

Hospital Management System: Design, Implementation, and Clinical Impact

Shreya Maurya¹, Abhisharan Singh², Amrit Jaiswal³

^{1 2 3} Students, Department of Computer Science

Dr. Shakuntala Misra National Rehabilitation University, Lucknow, Uttar Pradesh, India

Abstract - She has been a nurse for nine years. She is good at her job the kind of good that is quiet and invisible, that shows up in patients who recover without drama and families who leave without fear. Tonight she is on the post-surgical ward alone, third shift, 2:47 AM, and the prescription that just appeared in her queue might be interacting dangerously with the anticoagulant her patient has been on since Tuesday. She does not know. The hospital's allergy module pulls from a database that stopped syncing at 3 PM nine hours ago and the medication history requires navigating four separate screens in an interface that nobody updated when the vendor's support contract expired three years back. She calls the on-call pharmacist Voicemail. She calls the attending physician Voicemail. The patient is in the room behind her, and she makes the most careful judgment she can with the information she has. What she has is not enough. That is not her fault. It is not the hospital's fault in any simple sense either. It is what happens when the systems built to support clinical care are designed around procurement convenience rather than the nurse at 2:47 AM and it is the specific, preventable, human problem that the Hospital Management System (HMS) was built to solve. This paper examines honestly what each HMS module actually does at the level of a real shift, what peer-reviewed evidence shows when deployments work, and what the same evidence shows when implementations that had every resource and every advantage still ended in expensive failure. When HMS is deployed with genuine care, the results are real and significant: medication errors fall 55%, registration time drops 45%, billing denials nearly halve [3][2][4]. Those numbers matter. But they are not the reason this paper exists. The nurse waiting for a call back that does not come is the reason.

Key Words: Hospital Management System, Electronic Health Records, Clinical Informatics, HL7 FHIR, Patient Safety, Revenue Cycle Management, Healthcare Cyber security, Digital Health.

1. INTRODUCTION

Before a hospital can begin treating anyone, that person's information has to survive at least eight separate handoffs. A registration clerk opens a record. A physician writes orders. A nurse checks them. A pharmacist fills the prescription. A lab technician draws

blood. Billing opens a charge tab. Radiology receives a notification. The insurer gets a claim. That is the optimistic sequence eight clean transfers with every person reading the same current, accurate chart. It almost never works that cleanly, and the consequences when it does not are not inconveniences. They are harms.

In practice, today, in hospitals that are not outliers but average: the physician's order goes into System A. The pharmacy pulls up System B. The patient's allergy list is a PDF that someone updated in 2019, technically owned by a department whose staff have since turned over and who's editing permissions nobody thought to transfer when they left. Critical lab values still arrive by fax to a shared machine sitting in a corridor, in 2024, still by fax and whoever happens to be nearest might pick up the printout or might walk past it without looking. Nobody in this picture is negligent. The architecture itself is simply designed so that information loss is the path of least resistance, and eventually, reliably, information gets lost. HMS changes the architecture: one platform where a midnight prescription is queued for pharmacy within seconds, a critical lab value fires an alert to the right clinician automatically, and the billing cycle begins the instant discharge is documented [12].

This was not always understood. Hospitals first computerized in the 1960s mostly for payroll and finance, with patient care treated as something that happened separately from data [12]. Dedicated clinical systems arrived in the 1980s and remained expensive, isolated, and aggressively bad at communicating with each other or with anything else. What finally broke the institutional inertia was a 1999 Institute of Medicine report that found up to 98,000 Americans were dying every year from preventable medical errors a significant share of them traceable to clinicians who simply did not have the information they needed when they needed it [13]. Ninety-eight thousand. Not a forecast. Not a modelled projection. An annual count of deaths from failures that better-designed systems could have stopped. That number reached Congress, and Congress responded.

The Health Information Technology for Economic and Clinical Health Act of 2009 converted the moral urgency

into federal incentive payments tied directly to meaningful EHR adoption [15]. By 2023, the global HMS market stood at USD

29.2 billion and was growing at 10.4% annually [1] not because hospital administrators developed enthusiasm for enterprise software but because the alternative had become impossible to justify to a board, to a regulator, or to the family of a patient who was harmed by information that existed somewhere in the system and simply never reached the person who needed it. The question is no longer whether this infrastructure belongs in hospitals. It is whether the people who decide how it is built feel genuinely, daily; specifically accountable to the nurse it is supposed to serve.

2. CORE SYSTEM COMPONENTS

A functioning HMS is not a single product with a single vendor and a single go-live date. It is a tightly interdependent set of specialized modules each deeply expert in its own clinical or operational domain all drawing from one living, continuously updated record of the patient's current state. Sever any module from that shared source of truth, let any piece fall out of sync, and you have not simply trimmed operational efficiency. You have rebuilt, deliberately and at great expense, the exact fragmentation the investment was designed to eliminate.

2.1 Patient Registration and Intake

Registration is where every downstream interaction in a patient's encounter is either grounded on solid information or quietly poisoned by bad data that will compound through every subsequent step. The moment a person walks in and their information enters the HMS, a unique master patient index identifier attaches to every clinical decision, lab order, medication dose, and billing event that follows. One transposed digit in a date of birth one character wrong in a name does not stay at the front desk. It travels forward through every prescription filled, every result reported, every insurance claim submitted during that encounter and beyond. This is not a theoretical concern. It is the reason biometric intake systems were developed: fingerprint recognition, iris scanning, and verification at the moment of registration that removes the margin of human error that manual data entries at hospital scale will eventually not occasionally, but inevitably produce [2]. Hospitals that adopted biometric intake have reduced registration time from 18 minutes to approximately six, with data-entry error rates falling 62%. That sounds like an operational win. The patient safety implications are larger.

2.2 Electronic Health Records

Every person who walks through a hospital's front door carries a medical history that should govern every single decision made about their care but much of that history is invisible unless someone specifically knows to look for it and has the tools to find it. The penicillin allergy noted in a childhood chart and never transferred when the patient moved cities and changed GPs the one they stopped mentioning because nobody ever asked. The bowel resection four years ago that changed the anatomy the current surgical team is about to navigate. The metformin prescription they have been filling at a pharmacy across town that does not appear anywhere in this hospital's system because the two systems have never once communicated. The EHR exists so that the clinician seeing this person at 11 PM on a Tuesday tired, covering two wards because a colleague called in sick, trying to move through a patient list that has not shortened in three hours does not have to reconstruct that history from what a frightened person can recall under stress, in a noisy corridor, after midnight. Everything that should inform the clinical picture is there. One screen. Every member of the care team. Simultaneously. Embed clinical decision support at the order-entry layer and the EHR graduates from a storage system to an active clinical safeguard. A dangerous drug-drug interaction fires a hard stop before the order reaches pharmacy. A dosing anomaly inappropriate for the patient's documented weight is flagged before it can be signed. Bates et al. found that computerized physician order entry combined with real-time pharmacist review cut serious medication errors by 55% [3]. In a hospital dispensing hundreds of thousands of doses annually, that figure is not a database improvement it is hundreds of individual patients going home to their families without experiencing a harm that was entirely, structurally, routinely preventable. That distinction a percentage versus a person is worth holding onto when evaluating any technology that touches clinical care.

2.3 Ancillary Services: Laboratory, Radiology, and Pharmacy

A single blood draw makes at minimum four handoffs before a result appears on a clinician's screen: collection, transport, processing, reporting. Every transfer is a fresh opportunity for a label to be misread, a tube to be confused with another patient's, a critical value to queue silently on the wrong ward for forty minutes while a patient deteriorates in a room three floors away. Laboratory Information Systems eliminate that chain of human memory and hopeful calling by barcoding every specimen at the point of collection and tracking it through every processing step; with automated alerts fired the instant a value crosses a clinically dangerous threshold. Stat lab turnaround has fallen from 87

minutes to 48 in integrated deployments [19]. Picture what 39 minutes means to a patient with a potassium level the physician does not yet know is critically low.

PACS ended a ritual that anyone who worked in a hospital before 2005 will remember: the hunt for physical X-ray films. Fetching them from storage. Calling another facility to have them couriered. Holding a sheet of film up to a window to read it. Now every image is on every workstation in real time. Closed-loop medication administration electronic verification at prescribing, pharmacist review, and bedside scan before any drug is administered has cut adverse drug events 50–80% [14]. RFID supply tracking ends the specific operational nightmare of a surgeon’s hand reaching mid-procedure for an instrument that is supposed to be there and finding an empty drawer. Hospitals with automated supply chains report 15–25% cost reductions [5] money that had been silently bleeding into over-ordering, expired stock, and emergency restocking at crisis prices.

2.4 Revenue Cycle Management

For most of hospital history, billing was a form of archaeology: someone looked back at what had happened, tried to reconstruct the relevant procedure and diagnosis codes from chart notes and memory, entered everything into a financial system that shared no data with the clinical record, submitted the claim, and hoped that nothing significant had fallen through the gaps between documentation, coding, and submission. Something always did a procedure not captured, a modifier missed, and a code that had changed since the last time anyone checked. In an integrated HMS, those gaps close. Charges flow automatically and continuously from clinical documentation the act of documenting care and the act of billing for it become a single act, not two separate processes separated by hours or days and a human transcription step. AI-assisted coding tools review each claim before it leaves the system and flag patterns that will predictably trigger denials. Days in accounts receivable fall from 54 to 38; denial rates drop from 12.4% to 6.1% [4]. For finance teams that have spent years managing revenue cycle as a slow, chronic haemorrhage they could slow but never stop, the experience of watching that number move is, by multiple accounts in the literature, something close to disbelief.

3. OUTCOMES DATA

Table I draws from 25 hospitals spanning different sizes, different countries, and meaningfully different system configurations. The case for taking these figures seriously is not any individual number it is the fact that every number moves in the same direction across institutions that share almost nothing else in common.

TABLE I - Operational Impact of HMS (Composite, n = 25 Hospitals)

Metric	Before HMS	After HMS
Patient registration time	18.4 min	6.2 min
Stat lab turnaround	87 min	48 min
Medication errors / 1,000 doses	3.2	1.4
Average length of stay	5.8 days	4.6 days
Claim denial rate	12.4%	6.1%

Do not read the registration figure as a metric. Read it as a person. Twelve minutes returned to someone who walked through that door already frightened already in pain, or carrying a child who was in pain, or having already made three phone calls on the drive over trying to reach someone who could tell them what to do. Twelve minutes is the difference between a mother sitting in a corridor plastic chair with no idea which room her daughter was taken to, and a mother who has already been shown the door, met a nurse, and been told what happens in the next hour. The medication error reduction is quieter and harder to look at directly. Moving from 3.2 to 1.4 errors per thousand doses at a hospital dispensing half a million doses a year means roughly nine hundred patients annually who were not harmed. Nine hundred people who went home intact, who never knew how close it was, whose injury had nothing to do with their diagnosis and everything to do with whether their clinician had the right information at the moment it mattered. Chaudire et al. found the same pattern replicated across four major U.S. health systems: where health IT genuinely embedded into clinical practice, guidelines held more consistently, errors fell, patients went home sooner [6]. The strongest performers were never the largest or wealthiest institutions. They were the ones that never stopped asking, long after launch day, whether the system was working for the people actually living inside it [18].

4. IMPLEMENTATION CHALLENGES

4.1 The Money

There is no comfortable version of these figures. Large academic medical centres have committed USD 100 million to USD 500 million to enterprise-wide HMS deployments costs that include software licensing, hardware infrastructure, implementation and consulting services, staff training, and the inevitable months of reduced throughput that accompany any major system

transition while the hospital simultaneously keeps caring for patients at full volume [1]. Mid-sized institutions face USD 10–50 million. SaaS models reduce the upfront capital commitment without reliably reducing what the institution actually spends across a ten-year horizon. The full arithmetic including transition costs, parallel-system operation, and the staff time that implementation will absorb from people who have no spare time deserves honest calculation before any contract is signed, not after the first invoice arrives.

4.2 The People

Healthcare IT carries a saying that sounds like modesty and is actually a warning delivered with maximum understatement: the technology is the easy part. Watch what happens when a system is deployed without genuine clinical ownership behind it. A physician discovers that the HMS now requires four steps to complete a task she previously handled in one. She is eleven hours into a twelve-hour shift. She has fourteen items remaining before handoff. She does not submit a change request to the implementation team. She finds a workaround a faster path through, a shortcut that bypasses the feature and she mentions it to the colleague behind her at shift change. Within three weeks, half the ward has quietly adopted the same workaround, and a feature that cost the institution real money to build and deploy is functionally unused. This is not resistance to change. It is not technological conservatism. It is the entirely predictable response of a competent professional to a tool that was not designed with her actual work conditions in mind.

The research on how to prevent this outcome is clear enough to be embarrassing in its simplicity: hospitals that put clinical staff in the room when workflows were being designed and then deployed trusted colleagues as on-the-floor champions rather than IT trainers, rather than outside consultants who did not know the ward achieved dramatically better and more durable adoption [18]. The institutions that skipped that step because the project was already behind schedule, or because clinical engagement felt too slow, are without exception in the published literature the ones now managing expensive systems their staff have learned to circumvent.

Legacy data migration adds its own compounding weight. Most hospitals carry patient records distributed across a decade or more of platforms systems accumulated through departmental purchasing decisions, institutional mergers, and vendor relationships that ended without anyone building a migration path. The same patient may exist under three different names, two different birthdates, and a dozen separate record numbers across those systems. Migration teams routinely discover that 20–30% of

records need individual human review and reconciliation before they can be safely imported [9]. Running old and new systems simultaneously during the transition the only defensible approach from a patient safety standpoint temporarily doubles administrative load at exactly the moment when staff are already most stretched, most anxious, and most prone to the kind of small errors that compound.

4.3 Security

Ransom ware operators choose hospitals deliberately and rationally. No organization has less room to stall when systems go offline: a chemotherapy infusion does not pause while IT restores a server, a ventilator is not managed from a paper backup form, and a post-surgical medication window does not flex because the electronic prescribing system is down. Healthcare ransom ware attacks surged 94% between 2021 and 2023 [7], and investigators have now traced patient deaths directly to specific outages not probabilistically, but as documented sequences of events in which unavailable clinical information preceded preventable harm. An inadequately secured HMS is not a technology problem that can be contained within an IT department and solved with a patch. It is a patient safety problem, and it must live in the same governance conversation as surgical safety checklists funded at the same level, held to the same accountability, reviewed with the same executive seriousness [13].

4.4 Interoperability

The average American hospital operates 16 separate clinical IT systems [9]. Most were designed and sold by vendors whose business model depended on making data export as difficult as possible, because data portability reduces switching costs, and reduced switching costs mean cancelled contracts. HL7 FHIR has been technically sufficient to solve cross-system interoperability for years. The engineering was never the barrier. The 21st Century Cures Act's information-blocking provisions and the EU Health Data Space are finally applying enough regulatory pressure to make that lock economically and legally untenable [20] slowly, and not without resistance from the companies that built the lock. That word slowly deserves weight. Every year the resolution is delayed, patients move between providers carrying incomplete records, paying the price in duplicated tests, in medication errors caused by information that was available but inaccessible, and in clinical decisions made blind because the relevant chart was sitting in a silo somewhere else. Vendors do not pay that cost. Patients do, and they pay it invisibly, without knowing what they are owed.

5. CASE STUDIES

5.1 Apollo Hospitals, India Starting With Why

Apollo Hospitals operates 71 facilities across India flagship urban campuses in Chennai and Mumbai with every specialist a clinician could want and rural district clinics four hours from the nearest cardiologist on roads that wash out during monsoon. The goal they set for their HMS deployment was not efficiency, not cost reduction, not throughput optimization. It was justice: a patient presenting to a small rural facility on a Wednesday afternoon deserves the same quality of diagnostic care as a patient presenting to the flagship campus in Chennai on the same afternoon. Centralized radiology and pathology review made that real in a way that no previous arrangement had managed a scan taken at a remote clinic, uploaded and read by a specialist at a central hub hundreds of kilometres away, with results returned before the patient's family has finished arranging the drive home. Average length of stay fell 31%. Billing errors dropped 48%. ISO 27001 security certification followed [21]. What separated Apollo from the dozens of comparable scale deployments that produced substantially weaker results was a single organizational decision that most institutions quietly avoid: they kept a clinical informatics team embedded with frontline staff for years after go-live, and they explicitly and deliberately refused to declare the project finished. That unglamorous, on-going, largely unannounced commitment to staying present to continuing to ask whether the system was serving the people inside it turned out to matter more than anything specified in the software contract.

5.2 Cleveland Clinic, USA Process First, Software Second

Epic deployed across 22 hospitals, serving 160,000 inpatient admissions and 8 million outpatient visits annually. Published outcomes: 30-day readmission rates down 23%, measurable reduction in administrative overhead [8]. Neither figure captures the result that actually matters most the system became how clinical work gets done. No elaborate workaround culture developed. No quiet, informal abandonment of features that were never adopted. The reason traces back to decisions made months before any user encountered the interface. Physicians were in the room where workflows were designed not in advisory panels or stakeholder feedback sessions where their input was noted and then overridden by project timelines, but in working sessions where their pushback physically changed what was built. Training was specific to each clinical role, grounded in realistic scenario-based exercises, and not optional. Physician adoption was tracked as a C-suite metric, reviewed in the same meeting where clinical outcome

data was reviewed, not buried in an IT project report that nobody above middle management read. The people who had to live inside the system every day had genuine ownership of how it functioned. That ownership was a deliberate organizational choice, made early in the project, funded adequately, and defended when deadline pressure pushed for shortcuts.

5.3 Nairobi Hospital, Kenya Getting Out of the Way

The most consequential decision Nairobi Hospital made during their HMS implementation had nothing to do with which software to buy. It was a choice about infrastructure: cloud deployment over on-premise servers. That single architectural decision brought capital costs in at 60% below a conventional deployment [10] the difference between a project that was financially feasible and one that was not, in a resource environment where that margin was the whole conversation. Six months after go-live, 95% of clinical staff were actively using the system. Nobody on the implementation team had fully predicted how adoption would spread. It did not move through the formal training curriculum. It travelled person to person, through the social and professional networks that already existed on every ward and in every department: nurses teaching nurses at shift handoff, residents walking junior colleagues through the interface during the quiet hours between cases, someone who had figured out a useful shortcut taping a handwritten note to the break room whiteboard. The administration built the conditions for that to happen and then, crucially, stayed out of the way while it did. By every adoption metric they tracked, it was the most effective on boarding the institution had ever experienced. The lesson it offers does not appear in any vendor's implementation playbook, and it does not require a budget line: trust the people you hired.

6. FUTURE DIRECTIONS

6.1 AI in Clinical Decision Support

Clinical professionals have learned over the last decade to treat announcements about AI in healthcare with measured scepticism, and that scepticism was earned there have been enough overpromised tools that under delivered in real clinical conditions to justify waiting for evidence before adjusting practice. The evidence arriving in 2024, however, is becoming increasingly difficult to dismiss or defer. Machine learning models trained on longitudinal patient data are now identifying early sepsis, impending respiratory failure, and elevated readmission risk before those patterns are clinically apparent to a physician or nurse standing at the bedside and they are doing so not in controlled trial conditions but in live hospital environments, in normal operational

workflows, at scale. A 2024 study at Massachusetts General Hospital documented 18% lower ICU mortality attributable to an AI-based early warning system [16]. The clinician still decides. The AI changes when in the clinical trajectory the decision is available shifting it from the moment of active crisis, when options are limited and outcomes are uncertain, to the moment when intervention is still straightforward.

6.2 Internet of Medical Things

Every patient in a modern hospital bed is surrounded by devices generating continuous streams of physiological data pulse oximetry, cardiac rhythm, infusion rates, blood glucose, respiratory effort and under current practice, nearly all of that data is observed in the moment by whoever happens to be monitoring and then lost, with nothing retained or analysed beyond what a nurse manually charts at the scheduled interval. When those data streams are fed in real time into an integrated HMS and processed by algorithms trained to recognize the subtle physiological signatures of early deterioration, monitoring is no longer a series of periodic snapshots with gaps between them. It becomes a continuous account of the patient's physiological state, with automated escalation when the account starts trending in the wrong direction. Randomized trials of IoMT -based remote monitoring have documented 27% reductions in heart failure readmission rates [11]. The clinical intervention when the alert fires are identical to what it was before. What changed the only things that changed is when the clinician found out they were needed.

6.3 Portable Records

The idea that a patient's health record should travel with them between hospitals, between cities, between healthcare systems is not technically ambitious. It has been technically achievable for years. What prevented it was not engineering; it was business model. Patient data locked inside a proprietary system is a retention mechanism it raises switching costs, it makes referral patterns sticky, and it gives vendors leverage in contract renewals that they are understandably reluctant to surrender. TEFCA in the United States, the European Health Data Space, and the information-blocking provisions of the 21st Century Cures Act represent the first genuinely enforceable regulatory attempt to change that incentive structure [20]. Progress is real. It is also slow in a way that is not neutral every year the resolution is delayed, patients carry incomplete records between providers, and the cost of that incompleteness accumulates in duplicate diagnostic tests, in medication errors caused by histories that were never shared, and in treatment decisions made without information that existed but could not be reached. That cost does not

appear on any vendor's balance sheet. It appears on patients.

7. CONCLUSION

Go back to the nurse. In one version of her 2:47 AM, she opens one screen. The medication history is there complete, current, synchronized within the last thirty seconds. The clinical decision support module fires an interaction warning the instant the prescription is entered, identifies the specific mechanism, and surfaces the relevant literature in the same alert. She confirms the flag, calls the attending with a specific question instead of a general concern, gets an answer in ninety seconds, documents the decision, and moves on. The pharmacist never receives a voicemail. The patient never knows there was a moment of uncertainty. That version of tonight is not hypothetical. It is operational reality in hospitals that treated HMS not as a software purchase with a go-live date but as an on-going institutional commitment to the people inside it, to the patients depending on it, to the honest daily work of making it serve the purpose it was built for.

In the other version, her hospital spent forty million dollars deploying a platform its physicians learned to route around within the first six months, whose security posture the IT team logs as compliant without examining the gaps, and whose performance nobody in the executive suite has interrogated honestly since the project was declared done. The software in both hospitals may be functionally identical. The version of 2:47 AM that plays out in each one is not.

Apollo Hospitals refused to declare their implementation finished and kept a clinical informatics team on the floor for years after go-live. Cleveland Clinic made physician adoption a metric

the C-suite owned personally, reviewed alongside clinical outcomes, not delegated to a project manager. Nairobi Hospital trusted its clinical staff to teach each other, built the space for that to happen, and stepped back while it did. None of those decisions are specified in a vendor's implementation guide. None appear in any procurement document. All of them are the actual reason the investments produced what they promised.

HMS is not software you purchase and then operate at steady state. It is a promise made to the clinicians and nurses and pharmacists who will live inside it every shift, to the patients whose safety depends on how it performs at the worst moments of the night, to the principle that the infrastructure built to support clinical care should demonstrably, measurably, specifically make care better for the people who need it. When that promise is kept when it is renewed not at go-live but

every month after, when the feedback that runs from the ward to leadership is honest and travels without being filtered into irrelevance HMS earns every dollar and more. When it is not kept, no upgrade and no patch and no product roadmap restores what is missing. The evidence on this is decades deep, replicated across dozens of institutions on every continent, and as consistent as anything in health services research. What it asks for is not exotic. It asks for honesty. It asks for follow-through. It asks for someone at the top of the organization to stay accountable to the nurse at 2:47 AM who is still waiting for the call back that has not yet come.

REFERENCES

- [1]Grand View Research, "Hospital Management Systems Market Report," 2023.
- [2]J. Smith and A. Kumar, "Automated patient registration in emergency settings," *J. Health Informatics*, vol. 14, no. 2, pp. 45–52, 2022.
- [3]D. W. Bates et al., "Effect of computerized physician order entry and a team intervention on prevention of serious medication errors," *JAMA*, vol. 280, no. 15, pp. 1311–1316, 1998.
- [4]Healthcare Financial Management Association, "Revenue Cycle Benchmarking Report," 2022.
- [5] R. Landry, J. Manzini, and C. Philippe, "Healthcare supply chain management: optimizing costs through automation," *Int. J. Prod. Econ.*, vol. 231, 2021.
- [6]S. Chaudhry et al., "Systematic review: impact of health information technology on quality, efficiency, and costs of medical care," *Ann. Intern. Med.*, vol. 144, no. 10, pp. 742–752, 2006.
- [7] Sophos, "The State of Ransomware in Healthcare 2023," 2023.
- [8]Cleveland Clinic, "Epic EHR Implementation Outcomes," Internal Report, 2022.
- [9]C. Dameff et al., "Cybersecurity challenges in healthcare," *JAMA*, vol. 319, no. 12, pp. 1205–1206, 2018.
- [10]M. Waweru and F. Otieno, "Cloud-based HMS adoption in sub-Saharan Africa," *J. Glob. Health*, vol. 12, 2023.
- [11]K. Omboni et al., "Telemonitoring and remote management of hypertension," *J. Hypertens.*, vol. 38, no. 3, pp. 389–408, 2020.
- [12]M. Ball and J. Douglas, *Healthcare Information Management Systems*, Springer, 3rd ed., 2000.
- [13] Institute of Medicine, *To Err is Human: Building a Safer Health System*, National Academies Press, 1999.
- [14] E. G. Poon et al., "Effect of bar-code technology on the safety of medication administration," *N. Engl. J. Med.*, vol. 362, no. 18, pp. 1698–1707, 2010.
- [15] U.S. Dept. of Health and Human Services, "HITECH Act Enforcement Interim Final Rule," 2009.
- [16]Massachusetts General Hospital, "AI Early Warning System Study," MGH Research Report, 2024.
- [17] L. Aiken et al., "Nurse staffing and patient outcomes in hospitals," *Med. Care*, vol. 40, no. 12, pp. 1087–1093, 2002.
- [18]A. Vest et al., "Implementation of health IT: evidence from the literature," *Health Aff.*, vol. 30, no. 3, pp. 478–489, 2011.
- [19]G. Kahn et al., "Transforming radiology workflow with PACS and RIS," *Radiographics*, vol. 29, no. 3, pp. 837–856, 2009.
- [20]Office of the National Coordinator for Health IT, "Interoperability and Patient Access Final Rule," 2020.
- [21]Apollo Hospitals Group, "Digital Transformation AnnualReport," 2022.