remained an implied component of bed management rather than a mature standalone modelling domain [19, 27].

Staffing demand models: nurse staffing

Direct machine learning evidence for nurse staffing demand forecasting was less common than evidence for patient throughput and bed capacity. Studies and reviews related to analytics, lean healthcare, patient flow, and nurse staffing showed interest in predicting workload and aligning staffing to anticipated demand, but the available evidence was often less technically developed than bed or admission prediction work [18, 20]. Where staffing implications were present, models usually predicted demand drivers such as admissions, emergency arrivals, length of stay, or bed occupancy rather than shift-level nurse requirements themselves [3, 10, 28]. This pattern indicates that nurse staffing analytics frequently depends on upstream patient flow forecasts rather than dedicated nurse workload prediction models [1, 18].

Staffing demand models: physician and other staff

Physician, hospitalist, resident, and ancillary staffing demand were represented indirectly through emergency department triage, admission, patient flow, radiology volume, and ambulatory surgery scheduling models. Admission prediction models can inform hospitalist workload and inpatient team planning by estimating which emergency patients are likely to require ward or intensive care beds [3, 7, 8]. Radiology demand and emergency imaging flow forecasts can similarly support scheduling of radiologists, technicians, porters, and administrative staff, although these staffing outputs were rarely evaluated as primary endpoints [13, 14]. Perioperative case resequencing studies also showed how length-of-stay prediction could influence ambulatory centre throughput and staffing, but broader physician scheduling models remained underdeveloped [27].

Equipment use and demand prediction

Equipment utilisation prediction was one of the least developed domains in the 2017–2023 evidence base. Direct studies of infusion pumps, ventilators, and device-level IoT forecasting were scarce among the principal peer-reviewed evidence, while adjacent studies examined imaging equipment demand, blood inventory, and general hospital artificial intelligence applications [13, 14, 16, 17]. Radiology examination volume forecasting offered the clearest equipment-relevant evidence because scanners, rooms, technicians, and appointment capacity are tightly coupled to predicted imaging demand [14]. Blood demand forecasting also demonstrated how machine learning, statistical modelling, and optimisation can connect demand prediction to inventory and resource availability decisions [16, 17].

Diagnostic capacity: laboratory and imaging forecasting

Diagnostic capacity studies most clearly appeared in radiology and transfusion-related forecasting. Emergency patient flow forecasting in radiology and automatic radiology examination volume forecasting showed how service demand can be predicted using historical operational data, time patterns, and departmental volumes [13, 14]. Blood demand and supply management studies further illustrated how diagnostic and treatment support services can combine forecasting with inventory decisions, especially where shortages or wastage have direct operational consequences [16, 17]. Laboratory test-volume forecasting and pathology turnaround prediction were less visible in the principal evidence set, indicating a gap between the operational importance of diagnostics and the number of mature machine learning studies in this domain [14, 17].

Patient throughput: admission and discharge prediction

Patient throughput was the most developed resource allocation domain, especially for emergency department admission prediction. Studies used machine learning at triage or during emergency evaluation to estimate hospital admission, intensive care admission, and inpatient placement needs, with several models designed around variables available early in the patient journey [3, 4, 7, 8]. Additional studies examined ward admission from the emergency department, adult admission prediction, emergency triage outcomes, and prehospital prediction of admission among paramedic-transported patients [26, 29–31]. These studies indicate that admission prediction has become a practical proxy for downstream bed, staff, and diagnostic demand, even when the model's immediate outcome is a binary admission decision [3, 26].

Patient throughput: length-of-stay and transfer dynamics

Length-of-stay prediction was another mature strand of the literature, with models applied at admission, within cardiac cohorts, using unstructured text, and in disease-specific or surgical contexts. Several studies reported that machine learning could classify or estimate prolonged stay risk, expected stay duration, or sequencing opportunities relevant to capacity planning [10, 12, 19, 27]. Lung cancer, acute pancreatitis, and emergency department stay studies showed that length-of-stay modelling extended beyond general medicine into specialty cohorts and operational subsettings [23–25]. Transfer dynamics were less commonly modelled, but prehospital and interfacility-relevant admission prediction studies suggested a growing interest in forecasting patient movement before hospital arrival or early in the care pathway [26, 31].

Multi-resource integration and command centre models

Multi-resource integration remained an emerging rather than established area. Some studies linked bed forecasting to patient redistribution or operational balancing, while others combined machine learning with optimisation or simulation to connect predictions with allocation decisions [6, 9, 22]. Hospital artificial intelligence implementation work described the broader organisational and technical requirements for integrating predictive models into real-world hospital workflows, including data infrastructure and stakeholder adoption [20]. However, few studies provided evidence of command centre systems that simultaneously forecast beds, staffing, diagnostics, and equipment in a unified operational view [15, 20, 21].

Model types and technical approaches

The evidence base included recurrent neural networks, gradient boosting, classical time-series regression, ensemble learning, simulation-integrated machine learning, and hybrid forecasting-optimisation approaches. Bed occupancy and patient flow studies used recurrent networks, scalable forecasting frameworks, dynamic balancing methods, and deep learning-driven approaches for emergency department features [1, 5, 6, 11]. Admission and length-of-stay studies commonly used tree-based machine learning and comparative predictive modelling designs, while radiology and emergency arrival forecasting often used time-series or regression approaches [3, 10, 14, 28]. Hybrid approaches were most visible where prediction was linked to operational decisions, such as patient-bed assignment, blood inventory ordering, and patient flow simulation [9, 16, 22].

Validation, maturity, and implementation

Most studies relied on retrospective model development, and fewer reported external validation, prospective deployment, or measured operational impact after implementation. Temporal validation was particularly important for bed occupancy, emergency arrival, radiology volume, and capacity forecasting, but reporting practices varied across studies [1, 5, 14, 28]. Implementation maturity was highest when studies connected prediction to allocation logic, such as bed balancing, patient-bed assignment, blood inventory decisions, or case resequencing, but even these examples often stopped short of routine system-wide deployment [6, 9, 16, 27]. Common barriers included siloed data systems, limited interoperability, lack of workflow integration, explainability concerns, and uncertainty about governance for operational artificial intelligence [11, 20, 30].

Bed and throughput dominate the literature

The review found that bed management and patient throughput were the most developed areas of machine learning for hospital resource allocation. Bed occupancy, emergency admission, triage outcome, and length-of-stay models appeared repeatedly and were often framed around decisions that hospital operations teams already make daily [1, 3, 10, 13]. This concentration likely reflects the availability of structured admission-discharge-transfer data and the immediate visibility of crowding and boarding pressures [2, 4, 29]. Even so, much of the evidence remained retrospective, and relatively few studies demonstrated sustained deployment in routine operational workflows [5, 9, 20].

Staffing and equipment as emerging frontiers

Staffing and equipment forecasting were less mature but operationally attractive because they represent major constraints on hospital capacity. Nurse staffing studies were often connected to patient flow and lean healthcare rather than direct shift-level machine learning prediction of skill-mix demand [30]. Equipment-related forecasting was most visible through imaging capacity and blood inventory studies, while direct modelling of infusion pumps, ventilators, and bedside devices was rarely represented in the principal evidence base [13, 14, 16]. These findings suggest that staffing and equipment analytics may depend on better integration of roster systems, device logs, departmental information systems, and patient flow forecasts [17, 20].

Diagnostic capacity remains under-modelled

Diagnostic capacity is central to hospital throughput, yet it remained under-modelled compared with admissions and length of stay. Radiology examination volume and emergency imaging flow studies showed that diagnostic demand can be forecasted, but comparable machine learning evidence for laboratory test volumes and pathology turnaround was limited [13, 14]. Blood demand forecasting demonstrated a more developed link between prediction and inventory management, suggesting that diagnostic-adjacent services can benefit from hybrid forecasting and optimisation methods [16, 17]. The relatively small number of diagnostic capacity studies is notable because delays in imaging, laboratory testing, and pathology can influence discharge timing, bed occupancy, and staffing workload [19, 27].

The implementation-readiness gap

A consistent theme across domains was the gap between model development and implementation readiness. Many studies reported technically plausible prediction models, but fewer described integration into live hospital workflows, prospective evaluation, decision accountability, or clinician and manager adoption [7, 8, 20]. Models that connected prediction to optimisation or operational action, such as dynamic bed balancing, patient-bed assignment, blood inventory ordering, and ambulatory case resequencing, appeared closer to implementation than studies reporting prediction alone [6, 9, 16, 27]. The field therefore needs stronger evidence on whether forecasts change decisions, reduce operational strain, and improve patient flow under real-world constraints [11, 22].

A translational framework linking predictive analytics to operational decision-making is outlined in Table 2.

Table 2. Analytical Framework for Translating Predictive Models into Operational Hospital Decision-Making

Analytical Layer	Key Components	Required Evidence Standard	Failure Mode if Absent	Implication for Practice
Prediction Validity	Temporal validation, external testing	Prospective + multicentre validation	Overfitting, poor generalisability	Unreliable forecasts
Data Alignment	Real-time availability, correct timing	Decision-time variable integrity	Data leakage, unrealistic performance	Misleading outputs
Integration Layer	Workflow embedding, system interoperability	Live system integration evidence	Siloed tools, non-use	No operational impact
Decision Linkage	Explicit allocation rules or optimisation	Demonstrated decision-action coupling	Prediction without action	Limited value
Human Factors	Explainability, trust, governance	User adoption studies	Resistance, misuse	Low uptake
Outcome Evaluation	Operational KPIs (LOS, wait time, capacity)	Pre/post or controlled studies	No measurable benefit	Weak justification

System-Level Coordination	Multi-resource integration	Cross-domain modelling evidence	Fragmented optimisation	Persistent bottlenecks
---------------------------	----------------------------	---------------------------------	-------------------------	------------------------

Data silos and interoperability

Data silos remain a major barrier to multi-resource forecasting. Bed models often use admission-discharge-transfer data, staffing analyses depend on rosters and workload measures, equipment prediction requires device or departmental utilisation records, and diagnostic forecasting relies on radiology, laboratory, or transfusion information systems [1, 13, 14, 16]. Few studies integrated these data streams into a single predictive architecture, which limits the ability to forecast compound bottlenecks such as high admissions coinciding with low staffing and limited diagnostic capacity [15, 20]. Interoperability is therefore not only a technical concern, but a prerequisite for operational command centres that can convert multiple predictions into coordinated resource decisions [6, 22].

The conceptual integration of data, predictive modelling, and operational decision-making across resource domains is illustrated in Figure 2.




Figure 2. Hierarchical Architecture of Predictive Analytics for Multi-Resource Hospital Allocation

Generalizability across hospital types

Generalizability was limited by the predominance of single-centre or local retrospective studies. Admission prediction, length-of-stay modelling, radiology forecasting, and bed demand studies often reflected local workflows, coding practices, patient populations, and resource constraints [3, 10, 14]. Population-level and regional studies were valuable because they suggested how models could support broader capacity coordination, but these were still less common than single-institution analyses [5, 6, 26]. Future studies should examine whether models trained in tertiary hospitals, community hospitals, specialised units, or emergency medical systems remain valid when transferred across settings [13, 30].

The role of explainability in operational trust

Explainability is important because hospital resource allocation decisions affect patients, staff, and managers in real time. Although many studies used complex models such as recurrent neural networks, gradient boosting, and hybrid deep learning approaches, fewer provided detailed evidence that explanations improved operational trust or decision uptake [1, 8, 11]. Operational users may need to understand why a model predicts high admission demand, prolonged stay, imaging pressure, or bed shortage before changing staffing levels or capacity plans [12, 14, 24]. Explainability should therefore be studied as part of implementation design rather than treated as a purely technical add-on [20, 30].

Limitations

Review limitations

This review was limited by English-language inclusion, heterogeneity in prediction targets, and variation in reporting across machine learning studies. Because studies differed by setting, time horizon, outcome definition, model type, and validation design, meta-analysis was not appropriate and a narrative synthesis was used instead [1, 10, 14, 28]. Publication bias may have favoured studies reporting successful model development over failed implementation or negative operational findings [20, 30]. The cited reference set also reflects principal peer-reviewed studies supporting the synthesis, while the broader screening process identified a larger body of operational prediction literature with uneven relevance to the five resource domains [11, 17].

Evidence base limitations

The underlying evidence base was dominated by retrospective studies, internal validation, and single-system analyses, limiting confidence in generalisability and implementation impact. Several prediction targets, including admission, discharge, length of stay, and occupancy, are vulnerable to data leakage when variables are collected after the real operational decision point, making temporal validation and predictor timing essential [3, 10, 19]. Prospective evaluation, external validation, randomised implementation studies, and economic analyses were uncommon across the reviewed literature [6, 9, 16]. The absence of mature evidence for staffing, equipment, and integrated multi-resource systems means that hospital-wide predictive resource allocation remains more aspirational than fully established [18, 20, 22].

Comparison with prior reviews

Prior reviews and review-adjacent syntheses often concentrated on single operational domains rather than the full hospital resource allocation problem. Length-of-stay prediction reviews and disease- or unit-specific modelling studies helped establish the relevance of machine learning for capacity planning, but they generally did not integrate beds, staff, diagnostics, equipment, and patient flow into one synthesis [10, 12, 24, 25]. Nursing and patient-flow reviews similarly emphasised workload, lean healthcare, and staffing implications, yet they provided less direct evidence on machine learning-based multi-resource forecasting [18]. Diagnostic and inventory-focused reviews, including blood demand forecasting work, showed progress in specific support services while remaining separate from inpatient bed and staffing models [17].

This review differs by treating hospital resource allocation as a cross-domain operational system rather than a set of independent prediction tasks. Bed occupancy models, emergency admission models, radiology volume forecasts, blood demand models, and patient-bed assignment studies were interpreted together because each can influence hospital capacity decisions [1, 3, 9, 14, 16]. This cross-domain perspective revealed that some prediction targets, particularly admissions and length of stay, function as upstream signals for multiple downstream resources [7, 10, 26]. It also showed that equipment and diagnostic capacity are rarely integrated into bed and staffing models despite their practical importance for patient throughput [13, 17, 20].

The comparison with prior work highlights a translational gap between model development and routine hospital management. Many studies demonstrated retrospective predictive feasibility, but fewer linked forecasts to live operational workflows, command centre dashboards, staffing decisions, or allocation protocols [5, 20–22]. Hybrid studies that combined prediction with optimisation, balancing, or simulation appeared more implementation-oriented, but they

remained exceptions rather than the dominant evidence pattern [6, 9, 16]. Therefore, the major contribution of this review is not only identifying which models exist, but showing how far most remain from coordinated, multi-resource hospital decision support [11, 15, 18].

Recommendations

Researchers, journal editors, hospital administrators, and vendors should prioritise implementation-ready evidence, transparent reporting, and interoperable infrastructure for hospital resource allocation models. Future studies should move beyond retrospective prediction toward prospective, multicentre, temporally validated evaluations that report operational context, decision points, deployment status, and real-world effects on bed allocation, staffing alignment, diagnostic scheduling, equipment availability, and patient flow [1, 3, 10, 13, 16, 18, 20, 22, 31]. Hospitals and vendors should treat predictive analytics as an integrated decision-support capability rather than a set of isolated models, ensuring that forecasts from admission-discharge-transfer feeds, staffing rosters, diagnostic systems, equipment logs, and inventory platforms can be combined into a unified operational view [6, 9, 11, 14, 17].

Research gaps

Key research gaps include the lack of mature models that jointly optimise beds, staff, equipment, diagnostics, and patient movement; the scarcity of prospective clinical and economic evaluations; and limited attention to resilience, equity, and fairness in operational resource allocation. Existing studies usually forecast one domain at a time, such as bed occupancy, admission demand, radiology volume, blood inventory, or length of stay, while only a small subset links prediction to assignment, balancing, inventory ordering, or simulation-based operational decisions [1, 3, 6, 9, 10, 14, 16, 22]. Future research should test multi-resource forecasting systems prospectively, evaluate cost-effectiveness and workforce effects, and ensure that models remain reliable during surges, staffing shortages, and regional transfer pressures while avoiding biased allocation patterns [5, 21].

Implications

The main implication for research, operations, and policy is that predictive analytics for hospital resource allocation should be governed as a safety-critical decision-support system rather than a purely technical modelling exercise. Standardised benchmarks, common outcome definitions, temporal validation, fairness monitoring, auditability, and post-deployment evaluation are needed so that models for admission prediction, bed demand, length of stay, diagnostic volume, inventory, and patient flow can be compared and trusted across hospitals [1, 9, 7, 10, 14, 16, 20]. In operational practice, forecasts should support—not replace—human judgement, especially when allocation decisions involve emergency crowding, interhospital balancing, staffing shortages, diagnostic bottlenecks, or competing patient needs [3, 6, 13, 26, 28, 29, 31].

Conclusion

Machine learning models for hospital resource allocation increased in quantity and sophistication between 2017 and 2023. The strongest evidence was concentrated in patient throughput, emergency admission prediction, bed occupancy forecasting, and length-of-stay estimation.

Bed and throughput models are the most mature areas of the field, while staffing demand, equipment use, and diagnostic capacity remain comparatively underexplored. This imbalance matters because real hospital capacity depends on the coordination of physical beds, staff availability, diagnostic services, equipment, and patient movement.

The critical next step is implementation-focused research that tests whether predictive models improve decisions in live hospital settings. Prospective validation, external testing, workflow integration, economic evaluation, and multi-resource forecasting should become standard expectations for this field.

A coordinated push toward open data, shared benchmarks, interoperable infrastructure, and pragmatic trials will determine whether predictive analytics can fulfil its promise in hospital management. Without that shift, the field risks producing technically impressive models that remain disconnected from the operational decisions they are intended to support.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 18 Aug 2023

Revised: 14 Sep 2023

Accepted: 29 Oct 2023

Published online: 25 February 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

- 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.
<https://doi.org/10.1186/s12911-019-0742-7>.
- 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.
<https://doi.org/10.1186/s12911-022-01788-6>.
- 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>.

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.2018.2804791>.

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.2022.1123>.

Dijkstra S, Baas S, Braaksma A, Boucherie RJ. Dynamic fair balancing of COVID-19 patients over hospitals based on forecasts of bed occupancy. *Omega*. 2023;116:102801.
<https://doi.org/10.1016/j.omega.2022.102801>.

Fenn A, Davis C, Buckland DM, Kapadia N, Nichols 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.
<https://doi.org/10.1016/j.annemergmed.2021.02.014>.

Patel D, Cheetirala SN, Raut G, Tamegue J, Kia A, Glicksberg B, 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>.

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):785-802.
<https://doi.org/10.1007/s10729-022-09612-7>.

Turgeman L, May JH, Sciulli R. Insights from a machine learning model for predicting the hospital length of stay at the time of admission. *Expert Syst Appl*. 2017;78:376-85.
<https://doi.org/10.1016/j.eswa.2017.02.030>.

Harrou F, Dairi A, Kadri F, Sun Y. Effective forecasting of key features in hospital emergency department: hybrid deep learning-driven methods. *Mach Learn Appl*. 2022;7:100200.
<https://doi.org/10.1016/j.mlwa.2021.100200>.

Daghistani TA, Elshawi R, Sakr S, Ahmed AM, Al-Thwayee A, Al-Mallah MH. Predictors of in-hospital length of stay among cardiac patients: a machine learning approach. *Int J Cardiol*. 2019;288:140-7.
<https://doi.org/10.1016/j.ijcard.2019.04.067>.

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.
<https://doi.org/10.1177/1460458219892402>.

Becker AS, Erinjeri JP, Chaim J, Kastango N, Elnajjar P, Hricak H, et al. Automatic forecasting of radiology examination volume trends for optimal resource planning and allocation. *J Digit Imaging*. 2022;35(1):1-8.
<https://doi.org/10.1007/s10278-021-00551-3>.

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>.

Li N, Arnold DM, Down DG, Barty R, Blake J, Chiang F, et al. From demand forecasting to inventory ordering decisions for red blood cells through integrating machine learning, statistical modeling, and inventory optimization. *Transfusion*. 2022;62(1):87-99.
<https://doi.org/10.1111/trf.16769>.

Li N, Pham T, Cheng C, McElfresh DC, Metcalf RA, Russell WA, et al. Blood demand forecasting and supply management: an analytical assessment of key studies utilizing novel computational techniques. *Transfus Med Rev.* 2023;37(4):150768.
<https://doi.org/10.1016/j.tmr.2023.150768>.

Moreno-Fergusson ME, Guerrero Rueda WJ, Ortiz Basto GA, Arevalo Sandoval IA, Sanchez-Herrera B. Analytics and lean health care to address nurse care management challenges for inpatients in emerging economies. *J Nurs Scholarsh.* 2021;53(6):803-14.
<https://doi.org/10.1111/jnu.12700>.

Chrusciel J, Girardon F, Roquette L, Laplanche D, Duclos A, Sanchez S. The prediction of hospital length of stay using unstructured data. *BMC Med Inform Decis Mak.* 2021;21(1):351.
<https://doi.org/10.1186/s12911-021-01656-x>.

Klump M, Hintze M, Immonen M, Ródenas-Rigla F, Pilati F, Aparicio-Martínez F, et al. Artificial intelligence for hospital health care: application cases and answers to challenges in European hospitals. *Healthcare.* 2021;9(8):961.
<https://doi.org/10.3390/healthcare9080961>.

Redondo E, Nicoletta V, Bélanger V, Garcia-Sabater JP, Landa P, Maheut J, et al. A simulation model for predicting hospital occupancy for COVID-19 using archetype analysis. *Healthc Anal.* 2023;3:100197.
<https://doi.org/10.1016/j.health.2023.100197>.

Abuhay TM, Robinson S, Mamuye A, Kovalchuk SV. Machine learning integrated patient flow simulation: why and how? *J Simul.* 2023;17(5):580-93.
<https://doi.org/10.1080/17477778.2022.2091624>.

Shbool MA, Arabeyyat O, Al-Bazi A, Al-Hyari A, Salem A, Abu-Hmaid T, et al. Machine learning approaches to predict patient's length of stay in emergency department. *Appl Comput Intell Soft Comput.* 2023;2023:8063846.

Alsinglawi B, Alshari O, Alorjani M, Mubin O, Alnajjar F, Novoa M, et al. An explainable machine learning framework for lung cancer hospital length of stay prediction. *Sci Rep.* 2022;12(1):607.
<https://doi.org/10.1038/s41598-021-04623-3>.

Zeleke AJ, Palumbo P, Tubertini P, Miglio R, Chiari L. Machine learning-based prediction of hospital prolonged length of stay admission at emergency department: a gradient boosting algorithm analysis. *Front Artif Intell.* 2023;6:1179226.
<https://doi.org/10.3389/frai.2023.1179226>.

Strum RP, Mowbray FI, Zargoush M, Jones AP. Prehospital prediction of hospital admission for emergent acuity patients transported by paramedics: a population-based cohort study using machine learning. *PLoS One.* 2023;18(8):e0289429.
<https://doi.org/10.1371/journal.pone.0289429>.

Tully JL, Zhong W, Simpson S, Curran BP, Macias AA, Waterman RS, et al. Machine learning prediction models to reduce length of stay at ambulatory surgery centers through case resequencing. *J Med Syst.* 2023;47(1):71.
<https://doi.org/10.1007/s10916-023-01969-2>.

Tuominen J, Lomio F, Oksala N, Palomäki A, Peltonen J, Huttunen H, et al. Forecasting daily emergency department arrivals using high-dimensional multivariate data: a feature selection approach. *BMC Med Inform Decis Mak.* 2022;22(1):134.
<https://doi.org/10.1186/s12911-022-01904-9>.

Larburu N, Azkue L, Kerexeta J. Predicting hospital ward admission from the emergency department: a systematic review. *J Pers Med.* 2023;13(5):849.
<https://doi.org/10.3390/jpm13050849>.

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.
<https://doi.org/10.1111/1742-6723.14169>.

Raita Y, Goto T, Faridi MK, Brown DF, Camargo CA Jr, Hasegawa K. Emergency department triage prediction of clinical outcomes using machine learning models. *Crit Care*. 2019;23(1):64.
<https://doi.org/10.1186/s13054-019-2351-7>.