

ORIGINAL RESEARCH

Open access

# Agentic Artificial Intelligence System for Coordinating Hospital Discharge Tasks Using Care Plans, Medication Reconciliation Data, Transport Requests, Follow-Up Scheduling Rules, and Pending Order Status

Ethan Wright<sup>1</sup>, Chloe Bennett<sup>1\*</sup>, Jack Turner<sup>2</sup>

## Abstract

Hospital discharge delays are rooted in fragmented coordination across clinical, administrative, and logistical work. Discharge readiness depends on many interdependent tasks that must be completed in the correct sequence. No current system fully orchestrates discharge sub-tasks by reasoning across care plans, medication reconciliation, transport requests, follow-up scheduling, and pending orders in real time. Existing digital tools often support isolated functions rather than end-to-end coordination. This article proposes an agentic AI system that perceives discharge requirements, plans task sequences, executes coordination actions, and monitors completion under human supervision. The system is conceptual and should be understood as an emerging design pattern rather than an evaluated intervention. The proposed system includes a task decomposition engine, goal-oriented reasoning core, EHR data adapters, follow-up scheduling rules engine, transport request interfaces, and human override dashboard. These components would allow the agent to coordinate discharge logistics while preserving clinician authority. The system would function as a virtual discharge coordinator that continuously tracks task status and identifies stalled work. It could reduce idle coordination time, prevent overlooked tasks, and allow nurses and case managers to focus on higher-value clinical and relational work. Agentic AI could transform discharge management from a manual, interruption-prone process into an automated, accountable, and scalable workflow. Safe deployment would require interoperability, auditability, constrained autonomy, and rigorous evaluation.

**Keywords** EHR interoperability, Medication reconciliation, Care coordination, Agentic AI, Hospital discharge, Clinical operations

\*Correspondence:

Chloe Bennett

[chloe.bennett@gmail.com](mailto:chloe.bennett@gmail.com)

<sup>1</sup> Department of Healthcare AI and Informatics, Faculty of Medicine, University of Leeds, Leeds, United Kingdom

<sup>2</sup> Department of Digital Clinical Systems, Faculty of Engineering, University of Sheffield, Sheffield, United Kingdom

## Introduction

Hospital discharge is a clinically consequential and operationally complex transition that requires simultaneous attention to readiness, documentation, medications, follow-up, transport, and unresolved orders. Delays can arise when one dependent task remains incomplete, even if the patient is otherwise clinically ready to leave the hospital. Predictive work on discharge timing and discharge readiness shows that machine learning can support anticipation of discharge needs, but prediction alone does not coordinate the downstream actions needed to complete the transition [1, 2]. An agentic

discharge system would therefore need to move beyond forecasting toward structured coordination of the entire task chain [3, 4].

Current discharge coordination remains heavily dependent on nurses, case managers, pharmacists, physicians, transport teams, and scheduling staff who must communicate across fragmented systems. Manual workflows are vulnerable to omissions, duplicated effort, unclear ownership, and delays caused by unresolved handoffs. Studies on discharge prediction and discharge planning requirements indicate that workflow support must be aligned with multidisciplinary rounds, operational timing, and the lived constraints of inpatient teams [3, 5]. In this context, an agentic system would be valuable only if it complements existing roles rather than adding another monitoring burden [6].

Agentic AI introduces the possibility of systems that can interpret goals, generate plans, invoke digital tools, monitor task states, and adapt when exceptions arise. Recent discussions of autonomous agents in health care and clinical LLM agents suggest that AI systems are beginning to shift from passive information retrieval toward action-oriented support within constrained digital environments [7, 8]. Clinical workflow applications of large language models also point toward future systems that could summarize, triage, schedule, and coordinate work when embedded in safe governance structures [9, 10]. However, the discharge setting requires especially careful design because administrative actions can still affect safety, timing, and patient experience.

This article proposes an emerging agentic AI system for coordinating hospital discharge tasks using care plans, medication reconciliation data, transport requests, follow-up scheduling rules, and pending order status. The core thesis is that a supervised agent could autonomously maintain a dynamic discharge task graph, initiate permissible coordination actions, and escalate unresolved clinical dependencies to humans. Virtual EHR environments and standards-based decision support suggest that realistic testing of such agents should occur before deployment in operational care settings. The intended role is not to make discharge decisions, but to coordinate the work required after clinicians determine that discharge planning should proceed.

## Background

### The hospital discharge process and its bottlenecks

The discharge process involves a tightly coupled sequence of tasks, including confirming readiness, reconciling medications, completing orders, preparing instructions, arranging transport, and scheduling follow-up care. Bottlenecks emerge because these steps are distributed across teams and systems, and because some tasks cannot safely proceed until prior dependencies are resolved. Discharge prediction studies indicate that readiness can be anticipated, but they also show that operational action must follow prediction if the hospital is to reduce avoidable waiting [1, 2]. An agentic discharge coordinator would therefore need to model discharge as a dependency-sensitive workflow rather than a single endpoint [5, 11].

### Agentic AI and autonomous systems in healthcare

Agentic AI can be understood as a class of systems that pursue goals through planning, tool use, monitoring, and revision under defined constraints. In healthcare, such systems remain emerging because clinical environments require accountability, safety boundaries, and human supervision for actions that affect care. Surveys and conceptual analyses of AI agents in healthcare describe opportunities for autonomous workflow support, while also emphasizing the need for oversight, interoperability, and robust evaluation [7, 8]. LLM-based clinical agents further illustrate how language understanding and tool invocation could be combined, but discharge coordination would require domain-specific guardrails and operational integration [9, 10].

### Task decomposition and workflow reasoning

Task decomposition is central to agentic discharge coordination because the high-level goal of “prepare discharge” must be translated into concrete sub-tasks with owners, prerequisites, and completion conditions. A reasoning engine could represent these sub-tasks as a dynamic plan, allowing independent work such as transport preparation and appointment scheduling to proceed while gated tasks await clinician action. Prior work on clinical workflow optimization and discharge planning requirements supports the view that AI systems should assist with sequencing, prioritization, and handoff visibility rather than merely predicting discharge timing [3, 6]. In this design, the agent would maintain a structured representation of what must happen next, why it must happen, and who can authorize it.

## Data sources for discharge coordination

Discharge coordination requires data from care plans, medication reconciliation records, pending orders, transport systems, follow-up scheduling platforms, and clinical communication channels. Medication-related work is especially important because unresolved discrepancies can prevent safe discharge and require escalation to a prescriber or pharmacist. Machine learning studies on medication discrepancies and medication-error risk suggest that structured medication data can support earlier identification of reconciliation problems, but an agentic system would still need human confirmation for clinically meaningful changes [12, 13]. Interoperable medication records and EHR data standards would be necessary for the agent to retrieve current information without relying on manual transcription [14, 15].

## Human-in-the-loop models and safety in autonomous systems


Human-in-the-loop design is essential because a discharge agent would operate in a safety-critical environment where administrative coordination can still influence clinical outcomes. The agent should present its plan, rationale, evidence sources, and pending actions to clinicians, while allowing approval, modification, cancellation, or escalation at any point. End-to-end clinical AI implementation frameworks emphasize governance, monitoring, user trust, and workflow fit as prerequisites for responsible adoption [16]. Standards-based clinical decision support also provides a useful model for constraining recommendations and actions within auditable, institutionally governed rules [17, 18].

# System Architecture Overview

## High-level agent design

The proposed architecture could be centralized, with one discharge coordination agent maintaining the full task graph, or multi-agent, with specialized sub-agents for medications, transport, follow-up, and pending orders. In either configuration, the system would perceive the discharge state from clinical data sources, plan a task sequence, execute permitted coordination actions, and monitor outcomes through connected systems. Virtual EHR benchmarking environments suggest that such agents should be evaluated in simulated clinical contexts before they are trusted in live workflows [19]. The supervising clinician or case manager would remain the accountable human authority, with the agent acting as an always-on coordinator rather than an autonomous clinician.

**Figure 1** illustrates the proposed agentic AI discharge coordination architecture, showing how clinical and operational inputs are transformed into a dynamic task graph, constrained coordination actions, human-supervised escalation, and auditable discharge workflow outputs.



**Figure 1.** Agentic AI Architecture for Human-Supervised Hospital Discharge Task Coordination

## Core inputs and action space

The system's core inputs would include care plan goals, medication reconciliation status, transport requirements, follow-up rules, pending orders, and discharge documentation readiness. Its action space would be deliberately constrained to coordination tasks such as requesting transport, checking appointment availability, sending secure reminders, notifying responsible clinicians, and marking administrative sub-tasks as complete. FHIR-based interoperability and standards-based clinical decision support provide the conceptual foundation for connecting the agent to EHR data and executable policy artifacts [15, 17]. The system would not independently order medications, cancel clinical orders, or determine whether the patient is medically fit for discharge.

## Design principles

The architecture should be transparent, interruptible, auditable, workflow-compatible, and safe by design. Transparency requires that each plan step be explainable in relation to a care plan item, medication status, scheduling rule, or pending order. Interruptibility means that authorized users can pause the agent, override a decision, or redirect the plan when clinical circumstances change. These principles are consistent with implementation guidance for end-to-end clinical AI and with the need to assess clinical AI tools within governance structures that account for safety, accountability, and adoption [16, 18].

# Discharge Task Decomposition and Dependency Modeling

## Extracting discharge criteria and sub-tasks from the care plan

The agent would parse structured and unstructured care plan elements to derive required discharge tasks such as specialty follow-up, home services, equipment evaluation, pending laboratory review, patient education, and transport needs. LLM-based summarization studies show that clinical text can be transformed into structured discharge-relevant content, but coordination requires the additional step of converting that content into actionable tasks [20, 21]. The agent would map each extracted task to an owner, completion condition, relevant data source, and escalation pathway. Because extraction errors could propagate into workflow errors, every derived task should remain visible for human review before the agent executes consequential coordination actions.

## Medication reconciliation as a prerequisite

Medication reconciliation should be represented as a prerequisite for final discharge coordination because unresolved discrepancies can affect instructions, prescriptions, follow-up needs, and patient safety. The agent would compare admission medications, inpatient medication changes, and proposed discharge medication lists, then flag inconsistencies

for pharmacist or prescriber review. Predictive models for medication discrepancies and medication-error risk indicate that AI can help identify cases requiring closer attention, but such outputs should support rather than replace clinical judgment [12, 13]. A discharge agent could therefore treat medication reconciliation status as a gating signal while routing unresolved items to the responsible human.

## Modeling task dependencies and parallel execution

A dynamic dependency graph would allow the agent to distinguish tasks that can proceed in parallel from tasks that must wait for clinical completion. For example, follow-up appointment search and transport preparation could begin while the agent monitors pending orders, whereas final patient instructions may need to wait until reconciliation and discharge documentation are complete. Discharge prediction and workflow studies suggest that operational timing depends not only on readiness but also on coordination across multiple concurrent activities [3, 5, 11]. The agent would update the graph continuously as tasks complete, fail, become blocked, or require escalation.

**Table 1** clarifies how each discharge coordination input would be converted into dependency-aware tasks, constrained agent actions, and human escalation pathways within the proposed system.

**Table 1.** Agentic Discharge Coordination Logic: From Data Inputs to Task Dependencies, Permitted Actions, and Human Escalation

| Discharge Coordination Domain | Primary Data Inputs                                                                                       | Task Logic Represented by the Agent                                                                                | Dependency or Gating Function                                                                        | Permitted Agentic Coordination Actions                                                                             | Required Human Oversight or Escalation                                                               | Operational Value Added                                                               |
|-------------------------------|-----------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Care plan interpretation      | Structured care plan items; progress notes; discharge goals; specialty recommendations                    | Converts broad discharge goals into actionable sub-tasks with owners, completion conditions, and escalation routes | Determines which discharge activities are required before final coordination can proceed             | Extract task candidates; propose task list; map tasks to responsible role; display rationale                       | Human review required before consequential actions based on extracted text                           | Reduces missed discharge requirements and makes hidden care-plan dependencies visible |
| Medication reconciliation     | Admission medication list; inpatient medication changes; proposed discharge medications; pharmacist notes | Compares medication states and identifies unresolved reconciliation items                                          | Functions as a safety gate before final instructions, prescriptions, and discharge packet completion | Flag discrepancy status; notify pharmacist or prescriber; hold downstream administrative completion until reviewed | Pharmacist or prescriber must resolve medication discrepancies; agent cannot alter medication orders | Prevents unsafe downstream coordination when medication status remains unresolved     |
| Pending orders and consults   | Active orders; pending laboratory results;                                                                | Distinguishes completed, pending,                                                                                  | Blocks final discharge workflow                                                                      | Monitor order state; alert responsible                                                                             | Clinician decides whether                                                                            | Prevents premature administrative                                                     |

|                                        |                                                                                                                             |                                                                                               |                                                                                                                      |                                                                                                       |                                                                                                                      |                                                                                      |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
|                                        | imaging status; consult requests; discharge-dependent clinical tasks                                                        | clinically blocked, and newly added tasks                                                     | when unresolved clinical dependencies remain                                                                         | clinician; update task graph when status changes                                                      | pending items affect discharge readiness                                                                             | completion when clinical dependencies remain active                                  |
| Follow-up scheduling                   | Specialty follow-up rules; urgency categories; clinic availability; patient preferences; insurance or referral requirements | Identifies required appointment type, urgency, acceptable timing window, and fallback pathway | Can proceed in parallel once follow-up need is confirmed, but may require escalation if no slot exists               | Query appointment availability; prepare scheduling request; notify scheduler; flag capacity conflict  | Human scheduler or clinician resolves unavailable capacity, clinical urgency conflicts, or patient preference issues | Shortens appointment coordination delays and standardizes follow-up routing          |
| Transport coordination                 | Mobility status; destination; equipment needs; caregiver pickup plan; ambulance or wheelchair requirements                  | Converts transport needs into appropriate request category and timing estimate                | Can proceed in parallel with other administrative tasks but may be blocked by destination or readiness uncertainty   | Prepare transport request; check request status; notify transport team; update task completion state  | Human review required when transport mode, safety needs, or destination plan is uncertain                            | Reduces idle waiting after clinical readiness by preparing logistics earlier         |
| Discharge documentation readiness      | Discharge summary status; patient instructions; education materials; required forms; after-visit summary                    | Tracks whether required documents are drafted, reviewed, and ready for patient-facing use     | Final patient-facing discharge packet depends on medication reconciliation, order resolution, and clinician approval | Monitor document readiness; send reminder; display missing documentation item                         | Clinician remains responsible for content accuracy and final approval                                                | Prevents overlooked documentation gaps and improves visibility of final readiness    |
| Communication and handoff coordination | Secure messages; team assignments; role directory; escalation rules; unit workflow roles                                    | Routes the right notification to the right responsible role at the right time                 | Supports recovery from blocked tasks and failed coordination attempts                                                | Send secure reminders; notify task owner; escalate stalled dependency; document communication trigger | Human supervisor can pause, modify, or cancel notifications                                                          | Reduces duplicated calls, unclear ownership, and interruption-prone manual follow-up |

|                             |                                                                                                    |                                                                                       |                                                        |                                                                     |                                                                               |                                                                                       |
|-----------------------------|----------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|--------------------------------------------------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Audit and governance record | Action logs; data sources; rule triggers; approvals; overrides; escalations; completion timestamps | Creates traceable record of what the agent observed, planned, executed, and escalated | Applies across all agent actions as a governance layer | Log observation, rationale, action, approval, override, and outcome | Governance team reviews logs for safety, compliance, and workflow performance | Supports accountability, incident review, model monitoring, and iterative improvement |
|-----------------------------|----------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|--------------------------------------------------------|---------------------------------------------------------------------|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|

# Agentic Reasoning and Task Execution Engine

## Goal-oriented reasoning and plan generation

The reasoning engine would translate the discharge goal into an initial plan, revise the plan as new information appears, and maintain a current explanation of why each task is active, blocked, or complete. This core could combine rule-based planning for institution-specific policies with LLM-based interpretation of care plans, messages, and discharge summaries. Clinical LLM agent research suggests that tool-using models may support complex workflows, but also highlights the importance of evaluation in realistic clinical environments before deployment [9, 19]. In discharge coordination, the agent should reason over operational dependencies while leaving clinical determinations to human professionals.

## Action execution through APIs and digital interfaces

The execution layer would carry out permitted coordination actions through EHR APIs, transport request systems, scheduling platforms, and secure communication tools. FHIR-based interoperability could support retrieval of patient data, order status, medication information, and care plan elements, while standards-based decision support could encode the policies that constrain agent behavior [15, 17, 22]. The agent might request transport, check appointment availability, notify a clinician about a blocked order, or document that a nonclinical coordination task has been completed. Each action would be logged with its trigger, data source, rule basis, and human review status.

## Continuous monitoring and plan adaptation

After the initial plan is generated, the agent would monitor pending orders, transport status, appointment confirmations, medication reconciliation state, and discharge documentation readiness. If a task fails, a new order is entered, or the patient's status changes, the agent would revise the task graph and notify the appropriate human supervisor. Work on automated discharge summaries, LLM-assisted discharge documentation, and clinical workflow support suggests that AI systems can assist with information organization, but the agentic step is to connect that information to monitored actions [21, 23, 24]. Continuous monitoring would make the system a coordinator of evolving discharge work rather than a one-time checklist generator.

# Integration With Clinical Data Sources

## Care plan and pending order extraction

The discharge agent would require a reliable interface to the EHR for retrieving care plan items, active orders, pending consults, recent sign-outs, and discharge-related documentation. These data sources would allow the agent to distinguish between tasks that are already complete, tasks that remain clinically blocked, and tasks that are ready for administrative coordination. FHIR-based interoperability offers a conceptual pathway for representing patient data, orders, and care plan elements in a computable form that an agent could monitor without manual re-entry [15, 22]. Because care plans and

orders can change rapidly, the system should treat every extracted item as provisional until reconciled with the latest EHR state.

## Medication reconciliation feeds

Medication reconciliation feeds would allow the system to compare admission medication histories, inpatient medication changes, and proposed discharge medication lists. The agent could flag discrepancies, identify missing reconciliation status, and route unresolved medication issues to pharmacists or prescribers without independently altering medication orders. Work on medication-error prediction, medication discrepancy detection, and interoperable medication records suggests that structured medication data can support safer coordination when used within supervised workflows [12–14]. In this design, reconciliation status would function as a safety gate that determines whether downstream discharge instructions and coordination steps can proceed.

## Transport request and scheduling APIs

Transport and follow-up scheduling interfaces would connect the agent to operational systems outside the core clinical record. For transport, the agent could check whether the patient requires wheelchair assistance, ambulance coordination, caregiver pickup, or facility transfer support, then prepare the appropriate request for human-visible review. Research on transport decision support and machine learning for patient transportation indicates that logistics can be represented as a structured operational problem, although discharge-specific transport coordination would still require local workflow adaptation [25, 26]. For scheduling, the agent would query availability and propose appointments that satisfy clinical urgency, provider constraints, and patient-specific needs.

## Follow-up scheduling rules and guidelines

The follow-up scheduling rules engine would encode institution-specific policies, specialty pathways, urgency categories, provider availability, and patient preferences into a machine-readable policy layer. The agent could use these rules to identify the appropriate clinic type, appointment priority, and escalation pathway when no suitable slot is available. Work on post-discharge appointment adherence suggests that scheduling support should account for practical barriers and follow-up reliability rather than treating appointment booking as a purely clerical step [27]. Standards-based decision support provides a model for representing such rules as auditable artifacts that can be reviewed and updated by clinical governance teams [17, 18].

# Human-in-The-Loop Oversight and Escalation

## The supervisor console and alerts

The supervisor console would present the agent's current discharge plan, completed coordination tasks, blocked dependencies, pending human approvals, and recent alerts in a single operational view. Clinicians and case managers could approve, modify, pause, or cancel agent actions, with the interface emphasizing why each action is proposed and what data source supports it. Human-in-the-loop clinical AI frameworks emphasize that workflow integration, transparency, and clear accountability are essential for safe adoption [16]. The console would therefore be designed as a shared control surface rather than a black-box automation layer.

## Escalation triggers

Escalation triggers would define the conditions under which the agent must stop autonomous coordination and notify a human. Examples include unresolved medication discrepancies, conflicting discharge instructions, failed transport coordination, unavailable follow-up capacity, new pending clinical orders, or changes in the patient's condition. Medication reconciliation research and clinical workflow studies both support the need to distinguish routine coordination from safety-

relevant exceptions that require professional judgment [12-14]. The agent would be expected to escalate early when uncertainty affects discharge safety, documentation accuracy, or patient readiness.

## Learning from human corrections

Human overrides, edits, cancellations, and approvals would be captured as governance-relevant feedback for improving future coordination behavior. The system could learn local preferences, recurring bottlenecks, preferred escalation routes, and specialty-specific scheduling practices, while maintaining strict limits on autonomous clinical action. Clinical LLM and AI-agent research suggests that adaptive systems may become more useful when they learn from realistic workflow interactions, but such learning should remain monitored and auditable [8-10]. In this architecture, feedback would refine coordination patterns rather than permit unsupervised expansion of clinical authority.

## Safety Guardrails, Logging, and Explainability

### Operational boundaries and constrained actions

The discharge agent's operational boundaries would need to be explicit, enforceable, and visible to users. It should not order medications, determine medical readiness, alter clinical documentation, cancel clinical orders, or make independent care decisions. Its permitted actions would be limited to coordination, communication, scheduling preparation, transport request initiation, status tracking, and escalation of unresolved dependencies. This constrained design aligns with implementation guidance for end-to-end clinical AI, which emphasizes safety governance, workflow fit, and accountability before autonomous tools are embedded in care delivery [16].

### Explainable planning and reasoning traces

Each proposed plan step should include a human-readable rationale, the data source used, the rule or dependency that triggered the step, and the expected next action. For example, the agent might explain that follow-up scheduling is active because the care plan requires specialty review, while final discharge instruction preparation remains blocked because medication reconciliation is incomplete. Work on automated discharge summaries and LLM-generated clinical documentation illustrates the value of converting complex clinical information into readable summaries, but coordination requires traceability from each summary element to an operational action [20, 23, 28]. Explainability would therefore serve both usability and auditability.

### Comprehensive audit logging

The agent should maintain comprehensive audit logs for all observations, plan updates, user approvals, automated coordination actions, escalations, and overrides. These logs would allow retrospective review of whether the agent followed institutional rules, respected operational boundaries, and appropriately escalated exceptions. Studies of standards-based clinical decision support show that computable artifacts and rule-driven actions must be testable, inspectable, and governed over time [17, 18]. Audit logging would also support incident review, model monitoring, compliance assessment, and iterative improvement of the discharge coordination workflow.

## Evaluation Strategy

### Task completion and time metrics

Evaluation should examine whether the agent can correctly identify discharge sub-tasks, represent dependencies, initiate permitted coordination actions, and keep the task graph current as clinical information changes. Time-related measures could assess the interval between discharge planning milestones and task completion, but results should be interpreted as workflow indicators rather than proof of clinical benefit. Prior discharge prediction and discharge planning studies

provide useful foundations for defining readiness, workflow bottlenecks, and operational endpoints [1-3, 6]. Any evaluation should distinguish between the agent’s ability to coordinate tasks and the separate clinical judgment required to discharge the patient.

Safety and error rates

Safety evaluation should focus on incorrect task execution, missed dependencies, inappropriate escalation, failure to escalate, and conflicts between agent actions and current clinical status. The system should also be assessed for whether it preserves medication reconciliation safeguards, respects pending orders, and avoids actions outside its permitted scope. Medication safety research and automated discharge documentation studies show that AI can assist with complex clinical information processing, but also that errors in source interpretation could affect downstream workflows [13, 20, 24, 29]. For this reason, evaluation should include human review of representative agent plans, rationales, and audit trails before live deployment.

Clinical staff satisfaction and workload

Evaluation should include clinician, nurse, pharmacist, scheduler, and case manager perspectives on usability, trust, workload, interruption burden, and perceived accountability. A discharge agent could be technically capable but operationally unsuccessful if staff experience it as intrusive, opaque, or misaligned with local routines. Research on discharge planning requirements and clinical AI implementation emphasizes that adoption depends on workflow fit, perceived usefulness, and the ability to support rather than disrupt existing professional roles [6, 16]. Staff feedback should therefore shape interface design, escalation thresholds, and the boundaries of autonomous coordination.

Table 2 provides a deployment-readiness framework linking evaluation domains, measurable indicators, safety risks, governance safeguards, and implementation decisions for the proposed agentic discharge coordination system.

Table 2. Deployment Readiness and Evaluation Framework for a Human-Supervised Agentic Discharge Coordination System

| Evaluation Domain            | Core Question                                                                                                              | Suggested Measures                                                                                                                               | Failure Modes to Detect                                                                                                           | Governance or Safety Safeguard                                                                         | Evidence Needed Before Live Deployment                                              | Implementation Decision Supported                           |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------------------------|
| Task identification accuracy | Does the agent correctly identify discharge-relevant sub-tasks from care plans, orders, medications, and scheduling rules? | Sensitivity and specificity of task extraction; proportion of missed required tasks; proportion of irrelevant tasks proposed; reviewer agreement | Missing a required follow-up; creating unnecessary tasks; misreading care plan language; treating outdated information as current | Human-visible task list; source-linked rationale; review requirement before consequential coordination | Simulation testing against annotated discharge cases and expert-reviewed task lists | Whether the agent can safely generate discharge task graphs |
| Dependency modeling          | Does the agent correctly distinguish                                                                                       | Correct dependency classification;                                                                                                               | Allowing downstream coordination                                                                                                  | Dependency rules; medication reconciliation gate;                                                      | Retrospective testing using discharge                                               | Whether the agent can support real-                         |

|                              |                                                                                                    |                                                                                                                                                |                                                                                                                                      |                                                                                                        |                                                                                                        |                                                                           |
|------------------------------|----------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------|
|                              | active, blocked, parallel, and completed discharge tasks?                                          | blocked-task detection rate; graph update latency; agreement with case manager assessment                                                      | before medication reconciliation; failing to detect pending orders; incorrectly marking tasks complete                               | pending-order gate; continuous EHR state refresh                                                       | timelines and manually reconstructed dependency maps                                                   | time discharge workflow sequencing                                        |
| Constrained action execution | Does the agent remain within its permitted coordination authority?                                 | Number of actions outside scope; proportion of actions requiring reversal; successful completion of permitted tasks; API execution reliability | Ordering or altering clinical care; cancelling clinical tasks; sending inappropriate notifications; documenting beyond allowed scope | Hard-coded action boundaries; role-based permissions; pre-action policy check; automatic audit logging | Sandbox and virtual EHR testing showing no unauthorized clinical actions                               | Whether the system can be connected to operational tools                  |
| Escalation reliability       | Does the agent stop and notify humans when uncertainty, conflict, or safety relevance is detected? | Escalation sensitivity; false escalation rate; time from blocked state to alert; unresolved escalation backlog                                 | Failure to escalate medication discrepancy; alert fatigue from excessive escalation; routing alert to wrong role                     | Defined escalation triggers; severity tiers; responsible-role routing; override dashboard              | Scenario-based stress testing with medication, transport, scheduling, and pending-order exceptions     | Whether escalation thresholds are clinically and operationally acceptable |
| Human oversight usability    | Can staff understand, supervise, and override the agent without added cognitive burden?            | Usability scores; time to approve or modify plan; override frequency; perceived trust; interruption burden; role clarity                       | Staff ignore agent; unclear accountability; excessive clicks; alert fatigue; parallel manual workaround workflows                    | Supervisor console; explainable plan steps; pause/cancel/modify controls; visible accountability       | Usability testing with nurses, case managers, pharmacists, physicians, schedulers, and transport staff | Whether the agent fits local discharge routines                           |
| Coordination efficiency      | Does the agent reduce avoidable discharge                                                          | Time from discharge planning initiation to                                                                                                     | Faster but unsafe coordination; shifting                                                                                             | Balanced workflow and safety metrics; stratified monitoring                                            | Pilot evaluation in simulated or shadow                                                                | Whether operational value justifies                                       |

|                                 |                                                                                                |                                                                                                                                                     |                                                                                                                                            |                                                                                                 |                                                                                       |                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
|                                 | friction without compromising safety?                                                          | task completion; transport request lead time; follow-up scheduling completion; number of stalled tasks; discharge coordination idle time            | workload to another team; incomplete documentation; unresolved patient-facing barriers                                                     | by unit and patient complexity                                                                  | mode before active coordination                                                       | phased implementation                                                      |
| Medication safety preservation  | Does the system protect medication reconciliation as a required safety gate?                   | Unresolved discrepancies detected; reconciliation-related escalation rate; inappropriate downstream progression rate; pharmacist review concordance | Final discharge preparation before reconciliation; overlooked medication discrepancy; incorrect medication-status interpretation           | Medication reconciliation gate; pharmacist/prescriber escalation; no medication-order authority | Expert medication review of agent plans and reconciliation-triggered escalations      | Whether medication data can be safely incorporated into coordination logic |
| Auditability and accountability | Can every observation, rationale, action, approval, override, and escalation be reconstructed? | Completeness of audit logs; proportion of actions with source trace; override documentation rate; rule-trigger documentation                        | Untraceable action; unclear data source; undocumented human approval; inability to review adverse workflow event                           | Immutable logs; source references; timestamped approvals; governance review dashboard           | Audit review of simulated and pilot cases by clinical governance and compliance teams | Whether the system meets institutional accountability requirements         |
| Equity and access impact        | Does the agent coordinate discharge reliably across patient groups and discharge destinations? | Task completion by language, age, disability, destination type, insurance status, and social needs; follow-up scheduling success;                   | Unequal follow-up access; missed needs for patients with complex social barriers; biased prioritization of easier-to-coordinate discharges | Stratified monitoring; human review of high-risk discharge barriers; equity dashboard           | Pre-deployment fairness analysis and post-pilot subgroup review                       | Whether deployment may widen or reduce discharge coordination disparities  |

|                          |                                                                                              |                                                                                                                                      |                                                                                                                         |                                                                                     |                                                                                                     |                                                                            |
|--------------------------|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
|                          |                                                                                              | transport failure patterns                                                                                                           |                                                                                                                         |                                                                                     |                                                                                                     |                                                                            |
| Organizational readiness | Is the hospital environment technically and operationally prepared for agentic coordination? | API availability; data completeness; workflow mapping maturity; staff training completion; governance ownership; downtime procedures | Incomplete EHR feeds; inconsistent local rules; unclear ownership; inability to pause system during workflow disruption | Governance board; local rule validation; staged rollout; downtime and rollback plan | Readiness assessment across IT, nursing, pharmacy, case management, transport, and scheduling teams | Whether implementation should proceed, be delayed, or remain in simulation |

## Limitations

### Scope and system integration

The proposed agent is limited to administrative and logistical coordination, while all clinical decisions remain the responsibility of licensed professionals. In practice, reliable integration may be difficult because care plans, medication records, transport systems, scheduling tools, and pending order feeds may not be exposed through consistent APIs. FHIR and standards-based decision support provide useful foundations, but real-world EHR implementations may differ in data completeness, terminology, workflow configuration, and interface maturity [15, 17, 22]. These integration constraints could limit the agent’s ability to maintain an accurate real-time view of discharge readiness.

### Trust and adoption

Trust and adoption may be challenging because staff may resist an autonomous coordinator that changes familiar discharge routines or appears to redistribute responsibility. The system would need transparent reasoning, clear escalation behavior, easy override controls, and visible accountability to avoid becoming another source of cognitive burden. Clinical AI implementation frameworks emphasize that governance, stakeholder engagement, monitoring, and local adaptation are necessary for safe and sustainable adoption [16]. Therefore, the agent should be introduced gradually, with careful attention to how it affects professional roles, communication patterns, and patient-facing discharge processes.

## Conclusion

An agentic AI system for hospital discharge coordination could unify care plan interpretation, medication reconciliation status, transport logistics, follow-up scheduling, and pending order monitoring into a single supervised workflow. Rather than functioning as a passive checklist, the system would maintain a live task model that identifies what is complete, what is blocked, what can proceed, and what requires escalation.

Its main strength would be autonomous multi-task orchestration under clearly bounded authority. By adapting to new information in real time and keeping humans in control of clinical decisions, the system could support safer, more consistent, and more accountable discharge coordination.

Important challenges remain before such a system could be responsibly implemented. EHR interoperability, data reliability, staff trust, workflow alignment, explainability, and governance would all need to be addressed through rigorous design and evaluation.

Pilot implementations in high-volume medical-surgical units would be a logical next step for exploring feasibility and operational value. Such pilots should focus on whether the agent improves coordination reliability, reduces avoidable discharge friction, and supports staff without weakening clinical oversight.

## Acknowledgements

None

## Conflict of Interest

None

## Financial Support

None

## Ethics Statement

None

Received: 14 Oct 2025

Revised: 31 Oct 2025

Accepted: 08 Dec 2025

Published online: 25 February 2026

## 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

Safavi KC, Khaniyev T, Copenhaver M, Seelen M, Zenteno-Langle AC, Zanger J, et al. Development and validation of a machine learning model to aid discharge processes for inpatient surgical care. *JAMA Netw Open*. 2019;2(12):e1917221. <https://doi.org/10.1001/jamanetworkopen.2019.17221>.

McCoy TH Jr, Pellegrini AM, Perlis RH. Assessment of time-series machine learning methods for forecasting hospital discharge volume. *JAMA Netw Open*. 2018;1(7):e184087. <https://doi.org/10.1001/jamanetworkopen.2018.4087>.

Levin S, Barnes S, Toerper M, Debraine A, DeAngelo A, Hamrock E, et al. Machine-learning-based hospital discharge predictions can support multidisciplinary rounds and decrease hospital length of stay. *BMJ Innov*. 2021;7(2):e000399.

<https://doi.org/10.1136/bmjinnov-2020-000399>.

Ahn I, Gwon H, Kang H, Kim Y, Seo H, Choi H, et al. Machine learning–based hospital discharge prediction for patients with cardiovascular diseases: development and usability study. *JMIR Med Inform.* 2021;9(11):e32662.

<https://doi.org/10.2196/32662>.

Bishop JA, Javed HA, El-Bouri R, Zhu T, Taylor T, Peto T, et al. Improving patient flow during infectious disease outbreaks using machine learning for real-time prediction of patient readiness for discharge. *PLoS One.* 2021;16(11):e0260476.

<https://doi.org/10.1371/journal.pone.0260476>.

Sadel J, Grant NV, Burkhardt H, Kunze C. Optimizing hospital discharge planning: empirical insights and requirements of AI-based technologies from an explorative mixed methods field study. *JMIR Form Res.* 2026;10(1):e81824.

<https://doi.org/10.2196/81824>.

Karunanayake N. Next-generation agentic AI for transforming healthcare. *Inform Health.* 2025;2(2):73-83.

Xu G, Li X, Chen Y, Duan Y, Wu S, Yu H, et al. A comprehensive survey of AI agents in healthcare. *J Biomed Inform.* 2026;157:105045.

Mehandru N, Miao BY, Almaraz ER, Sushil M, Butte AJ, Alaa A. Evaluating large language models as agents in the clinic. *NPJ Digit Med.* 2024;7(1):84.

<https://doi.org/10.1038/s41746-024-01052-7>.

Tripathi S, Sukumaran R, Cook TS. Efficient healthcare with large language models: optimizing clinical workflow and enhancing patient care. *J Am Med Inform Assoc.* 2024;31(6):1436-40.

Wei J, Zhou J, Zhang Z, Yuan K, Gu Q, Luk A, et al. Predicting individual patient and hospital-level discharge using machine learning. *Commun Med (Lond).* 2024;4(1):236.

<https://doi.org/10.1038/s43856-024-00620-5>.

Abdo A, Gallay L, Vallecillo T, Clarenne J, Quillet P, Vuiblet V, et al. A machine learning-based clinical predictive tool to identify patients at high risk of medication errors. *Sci Rep.* 2024;14(1):32022.

<https://doi.org/10.1038/s41598-024-32022-x>.

Van De Sijpe G, Gijzen M, Van der Linden L, Strouven S, Simons E, Martens E, et al. A prediction model to identify clinically relevant medication discrepancies at the emergency department (MED-REC Predictor): development and validation study. *J Med Internet Res.* 2024;26:e55185.

<https://doi.org/10.2196/55185>.

Salgado-Baez E, Heidepriem R, Delucchi Danhier R, Rinaldi E, Ravi V, Poncette AS, et al. Toward interoperable digital medication records on Fast Healthcare Interoperability Resources: development and technical validation of a minimal core dataset. *JMIR Med Inform.* 2025;13:e64099.

<https://doi.org/10.2196/64099>.

Saripalle R, Runyan C, Russell M. Using HL7 FHIR to achieve interoperability in patient health record. *J Biomed Inform.* 2019;94:103188.

<https://doi.org/10.1016/j.jbi.2019.103188>.

Van Der Vegt AH, Scott IA, Dermawan K, Schnetler RJ, Kalke VR, Lane PJ. Implementation frameworks for end-to-end clinical AI: derivation of the SALIENT framework. *J Am Med Inform Assoc.* 2023;30(9):1503-15.

Soares A, Afshar M, Moesel C, Grasso MA, Pan E, Solomonides A, et al. Playing in the clinical decision support sandbox: tools and training for all. *JAMIA Open.* 2023;6(2):ooad038.

Isabelle M, Ip IK, Bakhtin M, Schneider L, Raja AS, Dutta S, et al. Assessment of 3 standards-based clinical decision support tools in an academic electronic health record using Clinical Quality Language, CDS Hooks, and Fast Healthcare Interoperability Resources: a retrospective evaluation. *JAMIA Open*. 2025;8(4):ooaf085.

Jiang Y, Black KC, Geng G, Park D, Zou J, Ng AY, et al. MedAgentBench: a virtual EHR environment to benchmark medical LLM agents. *NEJM AI*. 2025;2(9):Aldbp2500144.  
<https://doi.org/10.1056/Aldbp2500144>.

Ando K, Okumura T, Komachi M, Horiguchi H, Matsumoto Y. Is artificial intelligence capable of generating hospital discharge summaries from inpatient records? *PLOS Digit Health*. 2022;1(12):e0000158.  
<https://doi.org/10.1371/journal.pdig.0000158>.

Patel SB, Lam K. ChatGPT: the future of discharge summaries? *Lancet Digit Health*. 2023;5(3):e107-e108.  
[https://doi.org/10.1016/S2589-7500\(23\)00023-7](https://doi.org/10.1016/S2589-7500(23)00023-7).

Pavão J, Bastardo R, Rocha NP. A systematic review of the use of FHIR to support clinical research, public health and medical education. *Data Technol Appl*. 2025;59(1):134-54.  
<https://doi.org/10.1108/DTA-2024-XXXX>.

Williams CY, Subramanian CR, Ali SS, Apolinario M, Askin E, Barish P, et al. Physician- and large language model-generated hospital discharge summaries. *JAMA Intern Med*. 2025;185(7):818-25.  
<https://doi.org/10.1001/jamainternmed.2025.1158>.

Ganzinger M, Kunz N, Fuchs P, Lyu CK, Loos M, Dugas M, et al. Automated generation of discharge summaries: leveraging large language models with clinical data. *Sci Rep*. 2025;15(1):16466.  
<https://doi.org/10.1038/s41598-025-16466-x>.

Farhat H, Makhoul A, Gangaram P, El Aifa K, Howland I, Babay EP, Rekik F, et al. Predictive modelling of transport decisions and resources optimisation in pre-hospital setting using machine learning techniques. *PLoS One*. 2024;19(5):e0301472.  
<https://doi.org/10.1371/journal.pone.0301472>.

Hassan A, Benlamri R, Diner T, Cristofaro K, Dillistone L, Khallouki H, et al. An app for navigating patient transportation and acute stroke care in Northwestern Ontario using machine learning: retrospective study. *JMIR Form Res*. 2024;8(1):e54009.  
<https://doi.org/10.2196/54009>.

Lee SY, Eagleson RM, Hearld LR, Gibson MJ, Hearld KR, Hall AG, et al. Leveraging machine learning to enhance appointment adherence at a novel post-discharge care transition clinic. *JAMIA Open*. 2024;7(4):ooae086.

Rust P, Frings J, Meister S, Fehring L. Evaluation of a large language model to simplify discharge summaries and provide cardiological lifestyle recommendations. *Commun Med (Lond)*. 2025;5(1):208.  
<https://doi.org/10.1038/s43856-025-00824-1>.

Li W, Feng H, Hu C, Xu M, Cheng L. Accurate discharge summary generation using fine-tuned large language models with self-evaluation. *Sci Rep*. 2026;16(1).