Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

A FIRST GLIMPSE INTO HYBRID SIMULATION FOR INPATIENT HEALTHCARE ROBOTICS

Sebastian Amaya-Ceballos[1], Paula Escudero[1], Pamela Carreno-Medrano[2], and William J. Guerrero[3]

[1]School of Applied Sciences and Engineering, Universidad EAFIT, Medellin, Antioquia, COLOMBIA
[2]Dept. of Electrical and Computer Systems Engineering, Monash University, Melbourne, AUSTRALIA
[3]School of Engineering, Universidad de la Sabana, Chía, Cundinamarca, COLOMBIA

ABSTRACT

Burnout among healthcare workers in inpatient care is worsened by administrative tasks and staff shortages. Although service robots could support these environments, adoption remains limited, especially for non-clinical roles, and they lack structured methods for early-stage evaluation. This paper presents a conceptual hybrid simulation model as a foundation for pre-adoption analysis of robotic agents in inpatient care. The model integrates Discrete Event Simulation for patient flow and Agent-Based Modeling for human–robot interactions, designed using the ODD protocol and informed by expert input. It formalizes assumptions, roles, and workflow logic to support early experimentation and scenario testing. Initial simulations assess task distributions with and without robot integration, highlighting reduced documentation and assistance burdens for nurses. The conceptual model was validated through structured meetings with clinical staff and proved feasible for further development. The model includes assumptions, structure, and logic behind the simulation, serving as a basis for future model expansion.

1 INTRODUCTION

Healthcare workers (HCWs) in inpatient care settings face increasing demands due to high patient loads, a considerable administrative burden, and staff shortages. These demands contribute to higher levels of burnout, which may affect not only HCWs welfare but also the quality of patient care (Muir et al. 2022).

One of the primary contributors to burnout is the administrative burden (Herd and Moynihan 2021), which imposes an additional strain on HCWs and further exacerbates the stress they experience. The use of service robots has emerged as a potential tool to help reduce workload and improve workflow efficiency (Abdi et al. 2018). Literature has shown that there is a growing interest in the use of robotics in healthcare, driven by the decrease in costs and the reduction in human errors associated with the adoption of these technologies (Cresswell et al. 2018; Bates and Gawande 2003). Despite increasing interest in healthcare robotics, integrating robotic agents into healthcare settings remains challenging due to the complexity of healthcare environments. Small-scale demonstrations often struggle to scale up, and adoption largely depends on the acceptance of healthcare staff (Greenhalgh et al. 2017). While robotic technologies have found success in areas such as surgery, rehabilitation, and outpatient services, their application in inpatient environments remains limited (Morgan et al. 2022).

To ensure more effective deployment, healthcare institutions require tools to help evaluate how robotic agents can impact hospital operations and staff behaviour before real-world implementations. Simulation modeling offers a controlled environment to explore these impacts without disrupting ongoing care. Particularly, Discrete Event Simulation (DES) and Agent-Based Modelling (ABM) are common simulation modelling techniques in healthcare research (Mustafee et al. 2017). DES is the most preferred simulation modelling technique used in this context due to its ability to model patient flow and resource utilization (Günal and Pidd 2010). On the other hand, ABM is well suited for capturing individual behaviours and agent-environment interactions in dynamic settings (Grimm et al. 2010; Hsu et al. 2016). Hybrid approaches (such as a combination of ABM and DES) have also emerged as promising modelling tools leveraging the strengths of both techniques (Brailsford et al. 2019). While there are many simulation studies focused on hospital operations and others exploring robotic systems, few combine both in a single integrated framework, especially for inpatient settings involving collaborative human–robot work.

This paper aims to develop a conceptual simulation model to support the pre-adoption analysis of service robots in inpatient care, focusing on their role in supporting staff and improving workflow. The model integrates patient flow, healthcare staff decision-making, and system uncertainty and offers a foundation for future experimental analysis. This research represents the initial stage of a larger project aimed at determining the minimum conditions for successfully integrating robotic agents in hospital environments through a hybrid model, using as an initial case study, the hospitalization unit of Hospital Universitario La Sabana, in Colombia. This paper contributes to the field by offering a well-structured conceptual model, which formalizes and communicates the assumptions, structure, and logic behind the simulation. It supports transparency, reproducibility, and reusability, all critical aspects of rigorous simulation modeling.

The remainder of this paper is structured as follows: Section 2 reviews existing literature on simulation and robotics in healthcare. Section 3 presents the design and logic of the hybrid simulation model. Section 4 outlines initial simulation results and summarizes the expert-based face validation. Section 5 discusses the implications of this work and proposes directions for future research.

2 RELATED WORK

Simulation is one of the most common modelling techniques in healthcare, particularly for analyzing patient flow, resource allocation, and human behavior. DES and ABM are among the most commonly used techniques in this context, with DES being particularly well-suited for modelling structured workflows and queuing systems, and ABM enabling the representation of individual entities with autonomous behaviour that interact with each other and their environment (Mustafee et al. 2017; Günal and Pidd 2010; Hsu et al. 2016). These techniques are increasingly employed to address challenges in health technology assessment (HTA) (Salleh et al. 2017)

Among the many application areas of healthcare simulation, one line of work focuses on robot-enabled care. Recent studies span cost-effectiveness analyses of assistive robots for the elderly (Maresova et al. 2023), ABM evaluations of nanorobots for targeted drug delivery (Ntika et al. 2013), DES assessments of service robots and their clinical–financial impact in hospitals (Mukherjee and Sinha 2020), and hybrid ABM–data-science models that explore staff acceptance of socially assistive robots (Carlotta et al. 2018).

Recently, hybrid simulation models—particularly combinations of DES and ABM—have gained attention for their ability to combine system-level dynamics with individual behaviour representation. These models are increasingly used in complex healthcare systems, where capturing both operational workflows and human decision-making is essential (Brailsford et al. 2019; dos Santos et al. 2020). While most of the literature focuses on emergency departments or outpatient services (Viana et al. 2018; Li et al. 2020), few hybrid simulation studies have addressed inpatient care settings, and no evidence of the integration of robotic agents into inpatient workflows (Kar et al. 2022; Kittipittayakorn and Ying 2016).

Healthcare robotics research has traditionally concentrated on surgical, rehabilitative, or assistive applications, with limited focus on indirect support tasks such as transport or documentation in inpatient units. Only 3% of studies in this domain address these roles (Morgan et al. 2022). Moreover, literature shows that studies rarely combine hybrid techniques with conceptual modelling frameworks aimed at supporting pre-adoption evaluation. Conceptual models are essential for defining system boundaries, agent logic, and workflow structure, especially when the goal is to simulate socio-technical systems involving human–robot interaction. Following the need identified by Robinson (2020), this study contributes a structured conceptual model grounded in expert-informed logic to serve as the foundation for future experimentation and scenario evaluation.

3 CONCEPTUAL MODEL

Conceptual modeling is the most important stage of the simulation process. It involves developing a simplified yet representative abstraction of a system. One important issue in conceptual modelling is selecting an appropriate level of complexity; an overly detailed model can become unwieldy and lose its advantages, while an overly simplified one may yield results that fail to reflect the system's actual behavior (Robinson 2007). Therefore, conceptual modeling is a challenging process of simulation (Robinson 2020). One widely used conceptual modeling framework in ABM is the ODD (Overview, Design Concepts, and Details) protocol (Grimm et al. 2010). The conceptual model presented here is described following the ODD protocol since it includes the required level of detail to allow replicability without overexplaining the model mechanics.

This model aims to help identify key factors for robot integration, such as workflow efficiency, acceptance levels, and task delegation feasibility, along with their minimum requirements for viable adoption. In consultation with domain experts, the model's validity is assessed by its ability to replicate observed behaviors in hospital settings. Our model includes both static and dynamic entities. Static entities remain unchanged throughout the simulation and include the hospitalisation ward's layout and components (e.g, beds, nurse stations). The hospitalization floor includes 11 rooms (each holding up to four patients), a nurse station, and a medication station. All possible agent trajectories are mapped as paths within a discrete graph. Figure 1 provides a visual representation of the floor, including sample agent trajectories.

(a) Model hospitalization floor representation

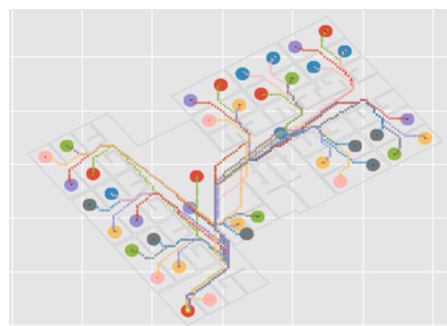

(b) Sample of possible paths in the model

Figure 1: Simulation of hospitalization floor

Dynamic agents—including doctors and nurses (referred to as Health Care Worker Agents, or HCWAs), patients, and robots—exhibit three types of interaction relationships: active (performing tasks on other agents), passive (agents on which activities are performed), and mixed (supportive agents that assist in procedures). The model integrates a patient arrival scheduler to simulate inflows. The state variables defining each agent's attributes and behavioral parameters are presented in Table 1

Human agents, including HCWAs and patient agents, are assigned a standardized set of state variables that represent their basic functions and interactions within the model. As illustrated in Figure 1, the nomenclature and values of these variables may vary by agent type. For example, patient states may include "resting" while nurse agents include "On Standby".

We assume that HCWAs work under fixed shift durations and move at a predefined walking speed. Doctor agents interact exclusively with the patient agent. They conduct periodic evaluation (e.g., patient assessment) and documentation tasks (e.g, record assessment details). Nurse agents are the most complex HCWAs in the model. They interact with patient and robot agents while executing tasks such as admission, medication administration, patient assistance, discharge, and documentation. Robot agents interact with nurses and patient agents and move at a defined walking speed. They perform two main tasks: 1 pre-filling of documentation related to the administration of medication to patients, and 2 responding to patient calls if a nurse agent is unavailable at the moment. Patient agents arrival times and length of stay are stochastically determined. They exhibit passive behaviour, undergoing all procedures without autonomous behavior, aside from requesting them at model-specified times.

Table 1: State Variables for Agents

Human agent	
Variable	**Description**
Remaining Time	Time until either the HCWAs transfers the shift or the patient agent requests discharge.
State	Represents the agent's current task or action.
Assigned Robot	Stores a reference to a robot assisting the agent and its assigned location.
Active Robot Interaction	Holds the robot currently interacting with the agent, if any.
Active Human Interaction	Stores the agent interacting with at a given moment.
Priority queue	A list of activities assigned to the agent, managed by a priority queue.
Doctor Agent	
Assigned Patients	A list of patients assigned to the doctor agent.
Transferred Workload	Indicates whether the agent has transferred their workload before ending their shift.
Path	A sequence of coordinates representing the agent's movement route.
Nurse agent	
Assigned Patients	A list of patients under the nurse agent's care.
Transferred Workload	Indicates whether the agent has transferred their workload before ending their shift.
Path	A sequence of coordinates representing the agent's movement route.
Max Robot Wait Time	The longest time a nurse agent will wait for a robot before continuing without assistance.
Current Robot Wait Time	The remaining wait time for a robot agent.
Documentation Progress	A counter tracking the number of pre-filled entries.
Next Patient for Medication	Identifies the next patient scheduled to receive medication.
Patient agent	
Personal Medication Frequency	Time interval that determines when the next medication is requested.
Personal Medication Duration	Time required for a nurse agent to prepare and administer medication to a patient.
Assigned Doctor Agent	Assigned doctor agent at admission.
Assigned Nurse Agent	Assigned nurse agent at admission.
Robot Agent	
Remaining Life	The time remaining before the robot is removed.
State	The robot's current state (e.g., idle, moving, assisting).
Active Human Interaction	The human the robot is currently interacting with.
Cooldown	Prevents immediate consecutive actions by implementing a delay.
Path	A sequence of coordinates outlining the robot's movement route.
Task	Task assigned to the robot by the control system.

3.1 Process Overview and Schedule

The proposed hybrid model integrates DES and ABM as two interacting layers. In this configuration, the DES layer maintains the ward's master timetable. It schedules patient arrivals, discharges, and procedure requests, which are assigned to HCWAs and robots. HCWAs perform both primary procedures and auxiliary tasks like documentation within stochastic time frames that reflect variability without requiring explicit modeling of each procedure. At each time step, HCWAs either process tasks, remain idle, or perform secondary actions such as moving or waiting for robotic assistance. Shift changes include a brief overlap, as outgoing staff complete ongoing tasks before leaving and transfer responsibilities.

The ABM layer governs decisions and spatial interactions among patients, HCWAs, and robots; operating at one-minute intervals, matching typical hospital procedure durations to capture macro-level operations without excessive detail. Patients are modeled as passive entities, while HCWAs and robots make localized decisions in real time. For instance, agents may skip tasks for patients about to be discharged or adapt when tasks are canceled or time is insufficient. Robots assist with medication, documentation, and patient support under a centralized allocation system that also handles cancellations. Agents in the ABM layer add flexibility that supports the use of a hybrid approach. Representing each HCWA and robot as an agent allows the model to accommodate additional behaviors and interactions in future work, such as nurse fatigue, learning effects, or adaptive robot coordination. Movement is guided by the D* algorithm, allowing agents to reroute in response to congestion. Nurses reprioritize their task queues at each step, selecting from available tasks. Decisions such as whether to involve a robot depend on its expected arrival time. These chain decisions influence overall system behavior and demonstrate how ABM enables the modeling of adaptive, interactive processes beyond the scope of schedule-based simulation alone.

Upon arrival, each patient is assigned a nurse and a doctor agent before requesting admission, which adds a task to the nurse agent's pending task queue. At this stage, both the doctor and nurse agents update their patient lists, while the patient agent updates the state variables indicating its assigned medical personnel. Following each request, the patient's state transitions to "waiting [procedure name]" (e.g., "waiting admission"). Once admitted, the patient enters a "resting" state (as well as at the end of each procedure), and the model stochastically schedules procedure requests throughout the patient's stay, storing them in the Pending Requests queue. These requests encompass medical evaluations by doctors, medication administration by nurses, and assistance from both nurses and robotic agents. When a procedure is requested, it is added to the corresponding agent's queue (Pending Tasks or Control Queue), and a subsequent request is scheduled for a later time. This cycle continues until the patient reaches the scheduled discharge time, at which point a discharge request is initiated. Upon discharge, the patient exits the model, and a new patient is scheduled for arrival. Table 2 outlines the duration and rescheduling parameters for these processes due to the lack of data, we used a triangular distribution to model task times based on expert estimates for the minimum and maximum values—a standard approach for such cases (Law 2016), we placed the mode at the midpoint between the minimum and maximum values provided by experts, centering the distribution while preserving the experts estimates.

HCWAs adhere to structured workflows, which are initialized at the beginning of each shift by the preceding set of HCWAs. This process occurs every 420 simulation minutes (ticks). Each shift starts with an initial briefing. Nurse agents subsequently perform an inventory check of the medication station before addressing requests in a priority order defined after consultation with domain experts: (1) medication administration, (2) admission and discharge processing, and (3) patient assistance and documentation. When executing a procedure, a nurse agent first navigates to the designated location. If the task involves medication administration, the agent initiates a robotic handshake to verify the robot's availability for assistance. The agent then initiates the procedure, updating its state to reflect the procedure name. For procedures involving additional agents, the patient's state is updated to "in [procedure name]" (e.g., "in admission"), and active interactions between the nurse and patient are recorded. In cases involving a robot, both the robot's and the nurse's active interaction states are updated, in addition to the robot agent's state. In contrast, doctor agents, after the initial briefing, immediately address tasks in their queue, prioritizing medical evaluations over documentation. Their state update mechanisms mirror those of nurse agents, except that robotic assistance (e.i, robotic handshake) is not incorporated into the doctor agents' workflow. Upon shift completion, any unfinished tasks are transferred to the incoming HCWA. If an HCWA is engaged in a task at shift end, they complete it before leaving the model.

Unlike HCWAs, robot agents do not maintain individual task queues. Instead, a centralized control system aggregates both medication and patient assistance requests issued by other agents and starts assigning them to available robots immediately upon simulation initialization. The system manages tasks by maintaining a record of requests and their corresponding robot assignments. For medication assistance, a nurse agent directs a robot to a specific location, and interaction is confirmed or aborted based on the robot's timely arrival. This procedure simulates pre-populated documentation by either prompting the nurse with relevant queries or automatically recording vital signs during medication administration. In patient call responses, the robot adopts a proactive role: when the call light is activated, both the nurse agent and the robot are notified. If the robot arrives first, it evaluates the urgency of the situation, escalating critical cases to the nurse agent, managing non-urgent requests autonomously, or reallocating tasks to non-specialized personnel. Situational urgency is modeled as a random variable to facilitate adaptive decision-making. In the event of a cancellation, all associated variables (including active interactions and path data) are promptly cleared to avoid conflicts and ensure model consistency. Robot agents are reset to an idle state whenever they fail to reach a patient's bed before the request is canceled.

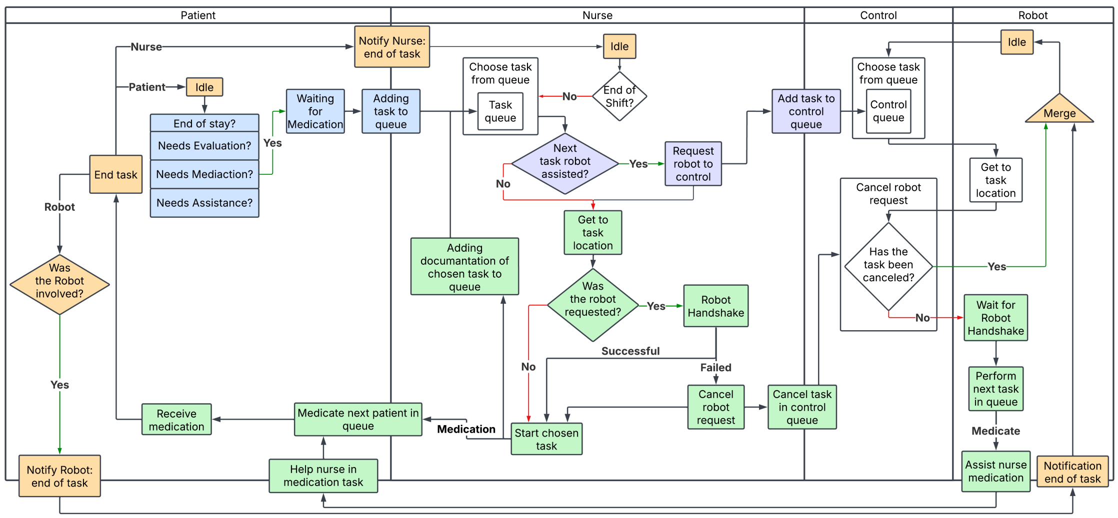

Figure 2: Medication pathway

Figure 2 depicts the medication pathway of the model, outlining the interactions and synchronization process involving the nurse, patient, and robot agents. When medication is due, the patient submits a request to the nurse agent, which adds the task to its pending queue (blue). In response to the patient's request, the nurse agent requests a robot's presence at the patient's bedside (purple). Upon completing any preceding task and arriving at the medication location, the nurse initiates a "robot handshake". If unsuccessful, the nurse proceeds without assistance, canceling the request for the robot. If successful, the nurse updates its Active Robot Interaction status to reference the robot, while the robot concurrently updates its *Active Human Interaction* status to reference the nurse. The nurse also sets the robot's state to "*helping-document*" (green). The robot pre-fills documentation fields and thus helps reduce the administrative workload for the nurse agent. Upon completing the medication process, all agents involved update their respective states (orange). A broader representation of the model, including all interactions, can be found in Amaya-Ceballos, S. (2025), while Figure 2 details this specific pathway.

3.2 Design Concepts

Our model is grounded in the principle of pre-adoption evaluation of service robots, seeking to understand their potential effects on workflow dynamics and healthcare worker burden prior to real-world implementation. One of the central design concepts in this model is emergence. Emergence refers to system-level behavior that arises from the local interactions among agents, rather than being explicitly programmed. In this context, emergent patterns are observed in how time is distributed across nurse tasks, how assistance is provided, and how resource constraints evolve under varying workloads. These outcomes are not predetermined but result from the decentralized decisions and task prioritizations made by nurse, doctor, patient, and robot agents.

Starting from this, and to maintain simplicity in this stage of the model development, decisions are dictated by agent interactions and model-defined rules (e.g., priority queues, task durations), rather than engaging in adaptive decision-making processes. For instance, upon detecting a stimulus, such as a request for a procedure or the depletion of remaining time, agents execute a predetermined response without evaluating alternatives or incorporating prior experience. The design choice emphasizes the emergent phenomena arising from straightforward interactions, ensuring both clarity in the agents' behaviors and ease of replication. Agent reactive behavior is contingent on environmental perception, making sensing a fundamental component of behavior.

While agent perception defines awareness, interaction rules define interactions. The model employs both direct and mediated agent interactions. HCWAs engage directly with patients during procedures, whereas patients initiate mediated interactions by adding requests to pending task queues, thereby prompting HCWA responses. Robots interact with nurse agents in both direct and mediated manners: they assist with documentation directly during medication processes, but in patient assistance, they signal a pre-assistance activity indirectly rather than engaging directly with nurse agents. Additionally, during patient assistance, robots interact directly with patients to provide necessary support. Mediated interactions also occur between nurse agents and robots; nurse agents add tasks to a centralized control queue accessible to all robot agents, facilitating indirect communication and coordination. Similarly, patients submit assistance requests to the same control queue, exhibiting mediated communication. Although the nature of these interactions is agent-rule dependent, their durations and frequencies are influenced by model parameters and exhibit stochastic variability.

To reflect the variability of real-life in process, the model integrates stochasticity for the timing and duration of procedures, patient arrivals, and discharges, as well as the adjusted durations of procedures involving robotic assistance. This introduces variability without explicitly modeling the underlying mechanisms, thereby capturing real-world uncertainties in patient flow and procedure execution. Additionally, medication duration and administration frequency are initially assigned randomly upon patient admission and then vary randomly around these initial values over time, reflecting individualized treatment trajectories without the complexity of a fully deterministic model.

3.3 Details

The model represents the hospitalization ward as a discrete graph, enabling agent navigation via the D* path-planning algorithm (Stentz 1994). Static agents derived from actual floor plans define distinct areas, ensuring authenticity. Staff agents—doctors, nurses, and robots—are instantiated sequentially, guided by predefined parameters (numbers of each staff type, target occupancy rate, and number of shifts), with role-specific state variables. HCWAs start at the nurse station in an idle state, bearing a 420-minute shift, no assigned patients or interactions, and an "information meeting" task. Nurses also receive a "take inventory" task. All numeric state variables initialize at zero, except for current wait time (set to the maximum robot wait), remaining stay (set to the shift length or robot lifespan), and an empty "Next Patient for Medication" field. Robotic agents begin idle with no tasks and an infinite operational lifespan. Patient arrivals are scheduled to meet the target occupancy (generally full). The model uses invariant, preloaded data rather than time-based inputs. This includes the hospital layout (a discrete graph defining all routes) and initial agent positions. Additional parameters—such as task durations, frequencies, and agent speeds—are derived from hospital records and refined through expert feedback. The model is evaluated by its primary performance measure: the average time distribution across agent states. With robot integration, a significant decrease is expected for directly supported tasks (Assisting Patient), while milder, indirect effects are anticipated for states like Documenting Records and On Standby due to shifting nurse availability and task prioritization.

Table 2: Summary of variables and time ranges

Category	Variable	Value/Range (minutes)
Patient Stay & Flow	Patient Stay Length	720 to 1440 min (12–24 hrs)
	Time Between Patient Arrivals†	120 to 300 min (2–5 hrs)
Shift Details	Shift Length*	420 min (7 hrs)
	Person Walking Speed*	15 steps/min
	Robot Walking Speed*	10 steps/min

Continued on next page

Table 2 – *Continued from previous page*

Category	Variable	Value/Range (minutes)
	Shift Transfer Meeting Time	30 to 40 min
	Inventory Time	30 to 70 min
	Documentation Time	5 to 20 min
Patient-Related Tasks	Medication Time	5 to 20 min
	Admission Time	20 to 30 min
	Evaluation Time	20 to 30 min
	Assistance Time	20 to 30 min
	Discharge Time	20 to 30 min
Task Frequencies	Evaluation Frequency	240 to 360 min (4–6 hrs)
	Medication Frequency	60 to 360 min (1–6 hrs)
	Assistance Frequency	960 to 1440 min (16–24 hrs)

*All variables use a triangular distribution in accordance with (Law 2016), except: * fixed values; † exponential distribution*

3.4 Submodels

The Fist submodel is the Robot handshake, in the medication procedure, a "robot handshake" confirms and initiates robotic support for nurse agents. When a nurse agent arrives at a patient's bedside or designated medication location, the control system checks if a robot agent is assigned. If an assigned robot is present, the handshake succeeds immediately. If the robot is assigned but has not yet arrived, the nurse waits up to the *Max Robot Wait Time*; the handshake succeeds if the robot appears within this interval. Otherwise, it fails, and the nurse proceeds unaided. If no robot is assigned, the handshake automatically fails. The *Max Robot Wait Time* is determined by empirical data and simulation requirements to ensure realistic waiting periods.

The second sub model is the Movement; HCWAs and robotic agents receive a target location for each task and check whether they are already at that location. If not, each agent queries a cache for a precomputed route; if none is found, it computes a path using the D* algorithm on a discretized floor plan, stores it in the cache, and loads it. Agents move at individual speeds (e.g., an HCWA travels 15 coordinate points per simulation tick; see Table 2) until the target is reached, at which point they transition to an idle state, awaiting subsequent tasks or operations (e.g., a robot handshake).

4 RESULTS

4.1 Validation of the Conceptual Model

The conceptual model was developed with expert input from Hospital Universitario de la Sabana to ensure an accurate representation of inpatient care workflows. Following the framework proposed by Robinson (2014), the model was assessed according to four key criteria: validity (fit for purpose), credibility (stakeholder confidence), utility (support for decision-making) and feasibility (can be coded as a computer model).

These four dimensions were assessed through a series of structured consultations involving an interdisciplinary team with expertise in robotics, healthcare, and simulation modeling. Validation included meetings with clinical experts and a dedicated session with the hospital's Head Nurse, three departmental head nurses, and two academic researchers from Universidad de La Sabana. Participants affirmed that the model captured the operational dynamics of the unit and showed potential for supporting pre-adoption evaluation of robot integration (see Table 3). The discussion also confirmed the feasibility of translating the conceptual model into a computational hybrid simulation framework. This consensus was clearly articulated by one of the consulting experts, who stated:

"As an expert in healthcare systems simulation with over six years of experience optimizing inpatient logistics in a university hospital, I verified that the discrete-event simulation model accurately reflects the service's real-world operations. The model captures the variability of activity durations, the spatial distribution of resources, and the dynamic interactions between patients and healthcare staff. It reasonably represents the complexity of the system and serves as a reliable decision-support tool for hospital management."

Table 3: Grouped quote snippets by common theme (quotes translated from Spanish)

Theme	**Snippets**
Reflection of the Workflow	"I do think it reflects the workflow; however, there are some additional variables missing..." "Answering the first question, whether it accurately reflects the workflow—of course it does..."
Institutional Variations & Staff Behavior	"Because each institution has different activities, different shift times... It's not the same if a shift is 7 to 11 hours..."
Integration with Hospital Software	"One question is whether it can be integrated so that information is uploaded directly..." "Is it easy to integrate with systems like 'Dinámica Gerencial,' 'Acción Samaritana,' etc.?"
Alleviating Nursing Workload	"If it's aimed at the nursing staff in general… it would reduce the physical and emotional burden..." "It can help with organizing inventories… generating alerts…" "Nurses move around a lot, as seen in time and motion studies…"
Usefulness for Decision-Making	"Do you consider the model useful for decision-making before implementing robots...? Yes… it can help reduce time, lessen workload, reduce accident rates..."
Cultural & Financial Aspects	"We would have to review the cultural aspect… the investment is very large… will it be administrative or handle a musculoskeletal load?"

4.2 Simulation Analysis

The model is implemented in Python using the open-source Mesa and GeoMesa libraries. Python's widespread adoption, cost-free nature, and ease of use facilitate accessibility and modifiability by hospital staff and future researchers. Mesa and GeoMesa are distributed under the MIT license and supported by active, large communities. Additionally, Python's compatibility with advanced machine learning libraries (e.g., PyTorch, TensorFlow, and JAX) will enable future extensions involving adaptive behavior in robotic agents through reinforcement learning. The model was initially executed locally with an interactive graphical user interface (GUI) (see Figure 1a). To support efficient experimentation at scale, the model was deployed on a high-performance computing cluster comprising four nodes, each equipped with dual Intel® Xeon® E5-2683 v4 processors (32 cores per node) and 64 GB of RAM, ensuring computational efficiency and scalability.

At this proof-of-concept stage, our aim was to verify that the hybrid model behaves plausibly rather than to reproduce every nuance of the real ward. We therefore ran 1000 replications of 50 shifts ($\approx$ two simulated weeks) with two doctors, four nurses and one robotic agent—enough to smooth the pie-chart distributions in Figure 3 and reveal rare events such as extended corridor blockages. The decision to use 1000 runs was guided by a 50-run pilot and Robinson's 5% half-width test (Robinson 2014): although most KPIs converged well before 1000 runs, nurse idle time would have required over 2600 runs to achieve a comparable error bound, so 1000 represented a pragmatic compromise ($\approx$ 5h on our cluster). For all subsequent scenario comparisons, we reverted to 50 runs per scenario–sufficient to assess relative differences without incurring undue computational cost.

Nurses spend considerable time walking (reflecting the decentralized nature of their tasks), followed by medication administration, admissions, and discharges, and Physicians predominantly perform diagnostic tasks. In contrast, patients experience considerable idle time, particularly while awaiting evaluations, admissions, medication, and assistance (revealing systemic inefficiencies in care delivery). The robotic agent supports nursing tasks with extended periods of movement and passive waiting, related to its limited speed and preventive design, which avoids interfering with human workflows. Another trial was conducted to assess the impact of varying the number of robots. This trial involved 50 runs for each of four distinct scenarios: a baseline with no robot, and then configurations with one, two, and three robotic agents, respectively. This trial focused on the activities involving the robotic agent, specifically medication documentation and patient assistance, as recommended by an expert. These tasks were selected to test the robot's integration, involving both passive and active interactions with other agents.

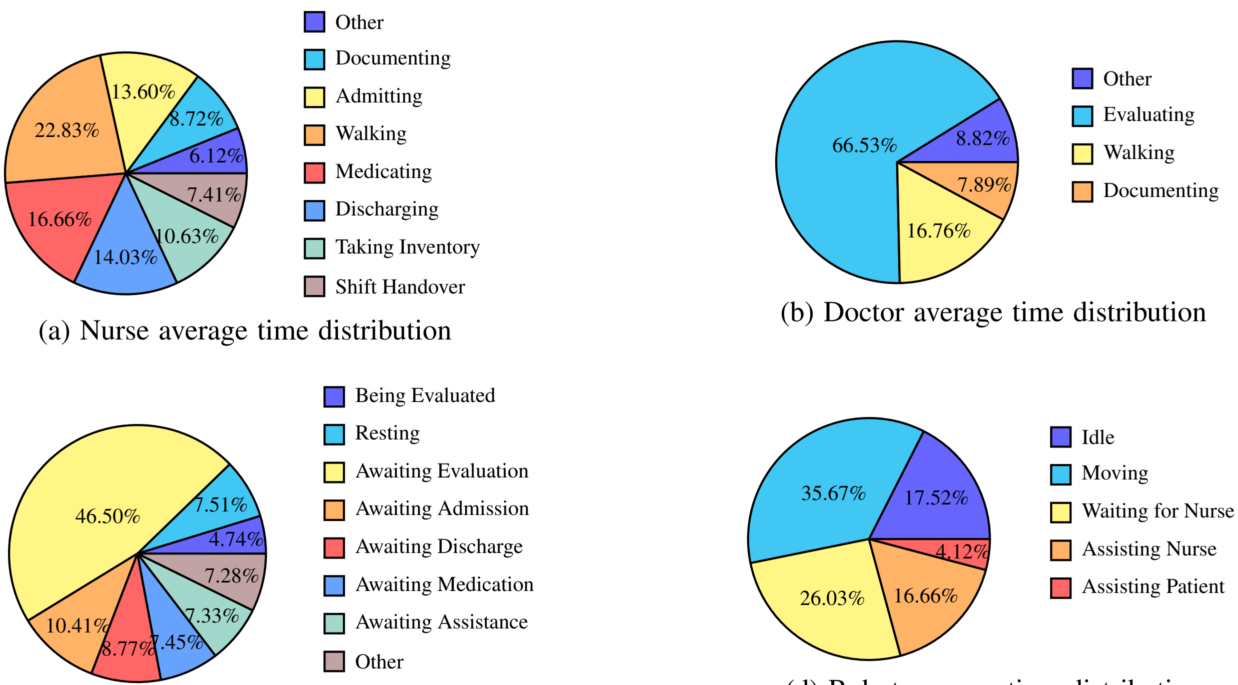

(a) Nurse average time distribution

(b) Doctor average time distribution

(c) Patient average time distribution

(d) Robot average time distribution

Figure 3: Average distribution of time across task

Figure 4 shows the confidence interval for the mean time per shift that nurses spend On Standby, Assisting Patient, and Documenting Records under four scenarios: no robot, one robot, two robots, and three robots. Introducing even a single robot cuts the time spent Assisting Patient by nearly half, and adding more robots yields further, though smaller, gains; this reduction is the only one that reaches statistical significance. At the same time, nurses' documentation workload steadily declines with each additional robot, while standby time rises slightly as robots take on more hands-on care tasks. These patterns suggest that the robot agents free nurses from direct patient assistance most effectively—because it is the only task where they have a direct impact—while their effect on composite tasks, such as record-keeping, is more modest. All results were reviewed and confirmed with domain experts.

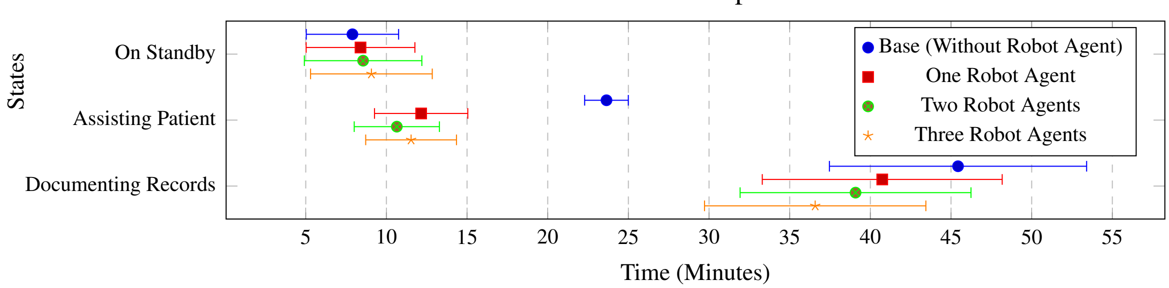

Figure 4: Mean time distribution across task types influenced by the robot agent

5 DISCUSSION & CONCLUSIONS

This study aimed to explore whether a hybrid DES–ABM model could serve as a tool to analyze the pre-adoption of service robots in inpatient care (focusing on their role in assisting staff and improving workflow) and help identify integration challenges and areas where robotic assistance can offer more benefit to healthcare workers without compromising the quality of patient care. The conceptual model presented here provides a basis for capturing interactions between agents, allowing emergent behavior to arise from human-robot collaboration. The model supports structured scenario exploration under controlled assumptions by explicitly modeling agent roles, decision logic, and task priorities.

In contrast to existing simulation studies (many of which focus on outpatient or surgical settings), this work addresses inpatient environments and proposes a hybrid DES–ABM approach grounded in expert input. Moreover, the model is implemented in Python, an open-source, general-purpose programming environment, allowing future researchers and healthcare providers to reproduce, modify, and extend the model without any usage restrictions. The result is a hybrid simulation framework that meets key conceptual modeling criteria (validity, credibility, utility, and feasibility) and can be readily adapted to other inpatient care settings due to its modular and context-independent design. This foundation makes it a valuable tool for pre-implementation exploration and decision support. However, the model is limited by its conceptual scope, serving primarily as a proof of concept rather than a comprehensive decision-making tool. It emphasizes agent interactions over hospital-level detail: patient arrivals use a simplified stream, and per-task prioritization omits system-level features like dynamic priority escalation. Agent behaviors are homogeneous, and spatial conflicts are handled via routing heuristics rather than collision-avoidance or negotiation. These simplifications reflect the early conceptual phase of the project and situate the model within an exploratory rather than predictive framework.

As this work is an early stage of a broader research project, future work will focus on translating the conceptual model into a fully operational simulation framework that reflects real-world configurations of inpatient care units. This includes refining the agent's behavior, calibration using empirical data, sensitivity analyses, and formal validation. Additional work will expand the robot agent's behavioral complexity and explore the integration of reinforcement learning techniques to simulate adaptive decision-making. In addition, the potential for applying the model in other inpatient contexts will be explored to better understand its adaptability and relevance for broader evaluations of robotic integration in healthcare.

REFERENCES

Abdi, J., A. Al-Hindawi, T. Ng, and M. P. Vizcaychipi. 2018. "Scoping Review on The Use of Socially Assistive Robot Technology in Elderly Care". *BMJ Open* 8(2).

Amaya-Ceballos, S. 2025. "Conceptual Model Representation of Inpatient Care Robotics Model". https://github.com/saamayac/conceptual_model_representation.

Bates, D. W., and A. A. Gawande. 2003. "Improving Safety with Information Technology". *New England Journal of Medicine* 348(25):2526–2534.

Brailsford, S. C., T. Eldabi, M. Kunc, N. Mustafee, and A. F. Osorio. 2019. "Hybrid Simulation Modelling in Operational Research: A State-Of-The-Art Review". *European Journal of Operational Research* 278(3):721–737.

Carlotta, B., H. Rainer, and U. Bechtold. 2018. "Socially Assistive Robots Diffusion in Elderly Care: A Pre-Adoption Study Through Agent-Based Modeling". In *2018 Portland International Conference on Management of Engineering and Technology (PICMET)*, 1–10.

Cresswell, K., S. Cunningham-Burley, and S. A. Sheikh. 2018, 03. "Healthcare Robotics a Qualitative Exploration of Key Challenges and Future Directions". *Journal of Medical Internet Research* 20.

dos Santos, V. H., K. Kotiadis, and M. P. Scaparra. 2020. "A Review of Hybrid Simulation in Healthcare". In *2020 Winter Simulation Conference (WSC)*, 1004–1015 https://doi.org/10.1109/WSC48552.2020.9383913.

Greenhalgh, T., J. Wherton, C. Papoutsi, J. Lynch, G. Hughes, C. A'Court, *et al*. 2017. "Beyond Adoption: A New Framework for Theorizing and Evaluating Non adoption, Abandonment, And Challenges to The Scale-Up, Spread, And Sustainability of Health and Care Technologies". *Journal of Medical Internet Research* 19(11):e367.

Grimm, V., U. Berger, D. L. DeAngelis, J. G. Polhill, J. Giske, and S. F. Railsback. 2010. "The ODD Protocol: A Review and First Update". *Ecological Modelling* 221(23):2760–2768.

Günal, M. M., and M. Pidd. 2010. "Discrete Event Simulation for Performance Modelling In Health Care: A Review Of The Literature". *Journal of Simulation* 4(1):42–51.

Herd, P., and D. Moynihan. 2021. "Health Care Administrative Burdens: Centering Patient Experiences". *Health Services Research* 56(5):751–754.

Hsu, S.-C., K.-W. Weng, Q. Cui, and W. Rand. 2016. "Understanding The Complexity of Project Team Member Selection Through Agent-Based Modeling". *International Journal of Project Management* 34(1):82 – 93.

Kar, E., T. Eldabi, and M. Fakhimi. 2022. "Hybrid Simulation in Healthcare: A Review of the Literature". In *2022 Winter Simulation Conference (WSC)*, 1211–1222 https://doi.org/10.1109/WSC57314.2022.10015418.

Kittipittayakorn, C., and K.-C. Ying. 2016. "Using The Integration of Discrete Event and Agent-Based Simulation to Enhance Outpatient Service Quality in An Orthopedic Department". *Journal of Healthcare Engineering* 2016(1):4189206.

Law, A. M. 2016. "A Tutorial on How to Select Simulation Input Probability Distributions". In *2016 Winter Simulation Conference (WSC)*, 103–117 https://doi.org/10.1109/WSC.2016.7822083.

Li, Y., Y. Zhang, and L. Cao. 2020. "Evaluation And Selection of Hospital Layout Based on an Integrated Simulation Method". In *2020 Winter Simulation Conference (WSC)*, 2560–2568 https://doi.org/10.1109/WSC48552.2020.9384091.

Maresova, P., L. Režný, P. Bauer, O. Fadeyi, O. Eniayewu, S. Barakovic *et al*. 2023. "An Effectiveness and Cost-Estimation Model for Deploying Assistive Technology Solutions in Elderly Care". *International Journal of Healthcare Management* 16(4):588–603.

Morgan, A. A., J. Abdi, M. A. Q. Syed, G. E. Kohen, P. Barlow, and M. P. Vizcaychipi. 2022, Dec. "Robots in Healthcare: A Scoping Review". *Current Robotics Reports* 3(4):271–280.

Muir, K. J., J. Webb-Jones, N. Farish, K. Barker, C. Miller-Davis, and S. Galloway. 2022. ""Room To Reflect": A Pilot Workplace Resiliency Intervention for Nurses". *International Journal of Environmental Research and Public Health* 19(12).

Mukherjee, U. K., and K. K. Sinha. 2020. "Robot-Assisted Surgical Care Delivery at a Hospital: Policies for Maximizing Clinical Outcome Benefits and Minimizing Costs". *Journal of Operations Management* 66(1-2):227–256.

Mustafee, N., S. Brailsford, A. Djanatliev, T. Eldabi, M. Kunc, and A. Tolk. 2017. "Purpose And Benefits of Hybrid Simulation: Contributing to The Convergence of Its Definition". In *2017 Winter Simulation Conference (WSC)*, 1631–1645 https://doi.org/10.1109/WSC.2017.8247903.

Ntika, M., P. Kefalas, and I. Stamatopoulou. 2013. "Multi-Agent System Simulation of Nano-Robotic Drug Delivery in Tumours of Body Tissues". In *2013 17th International Conference on System Theory, Control and Computing (ICSTCC)*, 777–782.

Robinson, S. 2007. "The Future's Bright The Future's... Conceptual Modelling For Simulation!". *Journal of Simulation* 1(3):149–152.

Robinson, S. 2014. *Simulation: The Practice of Model Development and Use*. Bloomsbury Publishing.

Robinson, S. 2020. "Conceptual Modelling for Simulation: Progress and Grand Challenges". *Journal of Simulation* 14(1):1–20.

Salleh, S., P. Thokala, A. Brennan, R. Hughes, and A. Booth. 2017. "Simulation Modelling in Healthcare: An Umbrella Review of Systematic Literature Reviews". *PharmacoEconomics* 35(9):937–949.

Stentz, A. 1994. "Optimal And Efficient Path Planning for Partially-Known Environments". In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, 3310–3317 vol.4.

Viana, J., T. B. Simonsen, F. A. Dahl, and K. Flo. 2018. "A Hybrid Discrete Event Agent Based Overdue Pregnancy Outpatient Clinic Simulation Model". In *2018 Winter Simulation Conference (WSC)*, 1488–1499 https://doi.org/10.1109/WSC.2018.8632282.

AUTHOR BIOGRAPHIES

SEBASTIAN AMAYA CEBALLOS is a Research Assistant of the Apollo Scientific Computing Centre at Universidad Eafit, Colombia. His current research primarily focuses on simulation modeling of human-robot interactions in healthcare settings. His email address is saamayac@eafit.edu.co

PAULA ESCUDERO is a Full Professor and Director of the research group in Analytics and Supply Chain of the School of Applied Sciences and Engineering at Universidad Eafit, Colombia. Her research is primarily focused on Agent-based and Discrete Event Simulation, with a particular emphasis on healthcare applications such as epidemiology and healthcare management. Her email address is pescuder@eafit.edu.co.

PAMELA CARREÑO MEDRANO is a Lecturer in the Department of Electrical and Computer Systems Engineering at Monash University, Australia. Her expertise lies at the intersection of human-robot interaction, computational modelling of human decision-making, and robot learning. Her current research is focused on making robots more intelligent and able to interact with humans in different situations. Her email address is pamela.carreno@monash.edu.

WILLIAM J. GUERRERO is an Associate Professor in the School of Engineering at Universidad de La Sabana, Colombia. His expertise is the design of innovative operations research models for healthcare systems logistics. His research is focused on designing mathematical models and metaheuristics. His email address is william.guerrero@unisabana.edu.co.