

CAR-EM: A Synthesis-Based Clinically Assistive Robot System for Emergency Medicine

SANDHYA JAYARAMAN, Computer Science and Engineering, University of California San Diego, La Jolla, California, USA

ANDREW VIOLETTE, Cornell University, Ithaca, New York, USA

U LAM LOU, Computer Science and Engineering, University of California San Diego, La Jolla, California, USA

SRUTI MANI, Department of Mathematics, University of California San Diego, La Jolla, California, USA

DIVYA PRAKASH, Computer Science and Public Health, Austin College, Sherman, Texas, USA

LESLIE OYAMA and **CHRISTOPHER COYNE**, Department of Emergency Medicine, University of California San Diego, La Jolla, California, USA

HADAS KRESS-GAZIT, Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York, USA

LAUREL RIEK, Computer Science and Engineering, University of California San Diego, La Jolla, California, USA

Emergency departments (EDs) are fast-paced, dynamic, safety-critical spaces where clinicians are overworked and underpaid. To support clinicians, researchers are exploring the contextualization and development of clinically assistive robots (CARs) that can assume non-critical tasks to reduce clinician overload. In this article, we introduce Clinically Assistive Robot System for Emergency Medicine (CAR-EM), collaboratively developed with ED clinicians. CAR-EM includes an autonomous robot and a task specification interface. It completes tasks by leveraging control synthesis, a framework that automatically transforms high-level tasks into control while providing guarantees and feedback. We conducted a feasibility study across two different hospital EDs, where interprofessional clinicians tasked the robot to perform patient assessments and item deliveries. Clinicians found the system easy to use, and particularly helpful to offload busywork. This work demonstrates control synthesis as a feasible tool to develop autonomy for robots in safety-critical spaces, and identifies considerations for failure interventions. We also discuss ethical considerations for deploying robots in hospitals, including healthcare worker displacement and work disruption. Thus, our work: (1) highlights the unique requirements of situating robots in real world hospital EDs, and (2) demonstrates a novel approach

Authors' Contact Information: Sandhya Jayaraman (corresponding author), Computer Science and Engineering, University of California San Diego, La Jolla, California, USA; e-mail: sjayaraman@ucsd.edu; Andrew Violette, Cornell University, Ithaca, New York, USA; e-mail: amv78@cornell.edu; U Lam Lou, Computer Science and Engineering, University of California San Diego, La Jolla, California, USA; e-mail: ulou@ucsd.edu; Sruti Mani, Department of Mathematics, University of California San Diego, La Jolla, California, USA; e-mail: s1mani@ucsd.edu; Divya Prakash, Computer Science and Public Health, Austin College, Sherman, Texas, USA; e-mail: dprakash21@austincollege.edu; Leslie Oyama, Department of Emergency Medicine, University of California San Diego, La Jolla, California, USA; e-mail: loyama@health.ucsd.edu; Christopher Coyne, Department of Emergency Medicine, University of California San Diego, La Jolla, California, USA; e-mail: cjcoyne@health.ucsd.edu; Hadas Kress-Gazit, Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York, USA; e-mail: hadaskg@cornell.edu; Laurel Riek, Computer Science and Engineering, University of California San Diego, La Jolla, California, USA; e-mail: lriek@ucsd.edu.



This work is licensed under Creative Commons Attribution International 4.0.

© 2026 Copyright held by the owner/author(s).

ACM 2573-9522/2026/4-ART63

<https://doi.org/10.1145/3797263>

leveraging guarantees and feedback from control synthesis methods to successfully implement context-specific CAR behaviors. Through this work, we aim to further research for safer and more reliable robots in real world, uncertain environments.

CCS Concepts: • **Computer systems organization** → **Robotics**; • **Human-centered computing** → **Human computer interaction (HCI)**;

Additional Key Words and Phrases: Human–Robot Interaction, Clinically Assistive Robots, Acute Care, Formal Methods, Control Synthesis

ACM Reference format:

Sandhya Jayaraman, Andrew Violette, U Lam Lou, Sruti Mani, Divya Prakash, Leslie Oyama, Christopher Coyne, Hadas Kress-Gazit, and Laurel Riek. 2026. CAR-EM: A Synthesis-Based Clinically Assistive Robot System for Emergency Medicine. *ACM Trans. Hum.-Robot Interact.* 15, 3, Article 63 (April 2026), 22 pages. <https://doi.org/10.1145/3797263>

1 Introduction

Emergency departments (EDs) serve as the frontline of healthcare for many patients presenting with a range of symptoms from minor (e.g., a common cold) to severe (e.g., gunshot wounds) [2]. The resulting volume and acuity of patients overwhelms existing staffing models in the ED [118]. A single clinician is often responsible for multiple tasks, ranging from item deliveries to performing complex diagnoses and procedures [49], while concurrently coordinating with external entities such as insurance companies and law enforcement [32]. As a result, ED clinicians report among the highest levels of burnout of all medical specialties [15, 118].

While only radical systemic changes to hospital operations and staffing will solve these fundamental problems [51], researchers in **Human–Robot Interaction (HRI)** are exploring the use of robots to lighten the load on overburdened ED clinicians [2, 52, 60, 86]. Both prior work [101, 102], and our own long-standing collaborations with clinical collaborators, suggest potential in offloading simple tasks to autonomous robots. For example, researchers have explored robots that support supply delivery [5, 28, 53, 93], room sanitation [1], and promote clinician teamwork to reduce preventable patient harm [101, 104, 105]. However, such robots need to be carefully contextualized in a manner that does not devalue the human elements of care by positioning as cheaper or more “efficient” sources of labor [12, 110]. These robots must not disrupt clinical work, add to clinicians’ cognitive load, or negatively impact patient care or satisfaction.

Clinically Assistive Robots (CARs) must: (1) perform tasks successfully, (2) interact appropriately with different populations (e.g., clinicians, patients, and others) and (3) communicate their status to relevant users. Further, there is a need for guarantees that ensure such robots will not exhibit undesirable behavior (e.g., entering patients’ rooms without permission), particularly due to the safety-critical nature of the ED.

To be beneficial for clinicians, CARs need to be “invisible in use” [100]. That is, they need to be contextualized to weave into the fabric of the ED without being an additional burden. Clinicians need to be apprised of the robot’s status, and by extension the status of their tasks, in a manner that does not disrupt their workflow. In case the robot fails to complete its task, clinicians need to be provided with adequate feedback to inform their next course of action.

While prior robots deployed in hospitals show that CARs have the potential to support clinicians [5, 89], there is limited research on how they can ensure safety and provide feedback to clinicians—a key requirement for safety-critical spaces. Also, to our knowledge, few researchers have evaluated autonomous mobile robots in real EDs [60]. Finally, most existing CARs only perform a single task

type (e.g., nonpatient-facing deliveries [89], companionship [92]), which limits their utility in the ED [107].

To address these gaps, we present **CAR System for Emergency Medicine (CAR-EM)**, a clinically assistive robot system developed alongside clinicians over the course of several years. CAR-EM consists of a task specification interface for clinicians, an autonomous mobile robot with an interaction tablet (named “Min” by one of our clinical participants), and a synthesis framework that transforms specifications given in temporal logic into Min’s control. We collaborated with clinicians to identify six tasks appropriate for CARs to perform within the ED. We selected these based on clinicians’ needs for the robot to take on tedious tasks that added to their work overload.

We automatically synthesize Min’s control to autonomously perform these tasks in the ED. Control synthesis is a set of techniques that take a high-level specification of robot behavior, typically written in temporal logic [56], and automatically synthesize a correct-by-construction controller that is guaranteed, under some assumptions, to satisfy the task. By leveraging these techniques, we allow users to reason about the task at a high-level rather than at the implementation level, and we can provide some guarantees on the robot’s behavior, which is valuable in any safety-critical space. When unexpected events occur, the robot might not be able to complete its high-level specification. Using the synthesis framework, we can reason and provide information to the clinician about why the robot failed.

To study CAR-EM in a real-world context, we conducted a feasibility study across two different EDs with 11 clinician participants (3 physicians, 6 nurses, and 2 ED techs). Overall, clinicians found the task specification interface easy to use, and the tasks Min performed useful. We also identified key considerations in robot feedback delivery, robot task appropriateness and ethical concerns for deploying CARs in hospitals.

Contributions and Novelty. The contributions of this work are three-fold. First, we demonstrate CAR-EM, to our knowledge the first synthesis-based, autonomous CAR system. Second, we present new knowledge about how to contextualize, design, implement, and deploy CAR systems in real-world EDs. Finally, we present insights and challenges from key stakeholders to support robotics researchers developing CARs.

CAR-EM represents a novel systems approach as it uses formal methods to represent clinicians’ desired tasks while meeting their requirements for feedback and guarantees. To our knowledge, CAR-EM is the first CAR to have explored implicit controller-level safety guarantees. We are also, to our knowledge, the first HRI research study evaluating CAR system deployments across multiple EDs. CAR-EM was developed over a yearslong collaboration with ED clinicians, ensuring that clinicians’ voices were included in every step of system ideation, design, implementation, and testing. Our findings indicate that control synthesis can pave the way for improved situational awareness for users, safer operation of robots, and more effective integration of CARs in real world settings.

2 Related Work

2.1 Robots in Hospitals

Many researchers have designed and deployed robots in acute care settings [86, 97]. For example, the Tug robot [14] and Moxi [108] both perform item deliveries (e.g., linens, meals, and lab samples). The COVID-19 pandemic saw a surge in the development of robots to support patient triage, screenings, infection control, item delivery, companionship, and telemedicine [48, 70, 90, 96, 97]. Researchers have also explored telepresence robots for ED settings [44, 70, 71], which have minimal autonomy and need a full-time robot operator. Other work has explored what it means to situate robots [102, 109] in the ED and socially navigate within it [106]. However, to our knowledge, few

autonomous mobile robots have been deployed in the ED aside from triage and item delivery robots [28, 96].

2.2 Formal Methods in HRI

HRI researchers have explored the use of formal methods techniques such as verification and synthesis to develop safe and correct robot behaviors [54, 80]. Particularly, researchers have used specifications and synthesis to facilitate end user programming [59, 80–84]. End user programming represents methods of democratizing programming for novice or non-programmer end users [26, 83]. HRI researchers have used program synthesis to create robot programs through visual, aural, tangible, and tactile interfaces [22, 42, 76, 79, 94, 95]. However, limited research explores how safety guarantees and feedback through formal methods can be used to address user requirements in real-world applications.

2.3 Control Synthesis

Control synthesis are methods of transforming high-level specifications of robot behavior into control [56]. High-level specifications are typically written using temporal logics [23] that capture required robot behavior. Our work builds on reactive synthesis from **Linear Temporal Logic (LTL)** [55, 57, 59, 116] and **Signal Temporal Logic (STL)** specifications [38, 39, 64, 65].

LTL formulas are composed of Boolean propositions, temporal, and logical operators as follows [24]:

$$\varphi ::= \pi \mid \neg\varphi \mid \varphi \vee \varphi \mid \circ\varphi \mid \varphi \mathcal{U} \varphi,$$

where “and” ($\wedge$) and “implies” ($\rightarrow$) can be constructed from “not” ($\neg$) and “or” ($\vee$), and “eventually” ($\diamond$) and “always” ($\Box$) from “until” ($\mathcal{U}$). LTL propositions are suitable to represent robot interactions at a high level.

STL is defined over continuous signals and predicate functions and thus captures physical safety well. An STL formula is defined as [66]:

$$\varphi ::= \text{True} \mid \mu \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \mathcal{F}_{[a,b]}\varphi \mid \mathcal{G}_{[a,b]}\varphi \mid \varphi_1 \mathcal{U}_{[a,b]}\varphi_2,$$

where φ is an STL formula, μ is a predicate over a continuous signal, $\mathcal{F}$ is “eventually”, $\mathcal{G}$ is “always”, $\mathcal{U}$ is “until”, and $a, b \in \mathbb{R}^+$ represent timing bounds. STL predicates are continuous functions and are well suited to represent continuous planning tasks such as navigation. CAR-EM achieves safety guarantees and provides feedback through a combination of LTL and STL synthesis.

3 System Overview

3.1 Design Requirements

Over the past several years, we have been collaborating with clinicians to identify contexts where CARs may support clinicians in the ED. We engaged in a series of collaborative design activities, including interviews and iterative low to medium fidelity prototyping, yielding four key design requirements for CAR-EM.

First, to contextualize CAR-EM’s tasks, we made an initial catalog of tasks and requirements that over 10 clinicians from earlier collaborations envisioned for future CARs. We analyzed excerpts of conversations where clinicians described the types of support they wished for. For example, one collaborator wished for a “to-do list robot” where the clinician provided the robot a list of tasks every once in a while. The robot would execute the tasks, and report back when done.

We grouped the tasks clinicians suggested into three task types: (1) patient assessments (vitals check, pain check, telemedical visit), (2) patient deliveries (food/necessities), and (3) pick-up (lab samples, procedure kits). We verified with two ED physician collaborators that these tasks and

requirements sufficiently encapsulated how clinicians envisioned CARs. The clinicians emphasized they wanted a low-effort system that provided clear, and easily accessible feedback. One clinician exemplified how Martti [67] (an ED telemedical system that provides clinicians remote access to medically trained interpreters) seldom showed up on time when they requested it, as the staff fetching the system were interrupted often. They wanted to know where the Martti was, how much more time it was going to take, and sometimes even the details of the hold-up.

Based on these insights, we designed low-fidelity sketches of the task specification interface, iterating over regular check-ins with the two clinicians. Once the design of the task specification interface was more concrete, we created Figma [33] prototypes of the interface end-to-end that the two clinicians regularly tested. We had a note taker in each session that clinicians tested the interface, and analyzed the notes to make key changes to the interface.

In one such iteration, one of the clinicians felt “pick-up” tasks may displace other ED clinicians (such as ED techs) whose responsibilities included picking up lab samples and setting up kits for procedures such as sutures. To avoid displacing labor, and disrupting clinical responsibilities and workflows, we removed pick-up tasks. Similarly, one clinician suggested splitting up the delivery of food/necessities into solids, liquids, and comfort items. This was to help clinicians keep track of patients’ intake, and prevent patients from receiving food/drinks when they were not allowed to. These insights revealed the importance of well-contextualized tasks.

Finally, we piloted the task specification interface with three clinicians (two physicians and one ED tech). The clinicians watched a tutorial video, and then created two tasks for the robot. After testing the interface, they answered a semi-structured interview that explored (1) the appropriateness of the system and tasks, (2) their impressions of the interface and its usability, (3) clinician–robot interactions and workflows, and (4) privacy concerns and failure management. Based on their insights, we formalized design requirements for CAR-EM as described below.

Careful Contextualization of Robot Tasks. The ED is part of a large, complex ecosystem of care delivery, and introducing new technologies into it is non-trivial [4, 102]. Every hospital and ED is different in terms of its policy, protocol, and built environment, and robot tasks and behaviors need to be adaptable to such differences [37, 103]. However well-envisioned, any task a robot performs will likely cause work reconfiguration [3, 34, 68]. It is critical to ensure the robot does not disrupt existing clinical workflows in harmful ways, or deliberately cause healthcare worker displacement. Thus, careful contextualization is important for CARs.

Our clinical collaborators identified two key task types that robots could help with: patient assessments (e.g., setting up a telemedicine visit, vitals check, pain check), and patient deliveries (e.g., food, water, comfort items). In the ED, both types of tasks happen frequently and across multiple patients simultaneously, but there is insufficient coverage with current staffing levels, and clinicians thought robots could lighten the load.

Simple System with Low Cognitive Demand. ED clinicians are constantly under high stress. They stated the most important thing for CARs is to not add to their existing cognitive load. They wanted (1) a system with straightforward task delegation, (2) the robot to be autonomous and invisible in use, and (3) feedback in case the robot failed.

Through iterative piloting, we co-designed the task specification interface for clinicians to delegate tasks to Min, and receive status updates with a few button clicks. We followed standard design principles such as self-explanatory icons [43], large high-contrast fonts, and help pages. We ensured the overall design aligned with existing apps and tools clinicians used for consistency [74]. In our pilots, we saw clinicians quickly acclimated to the interface and successfully created tasks.

To be “invisible in use” and minimize interrupting clinicians, CARs need to be autonomous with guarantees on their behavior to create trust against performing undesirable actions. Control synthesis methods suit this requirement as they can convert rich high-level user specifications




Fig. 1. Overview of the Clinically Assistive Robot System for Emergency Medicine (CAR-EM). CAR-EM includes an autonomous robot and a task specification interface, completing tasks by leveraging control synthesis, a framework that automatically transforms high-level tasks into control while providing guarantees and feedback.

into correct-by-construction robot controllers that guarantee tasks are completed (liveness), while avoiding undesirable actions (safety).

Need for Continuous and Meaningful Feedback. Clinicians wanted to access the status of the tasks they delegated to the robot. Thus, integrating the robot's status, and progression through tasks into the task specification interface would help them track their tasks, and plan accordingly in case of robot failure. Synthesized controllers have the capability to provide insight into the robot's status, which we used to design a status page that displayed the robot's current task and task step. Additionally, for patient assessment tasks, such as collecting vital signs, clinicians wanted to know when patients' readings fell outside of specific thresholds.

Limited Sensors. Clinicians highlighted the importance of having minimal sensors on the robot to avoid accidentally recording sensitive health information. So, our robot used a 2D lidar for navigation, and its tablet camera was only activated with the mock patients' consent during telemedical visits.

3.2 System Architecture

CAR-EM consists of a task specification interface, a mobile robot with an interaction tablet, and a synthesis framework (see Figure 1). It enables clinicians to request robot tasks via the task specification interface. The synthesis framework synthesizes control for the corresponding task (see Figure 2) and communicates feedback regarding the robot's status to the clinician via the task specification interface (see Figure 3). Mock patients could use the interaction tablet to interact with the robot.

3.2.1 Robot Platform. CAR-EM's mobile robot platform is called Min, and is a customized Turtlebot 2 with an RPLidar-A1 sensor. The main robot controller ran on a laptop mounted on Min's base. A secondary laptop connected to the robot's laptop as a **Robot Operating System (ROS)** client and sent navigation commands to the robot's base.

We mounted an iPad on the robot's body about 3 feet off the ground, to support interaction with people, who may encounter the robot (e.g., clinicians, staff, patients, bystanders). We attached a small LCD screen to the robot's laptop which displayed a neutral, friendly face with blinking eyes to give the robot a basic sense of personality. For item delivery, we attached a small drawer to the robot. Finally, we also gave the robot a scrubs-like costume with a stethoscope to indicate group membership [31].

3.2.2 Interfaces. Task Specification Interface. We built a clinician-facing tablet application per the design requirements described in Section 3.1. Clinicians could add tasks for the robot to perform,




Fig. 2. Example of using CAR-EM to perform a vitals check. Clinicians use the task specification interface to send the robot to a patient room. We pre-specified LTL and STL specification files that indicate (1) sensor and actuator ROS nodes necessary for implementation, (2) initial condition of the robot and environment, and (3) transition and liveness conditions of the robot and environment. The LTL and STL synthesizers use their respective pre-specified specifications to generate corresponding controllers. The controllers enable the robot to perform the task in accordance with the specifications. If any specification is violated during task execution, the task is interrupted with feedback being sent to the clinician.

monitor the robot's status as it progressed through tasks, and receive notifications of robot failure, patient status, and task completion. The app was deployed on an iPad Mini to be portable, so clinicians could carry it with them while they performed other duties.

Communication between the app and robot was facilitated via JSON messages over a web socket. We used different colors to indicate to clinicians when a task was going smoothly, if patients needed attention (e.g., a telemedical visit), or if the robot failed as seen in Figure 3. We implemented telemedical visits through Zoom integration.

Interaction Tablet Interface. For testing purposes, we built a prototype tablet application for the robot's body which could support interaction with mock patients. The application served two purposes: (1) it displayed the text that the robot spoke aloud (to support accessibility), and (2) it asked users to provide consent at every step of task progression, without which the robot could not advance to the next step. For example, after introducing itself, the robot would always ask patients if they were comfortable receiving care from it.

3.3 Robot Autonomy

We leveraged the expressive power of LTL and STL to capture requirements, and synthesis algorithms to create robot autonomy. We captured clinicians' task requests through LTL specifications, and the synthesized LTL controller seamlessly interacted with the STL-based controller that produced Min's low-level control.

Interaction. The robot's interactive component was synthesized from LTL specifications. In this work, we used the GR(1) [13, 77] fragment of LTL as the underlying formalism. We used the open source tool slugs [30] for LTL synthesis and LTLStack [115] in conjunction with slugs to obtain a ROS 1 implementation of the synthesized robot controller. The synthesized LTL controller acted as a supervisory controller, synchronizing communications to and from the clinician interface, robot tablet interface, and the STL controller.

The robot's interactions with mock patients consisted of a series of five steps: greeting the mock patient, verifying their identity for privacy purposes, presenting task-specific instructions, performing the task, and saying goodbye before leaving the mock patient's room. As the robot went through these steps, it sent updates to clinicians through the task specification interface. We encoded the LTL specification in a manner that ensured that the robot never performs a step




Fig. 3. The task specification interface allows clinicians to easily add tasks, monitor the robot's status, and check the patient's status with reasoning in case of unexpected events.

unless all previous steps have been completed. The corresponding LTL specification for the robot's interactions can be found in the Supplement.

The mock patient was given the option to decline interacting with the robot at any time. The LTL specification classified this as a task failure/interrupt, corresponding details of which the task specification module communicated to the clinician's task specification interface. Clinicians could then use this feedback as desired, while the robot returned to its station.

Navigation. The robot's navigation was synthesized from Event-Based STL specifications [39]. Capturing requirements on the motion of the robot using logic and synthesis has two advantages: (1) it allows for more complex requirements such as failure response, timing bounds, and obstacle avoidance and (2) there exist techniques for pinpointing and anticipating failures and their causes, which is crucial for interaction with clinicians. For example, when the robot's task was interrupted or failed, the LTL/STL specifications required it to autonomously return to its station after announcing its intent to do so. The STL specification for this application can be found in the Supplement.

The STL synthesizer converts the specification into an optimization problem, charting a path through the map based on a predefined roadmap of waypoints. As a safety precaution, the robot stops whenever it leaves the designated map or when it detects an obstacle that is less than 0.2 meters away. The robot used Gmapping and AMCL to perform localization with the RPLidar-A1 sensor and onboard wheel odometry.

LTL and STL interrupts were published to two different topics that were continuously monitored by the supervisory LTL controller. When an interrupt was recorded on either topic, the supervisory LTL controller sent the corresponding interrupt message to the task specification interface and terminated the task. To do this, it first triggered the robot to announce task termination, and then routed the robot back to its station. For the purpose of our evaluation, we only explored LTL interrupts to ensure a controlled experimental setup.

4 Feasibility Study: Methodology

We ran a feasibility study deploying CAR-EM in two different hospital EDs to evaluate its usability, and gather initial feedback from clinicians. Researchers served as mock patients in the study, as our goal was to evaluate the clinician-facing components of CAR-EM, and whether our LTL and STL representations adequately represented the behaviors clinicians envisioned. Clinicians tasked Min to perform several patient assessment and item delivery tasks. The robot autonomously navigated to the mock patient sitting in another location. Then, participants provided us feedback on CAR-EM and completed surveys. Our study was approved by the UC San Diego IRB.

Table 1. Participant Demographic Information

Participant #	Age Range	Gender Identity	Role	Experience Range
P1	35–44	Male	Physician	11–15 years
P2	35–44	Female	Nurse	6–10 years
P3	25–34	Male	Nurse	6–10 years
P4	25–34	Male	Tech	6–10 years
P5	35–44	Male	Nurse	11–15 years
P6	45–54	Female	Nurse	21+ years
P7	35–44	Male	Nurse	1–5 years
P8	45–54	Female	Tech	1–5 years
P9	25–34	Female	Nurse	6–10 years
P10	55–64	Female	Physician	21+ years
P11	25–34	Male	Physician	1–5 years

4.1 Participants

Eleven on-duty ED clinicians (3 physicians, 6 nurses, and 2 techs) across two hospital EDs (ED1 and ED2) participated in our study during a 25 minute break. The ED care team consists of (1) *physicians* who must quickly assess patients for severity of their conditions, take histories, perform physical exams, and create treatment plans while managing several patients concurrently, (2) *nurses*, who work with physicians to deliver care. They place IVs, make assessments, administer medications, monitor vital signs, detect changes in clinical status and patients’ clinical condition, and communicate with patients and their families, (3) *techs*, who work with the team to obtain serial vital signs, electrocardiograms, transport patients, and find supplies for procedures or resuscitation.

Two physicians participated at ED2, and nine participants participated at ED1 (See Table 1). Their ages ranged from 25 to 64 years, and their experience ranged from 3 to 25 years. All participants were recruited by word of mouth, gave informed consent, and were not compensated for their participation.

4.2 Measures

All participants completed the **System Usability Scale (SUS)** [20], **NASA Task Load Index (TLX)**, [40], and a demographics questionnaire. SUS assesses the usability of our system through factors such as the system’s complexity, ease of use, and technical consistency. TLX encapsulates participants’ experiences of using the system in terms of effort they expended, their frustrations with the system, and their success in using the system.

While using the robot, participants engaged in prompted think aloud questions to assess their situational awareness of the robot’s actions, such as “What do you think the robot is doing right now?” and “What would your care decision be at this point?” After the study concluded, we engaged in a semi-structured interview, and asked participants about their experience using the robot, the appropriateness of the robot’s tasks, the design and performance of the clinician interface, and broader questions about the rising adoption of AI and other technologies in healthcare.

We had one researcher follow the robot around throughout the entire study. This was for two reasons. First as a safety precaution, since we were testing CAR-EM in fully functioning EDs with real patients. Second, the researcher noted an overview of each run, concerns they had during each run, and if and how they had to intervene to mitigate foreseeable robot failures.

4.3 Analysis

We analyzed data from the think aloud questions and semi-structured interviews using **Reflexive Thematic Analysis (RTA)** [16, 17]. Two researchers independently analyzed each transcript to generate codes. We then iteratively refined codes and resolved inconsistencies through discussions to identify themes based on shared patterns of meaning [18]. Finally, we used the themes to identify broader implications of CARs on clinical workflows, labor displacement, and changes to care. We did not calculate inter-rater reliability per current best practices in qualitative analysis, since we aimed to identify salient concepts and recurrent themes [19, 72]. We used the researcher's notes to classify and count interventions needed to mitigate robot failure.

4.4 Procedure

We conducted our study in two EDs, in a suburban (ED1) and an urban (ED2) area. Both serve a large number of patients, and are nearly always at capacity, with every patient room and hallway bed occupied. In both EDs, participants were seated in a makeshift "robot station," out of sight of the robot or mock patient. We did this to simulate a real situation where a clinician might use the robot to perform a task while also performing other duties. The mock patient sat around the corner from the robot station, and the robot autonomously navigated between the two during the procedure.

After providing informed consent, participants watched a video tutorial to orient them with the task specification interface and robot. Then they completed a practice task (an item delivery), which gave them a chance to use the task specification interface, watch the robot execute the task, and ask follow-up questions.

Due to clinicians' limited availability, they performed two main tasks. Each began with a realistic patient case followed by an opportunity to use the robot to assist them (see Supplement). The first task involved conducting a pain assessment, and the second task was either a telemedical visit (for physicians or nurses) or a delivery task (for ED techs). Clinicians used the task specification interface to request tasks, then, the robot navigated to the mock patient, performed the task, and returned to its station. To simulate task interrupts, the mock patient (researcher) was instructed to always decline the first task when the robot asked for consent. This allowed us to explore how clinicians expected themselves, the robot, and the task specification interface to respond to unexpected interruptions to the robot's task. Clinicians never had direct line of sight of the mock patient and were frequently interrupted by colleagues and other responsibilities. However, they could check the task specification interface to monitor the robot at all times.

During both tasks, and following the procedure, participants answered questions and completed questionnaires. The entire study took about 25 minutes to complete.

5 Feasibility Study: Findings

We describe our findings in this section, reporting on the evaluation of CAR-EM (Section 5.1), and the six themes we identified through RTA (Sections 5.2–5.7).

5.1 System Evaluation

Task Performance. The robot successfully completed all three tasks (one practice task and two main tasks) for all 11 participants, performing a total of 33 tasks. A successful task involved appropriately receiving the clinician's task, navigating autonomously to the task destination, performing the task, returning to the robot's station, and continuously updating the clinician's interface with the robot's status.

As our study was conducted in real EDs, in some cases a researcher engaged in a minor navigation intervention to briefly “nudge” the robot if they foresaw any disruption to clinical work. For example, in four instances the robot attempted to navigate past an unforeseen obstacle or slightly outside the hallway where it was authorized to operate. In these cases, the researcher nudged the robot back to navigate within authorized bounds. In two instances, the robot’s navigation path was in the way of active emergencies. Here, the researcher temporarily stopped the robot by standing in its way as an obstacle, or requested a right of way by saying “Excuse me!” on behalf of the robot.

Task specification interface evaluation. We evaluated how easily participants navigated the task specification interface. In all 11 runs, the clinician was able to assign the robot tasks fully independently. When asked think-aloud questions like, “What do you think the robot is doing now?” 5/11 clinicians used the interface intuitively to gauge the robot’s status. Three clinicians “imagined” what the robot may be doing without checking the interface.

SUS and TLX. Participants had a mean SUS score of 89.25/100 (s.t.d. = 8.17), indicating high usability [11]. They had a mean unweighted raw-TLX score of 12.80 (s.t.d. = 9.12) indicating low perceptions of task load.

5.2 General Impressions

Overall, clinicians responded positively toward CAR-EM. Several enjoyed Min’s appearance. “The robot [was] very cute.” –P6. “I like the smiling and the blinking, that’s cool. I definitely like the cover you guys put on it—definitely brings up a more human touch.” –P1. Others commented on the utility of the robot: “[Min is] returning time to clinicians to perform the joyful parts of their jobs which is interacting with patients” –P1.

All clinicians found the task specification interface easy to use. “The app is very straightforward to use. [It] feels similar to other apps that we’re already used to using. Finding everything was pretty easy. The pictures made it easy to identify where you needed to go.” –P11 “I think the app was really easy to use and very self explanatory. Even if I didn’t have the tutorial, I think I would have been able to figure it out, which is great.” –P2.

5.3 Robot Role

Most participants viewed Min as an assistant or a tool, rather than as a replacement or stand-in. P5 highlighted Min’s role and said “[It’s like an extra] person on the floor to communicate with the patients, be in contact with the patient so that they know [Min is] not one of us, like the nursing staff or doctor. [But, Min is] assessing them and letting them know [the clinician is] taking care of them.”

All clinicians felt supporting busywork was the most useful application for the robot. “It’s tough to do some of these simple tasks, just because they take time. So, having the robot do some of these things is great.” –P1. “[Item deliveries] will be so helpful because when you go to evaluate your patients and they ask you for some comfort items like [a] blanket, or simple water, or snack, then I usually either have to grab [them] myself, which takes me away from clinical task[s]” –P2.

5.4 Task Appropriateness

Most participants felt the robot’s tasks were appropriate given the clinical contexts we presented. For example, “[The telemedical visit] is a perfect example [use case for the robot], when we discharge patients.” –P1. “Simple tasks like these, yeah, this makes sense [for the robot to perform].” –P3.

However, a few participants raised concerns of situations where the robot could be used inappropriately, such as for medication management or hands-on patient-facing tasks. “We’re supposed to reassess our patients [when we give them medication]. If we gave [something non-sensitive like] Tylenol, I would say yes [to using the robot]. But if we gave some type of sedative [...] it would

probably be good [for a clinician] to reassess and make sure the patient is still breathing.” –P5. P3 also noted that some tasks are entirely inappropriate for the robot to perform “Anything you need to be hands on with, [like resuscitation]. And I guess delivering meds, right? That’s not [...] appropriate.”

Some clinicians discussed evaluating whether patients would be receptive to receiving care from a robot before sending it to perform tasks. “In this ER in particular, we have a lot of older adults. I think that they would be turned off by talking to a robot, just ‘cause that’s the generation.” –P9. P4 highlighted the importance of being mindful of patients’ varying emotional states. “If a patient is aggressive or, like, perturbed already, I’m not gonna send a robot in.”

Two clinicians mentioned that tasks such as telemedical visits were only done by physicians and nurses, and thus it may be inappropriate for other clinicians, such as techs, to perform them through the robot. P8, an ED tech, said “Sometimes it can be like a third party [for a tech to do a televisit]. If I send the robot in and [the patient doesn’t] have questions for me and they want the nurse, then you need to [track] down the nurse. I don’t know anything about [the patient’s treatment plan]. I’m not gonna send [the robot].”

5.5 Impacts on Clinical Workflow

Participants commented on how they saw CARs impacting how they performed their duties, and the overarching effects such robots may have on the ED at large. “Getting simple tasks down could be really helpful. As long as it’s not slowing down workflow, [it’s] increasing workflow, and patients feel more taken care of instead of less taken care of and ignored.” –P4. “You’re putting in, delegating all these little tasks [to the robot]. [In] the end, I think it’s just gonna help with the whole flow of everything else [in the ED].” –P7.

This sentiment tied into supporting patient throughput, a common concern in the ED, which several clinicians highlighted. “Especially with this ER, you gotta get people out” –P9. “I think [telemedical visits] would ease up some stress for me because, we have to clear rooms as soon as we can.” –P11.

Clinicians also noted that using the robot could potentially reduce walking distances for clinicians. “As a nurse you kind of run around all day, it’d be nice to just get a simple task like that, just sit down and just talk.” –P3.

However, clinicians were also wary of how the robot may disrupt their workflows. “I think [it] would disrupt my workflow if the robot wasn’t well received. Because then I went through the time of using the app, getting [Min] to [start the task] and then waiting for it to come back when I could have just done [the task] myself.” –P11. “I’ll be frustrated if there’s a pause in the [execution] of the robot’s task and having a feeling like I need to run back and check what’s going on with the robot.” –P1.

5.6 The Human Element of Care

Some clinicians said they would not feel comfortable delegating tasks to a robot. “Maybe that’s just how I am as a nurse, so I’d rather [do it.] It’s really hard to delegate tasks, especially to a robot.” –P9. “ [The clinician] just [doesn’t] know what’s happening. I guess the interaction [between the robot and the patient] makes me nervous.” –P4.

Clinicians also emphasized that laying eyes would help them get context about patient conditions that they would miss if a robot did the task instead of them. “Even though I’m just bringing a patient water, I get to look, and as a nurse you could [notice something’s off]. Like, wait a minute, take a look at their monitor. That’s where a lot of things are caught. If it’s just a robot [the nurse is] not gonna be able to see the entire context of the patient. Because, as soon as you walk into a room as a clinician, your immediate thing is a personal impression. Every time you look at [a patient],

how sick do they look?” –P4. “Personally as a nurse, I’d like to lay eyes. We’re supposed to reassess our patients.” –P5.

Participants also reflected that human touch is a large part of care provision. “[The robot] takes away from human interaction [...]. You might deliver someone food and then they might have another question or concern, and sometimes they [...] use those times in between to talk to the nurse or the tech to ask more questions, get more help, or to request a physician, which might not always be available for the robot.” –P11. “I’m not quite sure yet how I feel [about AI/technology in healthcare]. Part of me feels like it could be really beneficial and awesome and part of me is like ugh, like this takes away from the personal nursing interaction.” –P9. “A part of care is in person [...] the patient needs to feel like they’re being taken care of, and that’s a constant battle for us. Cause they already are waiting for hours, [and] if you’re ignoring them [further by using a robot, that] does take away that human compassion aspect.” –P4.

5.7 Situational Awareness via the Task Specification Interface

When the robot performed tasks smoothly, clinicians were happy with the feedback our tablet interface provided. “I think that the prompt that it brings up every time it’s changing to the next stage. [Like] introducing [itself], starting the task, task completed, I think that was good to know because you’re not visually seeing what it’s doing.” –P2. “I like that I knew what was going on. [...] I need a lot of things done [in my job]. I ask for them to get done, and then I have no idea if they’ve been done, if they’re in progress, if there’s a delay or [more]. So I love that it says, like verifying or, in process [...]. So that was really helpful to having some sense of status.” –P10.

However, when things went wrong, clinicians wanted more situational awareness to know why, and to recover from it. “What would be kind of cool is if there is an obstruction and [the robot] is paused for any reason, having the ability to televideo in at that point.” –P1. “[I want a transcription of] just whatever the patient says. I [prefer] what I [can] see [for myself]. So maybe have some sort of microphone capability that was recording.” –P7.

6 Discussion

Overall, CAR-EM was successful in meeting clinicians’ expectations for the system. The robot autonomously performed all runs successfully, and the system functioned smoothly end-to-end. Clinicians enjoyed the robot’s embodiment, and found the system both useful and straightforward to use. They gave us rich feedback on their preferences regarding CAR-EM’s feedback mechanisms, their sense of situational awareness while using it, and discussed some ethical considerations for moving forward. Our experiences deploying the robot across two EDs uncovered findings about peoples’ perceptions of the system across the different locations. We also learned how such systems may integrate into existing clinical work patterns. We discuss these topics in detail below.

6.1 Invisible in Success, Informative in Failure

In our study, we observed that when a robot is perceived as capable of doing something correctly, clinicians want it to be autonomous, and not think about it (i.e., the robot is “invisible in use”). However, when something goes wrong, they want as much information as possible as to why, and be able to control the robot to mitigate failure. For example, when the mock patient declined to receive a pain assessment from the robot, P1 said “I would like to know why they declined to participate in the task. Let’s see here, maybe I can televideo visit with them to talk to them?”

Our use of synthesized control ensured that CAR-EM was designed from the ground up to meet this primary requirement: that the robot be invisible in success and informative in failure. A core capability of synthesized control is knowing when the robot succeeded, or knowing why it failed. This ensures that when the robot provides feedback, it is not possible for the feedback to be

erroneous. However, the robot could fail silently during network issues, for which the system should be designed to have backups. During our study, we observed two open questions on translating synthesis provided feedback to end-users that are interesting to explore in future work: (1) when and how do we provide failure information to clinicians? and (2) how do we enable clinicians to act on it?

Clinicians in our study had mixed perspectives on when the robot should alert them in case of a task failure. For example, some clinicians felt that the robot interrupting them frequently would outweigh the help it could provide by offloading tasks. Other clinicians felt that minor interruptions, like the robot letting them know when a patient declined interacting with it, would be useful in supplementing their understanding of what a patient is comfortable with. This highlights the importance of carefully designing how to deliver failure information, as doing so is not trivial, and can lead to alert fatigue and burnout [7, 91].

Second, our study demonstrated the nuances of determining how best to let clinicians “take over” for the robot, and in what circumstances. For example, the robot getting stuck somewhere while delivering jello may not be something worthy of handing control to a clinician, but if it is on its way to initiate a televisit, it is. Recent robotics research has explored interruption in depth [10], including clinical interruptions [69]. This can provide valuable insights for creating systems like CAR-EM in future work.

6.2 Physical Touch

In our study, some clinicians expressed that anything that involves touching a patient should be done by a clinician, not a robot. This aligns with work by Wright [117], who studied the deployment of the Hug patient-transfer robot deployed at a nursing home in Japan [88]. While theoretically a great idea (saving care workers from back injuries), within days it had been relegated to the closet. Neither residents nor care workers liked it—it was uncomfortable for residents, and care workers found it to decrease the single most important characteristic of their job: “caring with one’s own hands” to maintain both kinship and skinship with residents.

While none of CAR-EM’s current tasks involve touching patients, our clinical collaborators frequently suggested it perform tasks like supporting patient mobility, and placing ECG electrodes, pulse oximeters, and such on patients’ bodies in future iterations. Further research is needed to evaluate the appropriateness of such tasks within the broader ethical, psychological, and social frameworks of care provision.

Interestingly, clinicians were also against patients touching the robot, due to concerns about sanitation “If the patient is [dirty] they can’t touch [the robot]. That feels unsanitary.” –P9, and safety, “Someone’s gonna pick [Min] up and move it. Don’t hurt Min, dude! [hugs Min]” –P5, aligning with prior work [103]. These suggest that future CARs should consider moving away from touch-screen based HRIs to limit the spread of infection, and weighing down CARs in areas of high foot-traffic to avoid their destruction/misuse.

6.3 Ethical Considerations

Potential Worker Displacement. In our own research processes contextualizing tasks for CAR-EM, we identified tasks (e.g., delivering lab samples) that could displace some clinicians. For example, ED techs have a large number of diverse roles that even physicians and nurses may not be fully aware of, and thus may propose some of those roles to be displaced by automation. Similarly, while clinicians liked Min and thought it could have useful purposes, we also heard sidebar comments like, “Oh, you’re here to take my job away,” and, “I can’t believe [the hospital system] is spending money on this instead of paying us better.” These concerns highlight the invisibility of care work [98], and how robots could displace it.

Recent works in HRI have noted the invisibility of labor [61, 63, 73, 99], and the increased invisibility of marginalized groups [62]. Research exploring articulation work also show that clinicians make continuous adjustments, and reconfigure their professional roles to accommodate new technologies [58]. They engage in “invisible,” and unpaid labor while interfacing with robots, such as by clearing hallways for robots, and adjusting their schedules around robots.

This invisibility is reinforced with the recent proliferation of AI chatbots [6, 21] aimed to provide “cheaper” more “empathic” care. These concerns are particularly pressing in high-stakes, emotionally charged environments such as acute care, where patients and their loved ones seek human compassion, and reassurance. They also reflect more complex and nuanced socioethical questions of how CARs may support worker-centered HRI [87], rather than displace or disrupt clinical workflows. Thus, future research could approach CARs through a care ethics lens [46, 112], ensuring deliberate attention to the potential downstream effects of such technologies [47, 113].

Changes to Power. Concerns of worker displacement are reflective of broader systemic issues in healthcare as a result of increasing adoption of AI and technologies. While hospital delivery robots like TUG and Moxi have been reported to have been well-received among staff, [14, 25, 108], recent work [111] suggests resentment may be brewing, with nurses in particular displeased with how these robots are deployed in hospitals. Ironically, one of our participants who took a selfie with Min was wearing a button from a recent “Trust Nurses not AI” protest from the **National Nurses Union (NNU)**, which highlighted the importance of trusting human healthcare rather than AI [8]. NNU is the largest nursing union in the United States, and in 2024 organized protests at two prominent US hospitals to draw attention to the dangers of AI in nursing [8].

Such protests suggest clinicians’ powerlessness toward the deployment of technologies that affect their livelihood, workplace, and the misplacement of patients’ trust. Given this climate, and the fact that aside from Moxi, few robots have been collaboratively designed with nurses [9], these concerns are not surprising. They offer us, as researchers, a chance to reflect on the broader context of automation within complex socio-technical work contexts, particularly at this moment of proliferation of generative AI technologies being rapidly introduced into clinical workflows. We need to ask questions about power structures in this context, aligning with recent work in the field [35, 41, 62, 114], and engage in reflexivity of our own roles as robotics researchers in how we move forward.

Changes to Patient Care and the Importance of “Laying Eyes.” While CARs may reduce some of the difficulties patients face, including long wait times [29, 50], they also may displace important opportunities for observing patients. Our participants discussed the importance of a clinician “laying eyes” on a patient, i.e., noticing key changes in their condition. Prior research has shown that experienced critical care clinicians often rely on intuition and quick thinking to identify patients’ conditions while laying eyes [27, 85]. Thus, it is important to be careful not to displace such opportunities, lest there are deleterious patient care effects.

Using Robots for Even More Multitasking. Unsurprisingly, clinicians identified more ways by which they could multitask with the help of the robot. “If I’m stuck in a patient’s room, [...] watching the patient for some time, I can check on [a different patient]. Just [to] make sure that they don’t need anything.” –P9. Prior HRI research shows that clinicians will use robots to multitask regardless of a researcher’s best intentions [69, 70]. Thus, it is good to plan for potential multitasking, such as by designing systems that are mindful of how they interrupt clinicians.

Despite this, several clinicians noted that using the robot would streamline their workflows. “[If I use the robot,] I can stay at my workstation and continue to write notes, and start the Zoom call. ‘Do you have any questions? [...] I’ll send you home with meds’ and I can put the order in right there on my computer [instead of walking around and getting interrupted]. So it’s almost faster to get to that disposition.” –P10. Thus, there exists a fine balance between using the robot for

more multitasking versus using the robot to streamline workflow that comes down to clinicians' individual attitudes.

6.4 Differences in Receptiveness across EDs

We noticed that across the two EDs that we tested Min in, we received different levels of receptiveness from both clinicians and other bystanders. One reason for this may be socio-economic differences between the locations of the two EDs. "I think my only fear is that sometimes depending on the patient population [reception towards the robot may be different]. I think here at [Hospital 2], this robot might be refused a lot more or [even] harmed. But I can definitely see that in [Hospital 1], this might be something that's a lot more helpful. It might be more well received there compared to here." –P11. Another reason may be that ED2 had very narrow hallways compared to ED1, and so the robot served as an additional navigational obstruction for clinicians.

Additionally, in Hospital 1 we had a clinical research advocate who supported us in conducting our study by introducing us to several clinicians and garnering interest. On the other hand, in Hospital 2, participant recruitment was somewhat more opportunistic in nature in the absence of the clinical research advocate. Study methodologies for deploying robots in workplaces may benefit from having a research advocate who is "one of the people" in the loop.

6.5 Lessons from a Technical Standpoint

We learned that through synthesis, we can address clinicians' core requirements for guarantees and feedback on robot operation. We demonstrated our system working well with off the shelf hardware (TurtleBot) and localization algorithms (AMCL) in the complex and dynamic space that is the ED, furthering support for efforts to facilitate low cost, open-HW/SW robot deployments [70]. While we did not explore feedback provision in depth, our system lays the groundwork for improving feedback in future work. We also learned that social navigation is a necessity for CARs in the ED, supporting other prior HRI research [78, 106]. For example, sometimes our robot wrangler had to say "excuse me" on behalf of the robot while it was navigating. This is a complex socio-technical space which warrants further investigation, particularly with how robots respond to ED bystanders, narrow hallways, and critical moments.

Deploying Synthesis-Based Systems in the Real World. Our implementation of CAR-EM to realize end-user specific requirements in EDs highlights a novel approach towards the safe, user-centric design of real-world robot systems. Representing user requirements as formal specifications that are synthesized into robot control implies that CAR-EM's control is correct-by-construction. The correct-by-construction guarantees that control synthesis provides can be leveraged by researchers developing similar robots that operate in safety-critical environments that demand guarantees for safe robot operation. Our study shows that synthesis outputs can be further used to provide users feedback regarding the robot's status—a property that can be explored in contexts that demand explainability. Thus, exploring the potential limitations of control synthesis as highlighted in our feasibility study, especially in environments as variable as the ED, could add valuable insight into practical applications and areas of system refinement.

Our demonstration also highlights the feasibility of implementing social robot behaviors such as interactions through the use of LTL and STL. Future research could explore the use of verification and synthesis to develop appropriate social robot behaviors for real-world contexts and applications.

6.6 Limitations and Future Work

Our study faced a few limitations that we intend to address in future work. First, our participants were recruited on a voluntary basis, and thus participated only if they had overall positive

perceptions of technology. Also, they were limited in geographic diversity, and future research will benefit from toward more generalizable results.

Second, while from a technical perspective our system is fully autonomous, we acknowledge the complex context of the ED requires human supervision of the robot's functioning. Full autonomy is highly context-, and ED-specific. This is even reflected in the setup procedures of Moxi, a fully autonomous CAR, where the robot's maps, and functions are first fine-tuned to align with the specific hospital of deployment. We will continue to fine-tune CAR-EM in future work.

Third, we did not study the patient-facing aspects of CAR-EM in this work, and conducted our study with mock patients instead. This was because studying the patient-facing aspects of CARs involves a deeper exploration of what it means for robots to care in high-stakes environments, which is beyond the scope of this work. We are currently conducting collaborations with ED patients to design the patient-facing aspects of CAR-EM. We are exploring the nature of interpersonal relationships and boundaries with the robot, and patient perspectives of robot interactions in the absence of human interactions. In conjunction, we are designing patient interactions with CAR-EM, including how the robot should be introduced to patients, and how it may deliver clinical content such as pain assessments in an accessible manner.

Finally, while we studied the sociotechnical aspects of providing feedback and situational awareness to clinicians, we have yet to explore how clinicians may act upon the feedback. Moving forward, we will explore how different types of failure feedback may be prioritized, and communicated to clinicians [45]. We will study how clinicians make sense of the feedback, and explore how they may act to mitigate failures.

Failure mitigation for CAR-EM can take two forms: autonomous failure mitigation, or by asking clinicians for help. Formal methods researchers have explored failure mitigation through specification repair [75], and *a-priori* failure warnings [39]; these approaches could enable CARs to mitigate failure autonomously. However, there is limited research exploring when and how CARs should ask clinicians for help. Further research is needed to determine how control handovers (e.g., teleoperation [36]), and other mitigation strategies in such circumstances should be implemented.

In the future, we will conduct week-long longitudinal studies with CAR-EM to explore workflow reconfigurations, and clinician- and patient-facing robot relationships in greater depth. Through this planned research, we hope to contribute toward developing safe, privacy-sensitive robot systems that satisfy the context-specific needs of acute care.

6.7 Conclusion

In this article, we present CAR-EM, a synthesis-based clinically assistive robot system for acute care collaboratively designed with clinicians. We present a first evaluation of this system with clinicians in real-world ED environments, which unravels new knowledge that informs design and implementation considerations for future CARs. Through these contributions, we hope researchers explore the intersections of formal methods and HRI toward correct and safe robot systems that satisfy the context-specific requirements of safety-critical spaces.

References

- [1] Layla Abdullah and Fadi El-Hassan. 2022. Evaluating the efficacy of autonomous robotic systems for enhanced sanitation in healthcare environments. *International Journal of Responsible Artificial Intelligence* 12, 12 (2022), 1–8.
- [2] Maysam Ali Abdulwahid, Andrew Booth, Janette Turner, and Suzanne M. Mason. 2018. Understanding better how emergency doctors work. Analysis of distribution of time and activities of emergency doctors: A systematic review and critical appraisal of time and motion studies. *Emergency Medicine Journal* 35, 11 (2018), 692–700.
- [3] Daron Acemoglu and Pascual Restrepo. 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33, 2 (2019), 3–30.

- [4] Sara L. Ackerman, Kathleen Tebb, John C. Stein, Bradley W. Frazee, Gregory W. Hendey, Laura A. Schmidt, and Ralph Gonzales. 2012. Benefit or burden? A sociotechnical analysis of diagnostic computer kiosks in four California hospital emergency departments. *Social Science & Medicine* 75, 12 (2012), 2378–2385.
- [5] Aethon. 2022. Hospital Robots—Solutions for Healthcare. Retrieved October 1, 2024 from <https://aethon.com/mobile-robots-for-healthcare/>
- [6] Hippocratic AI. 2024. Hippocratic AI. Retrieved October 12, 2024 from <https://hippocraticai.com/>
- [7] Jessica S. Ancker, Alison Edwards, Sarah Nosal, Diane Hauser, Elizabeth Mauer, and Rainu Kaushal and With the HITEC Investigators. 2017. Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system. *BMC Medical Informatics and Decision Making* 17 (2017), 1–9.
- [8] California Nurses Association. 2024. California Nurses Association Demand Patient Safeguards Against Artificial Intelligence Technology. Retrieved October 1, 2024 from <https://www.nationalnursesunited.org/press/cna-demand-patient-safeguards-against-artificial-intelligence>
- [9] Grace Titilayo Babalola, Jenna-Marie Gaston, Joseph Trombetta, and Stephanie Tulk Jesso. 2024. A systematic review of collaborative robots for nurses: Where are we now, and where is the evidence? *Frontiers in Robotics and AI* 11 (2024), 1398140.
- [10] Siddhartha Banerjee, Andrew Silva, and Sonia Chernova. 2018. Robot classification of human interruptibility and a study of its effects. *ACM Transactions on Human-Robot Interaction* 7, 2 (2018), 1–35.
- [11] Aaron Bangor, Philip T. Kortum, and James T. Miller. 2008. An empirical evaluation of the system usability scale. *International Journal of Human–Computer Interaction* 24, 6 (2008), 574–594.
- [12] Emily Bender and Alex Hanna. 2024. Episode 37: Chatbots Aren’t Nurses (Feat. Michelle Mahon). July 22, 2024—Mystery AI Hype Theater 3000.
- [13] Roderick Bloem, Barbara Jobstmann, Nir Piterman, Amir Pnueli, and Yaniv Sa’ar. 2012. Synthesis of reactive (1) designs. *Journal of Computer and System Science* 78, 3 (2012), 911–938.
- [14] Richard Bloss. 2011. Mobile hospital robots cure numerous logistic needs. *Industrial Robot: An International Journal* 38, 6 (2011), 567–571.
- [15] A. Boutou, G. Pitsiou, E. Sourla, and I. Kioumis. 2019. Burnout syndrome among emergency medicine physicians: An update on its prevalence and risk factors. *European Review for Medical and Pharmacological Sciences* 23, 20 (2019), 9058–9065.
- [16] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101.
- [17] Virginia Braun and Victoria Clarke. 2012. *Thematic Analysis*. American Psychological Association.
- [18] Virginia Braun and Victoria Clarke. 2013. *Successful Qualitative Research: A Practical Guide for Beginners*. SAGE Publications.
- [19] Virginia Braun and Victoria Clarke. 2021. One size fits all? What counts as quality practice in (reflexive) thematic analysis? *Qualitative Research in Psychology* 18, 3 (2021), 328–352.
- [20] J. Brooke. 1996. SUS: A quick and dirty usability scale. *Usability Evaluation in Industry*. Patrick W. Jordan, B. Thomas, Ian Lyall McClelland, and Bernard Weerdmeester (Eds.).
- [21] Julia E. H. Brown and Jodi Halpern. 2021. AI chatbots cannot replace human interactions in the pursuit of more inclusive mental healthcare. *SSM - Mental Health* 1 (2021), 100017.
- [22] David Johan Christensen, Rune Fogh, and Henrik Hautop Lund. 2014. Playte, a tangible interface for engaging human-robot interaction. In *Proceedings of the 23rd IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, 56–62.
- [23] Edmund M. Clarke. 1997. Model checking. In *Proceedings of the 17th Conference on Foundations of Software Technology and Theoretical Computer Science*. Springer, 54–56.
- [24] Edmund M. Clarke, Jr, Orna Grumberg, Daniel Kroening, Doron Peled, and Helmut Veith. 2018. *Model Checking*. MIT Press.
- [25] Bonnie Clipper, Joyce Batcheller, Andrea L. Thomaz, and Agata Rozga. 2018. Artificial intelligence and robotics: A nurse leader’s primer. *Nurse Leader* 16, 6 (2018), 379–384.
- [26] Enrique Coronado, Fulvio Mastrogiovanni, Bipin Indurkha, and Gentiane Venture. 2020. Visual programming environments for end-user development of intelligent and social robots, a systematic review. *Journal of Computer Languages* 58 (2020), 100970.
- [27] Mark Anthony Dalton. 2022. *Understanding the Process of Nurses Recognition and Response to Patient Deterioration*. Liverpool John Moores University.
- [28] UTMB Angleton Danbury. 2023. Moxi Case Study. Retrieved September 27, 2024 from https://static1.squarespace.com/static/606f1bb0f7e05e3329035ff8/t/6532d225cb1af10b0133130b/1697829419700/UTMB_Case_Study_Web.pdf
- [29] Polly J. Davenport, Stephen J. O’Connor, Jeffery M. Szychowski, Amy Y. Landry, and S. Robert Hernandez. 2017. The relationship between emergency department wait times and inpatient satisfaction. *Health Marketing Quarterly* 34, 2 (2017), 97–112.

- [30] Rüdiger Ehlers and Vasumathi Raman. 2016. Slugs: Extensible gr (1) synthesis. In *Proceedings of the 28th International Conference Computer Aided Verification (CAV '16)*. Springer, 333–339.
- [31] Friederike Eyssel and Dieta Kuchenbrandt. 2012. Social categorization of social robots: Anthropomorphism as a function of robot group membership. *The British Journal of Social Psychology* 51, 4 (2012), 724–731.
- [32] Jahan Fahimi and Lewis Goldfrank. 2019. Principles of social emergency medicine. *Annals of Emergency Medicine* 74, 5 (2019), S6–S10.
- [33] Figma. 2024. Figma. Retrieved October 12, 2024 from <https://www.figma.com/>
- [34] Sarah E. Fox, Samantha Shorey, Esther Y. Kang, Dominique Montiel Valle, and Estefania Rodriguez. 2023. Patchwork: The hidden, human labor of AI integration within essential work. *Proceedings of the ACM on Human-Computer Interaction* 7, CSCW1 (2023), 1–20.
- [35] Pratyusha Ghosh, B. Leido, and Laurel D. Riek. 2026. Disability justice in human-robot interaction: Reflections on paternalism, autonomy, and care for more equitable futures. In *Proceedings of the 2026 ACM/IEEE International Conference on Human-Robot Interaction*.
- [36] Pratyusha Ghosh, Sachiko Matsumoto, Vivek Gupte, Robert Bloom, Alex Chow, Nandini Desai, Donovan Le, Tania Morimoto, and Laurel D. Riek. 2026. Input matters: How telepresence control devices affect performance, sense of control, and user experience. In *Proceedings of the 2026 ACM/IEEE International Conference on Human-Robot Interaction*.
- [37] Michael J. Gonzales, Vanice C. Cheung, and Laurel D. Riek. 2015. Designing collaborative healthcare technology for the acute care workflow. In *Proceedings of the 2015 9th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth)*. IEEE, 145–152.
- [38] David Gundana and Hadas Kress-Gazit. 2021. Event-based signal temporal logic synthesis for single and multi-robot tasks. *IEEE Robotics and Automation Letters* 6, 2 (2021), 3687–3694.
- [39] David Gundana and Hadas Kress-Gazit. 2022. Event-based signal temporal logic tasks: Execution and feedback in complex environments. *IEEE Robotics and Automation Letters* 7, 4 (2022), 10001–10008.
- [40] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (task load index): Results of empirical and theoretical research. *Advances in Psychology* 52 (1988), 139–183.
- [41] Yoyo Tsung-Yu Hou, EunJeong Cheon, and Malte F. Jung. 2024. Power in human-robot interaction. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 269–282.
- [42] Ayanna M. Howard, Chung Hyuk Park, and Sekou Remy. 2011. Using haptic and auditory interaction tools to engage students with visual impairments in robot programming activities. *IEEE Transactions on Learning Technologies* 5, 1 (2011), 87–95.
- [43] Icons. 2024. Free Icons, Clipart Illustrations, Photos, and Music. Retrieved October 2, 2024 from <https://icons8.com/>
- [44] Tariq Iqbal, Michael J. Gonzales, and Laurel D. Riek. 2015. Joint action perception to enable fluent human-robot teamwork. In *Proceedings of the 2015 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 400–406.
- [45] Rabeya Jamshad, Arthi Haripriyan, Advika Sonti, Susan Simkins, and Laurel D. Riek. 2024. Taking initiative in human-robot action teams: How proactive robot behaviors affect teamwork. In *Proceedings of the Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 559–562.
- [46] Sandhya Jayaraman, Pratyusha Ghosh, Soyoon Kim, Soham Satyadharma, Angelique Taylor, Christopher Coyne, and Laurel D. Riek. 2026. Designing care-fully: Robots for acute cancer care. In *Proceedings of the 21st Annual ACM/IEEE Conference on Human Robot Interaction (HRI)*.
- [47] Sandhya Jayaraman, Deep Masanam, Pratyusha Ghosh, Alyssa Kubota, and Laurel D. Riek. 2026. Robots for care: Exploring downstream socio-ethical effects and upstream interventions. In *Proceedings of the 21st Annual ACM/IEEE Conference on Human Robot Interaction (HRI)*.
- [48] Sandhya Jayaraman, Elizabeth K. Phillips, Daisy Church, and Laurel D. Riek. 2024. Privacy and utility perceptions of social robots in healthcare. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100039.
- [49] Amy Johnston, Louisa Abraham, Jaimi Greenslade, Ogilvie Thom, Eric Carlstrom, Marianne Wallis, and Julia Crilly. 2016. Staff perception of the emergency department working environment: Integrative review of the literature. *Emergency Medicine Australasia* 28, 1 (2016), 7–26.
- [50] Peter Jones and Karen Schimanski. 2010. The four hour target to reduce emergency department ‘waiting time’: A systematic review of clinical outcomes. *Emergency Medicine Australasia* 22, 5 (2010), 391–398.
- [51] Joshua W. Joseph and Benjamin A. White. 2020. Emergency department operations: An overview. *Emergency Medicine Clinics of North America* 38, 3 (2020), 549–562.
- [52] Rongsheng Kee, Jonathan C. Knott, Suelette Dreyfus, Reeva Lederman, Simon Milton, and Keith Joe. 2012. One hundred tasks an hour: An observational study of emergency department consultant activities. *Emergency Medicine Australasia* 24, 3 (2012), 294–302.

- [53] Ralf Kittmann, Tim Fröhlich, Johannes Schäfer, Ulrich Reiser, Florian Weißhardt, and Andreas Haug. 2015. Let me introduce myself: I am Care-O-bot 4, a gentleman robot. In *Proceedings of the Mensch und Computer*.
- [54] Hadas Kress-Gazit, Kerstin Eder, Guy Hoffman, Henny Admoni, Brenna Argall, Rüdiger Ehlers, Christoffer Heckman, Nils Jansen, Ross Knepper, Jan Křetínský, et al. 2021. Formalizing and guaranteeing human-robot interaction. *Communications of the ACM* 64, 9 (2021), 78–84.
- [55] Hadas Kress-Gazit, Georgios E. Fainekos, and George J. Pappas. 2009. Temporal-logic-based reactive mission and motion planning. *IEEE Transactions on Robotics* 25, 6 (2009), 1370–1381.
- [56] Hadas Kress-Gazit, Morteza Lahijanian, and Vasumathi Raman. 2018. Synthesis for robots: Guarantees and feedback for robot behavior. *Annual Review of Control, Robotics, and Autonomous Systems* 1, 1 (2018), 211–236.
- [57] Hadas Kress-Gazit, Tichakorn Wongpiromsarn, and Ufuk Topcu. 2011. Correct, reactive, high-level robot control. *IEEE Robotics & Automation Magazine* 18, 3 (2011), 65–74.
- [58] Antonia Krummheuer and Kristina Tornbjerg Eriksen. 2025. The invisible work of human-robot collaboration: When streamlined processes meet the complexity of real-world practices. *Interactions* 32, 2 (2025), 36–41.
- [59] Alyssa Kubota, Emma I. C. Peterson, Vaishali Rajendren, Hadas Kress-Gazit, and Laurel D. Riek. 2020. JESSIE: Synthesizing social robot behaviors for personalized neurorehabilitation and beyond. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 121–130.
- [60] Maria Kyrarini, Fotios Lygerakis, Akilesh Rajavenkatanarayanan, Christos Sevastopoulos, Harish Ram Nambiappan, Kodur Krishna Chaitanya, Ashwin Ramesh Babu, Joanne Mathew, and Fillia Makedon. 2021. A survey of robots in healthcare. *Technologies* 9, 1 (2021), 8.
- [61] Hee Rin Lee. 2024. Contrasting perspectives of workers: Exploring labor relations in workplace automation and potential interventions. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 1–17.
- [62] Hee Rin Lee, EunJeong Cheon, Chaeyun Lim, and Kerstin Fischer. 2022. Configuring humans: What roles humans play in HRI research. In *Proceedings of the 2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 478–492.
- [63] Hee Rin Lee, Sarah Fox, EunJeong Cheon, and Samantha Shorey. 2025. Minding the Stop-Gap: Attending to the “temporary,” unplanned, and added labor of human-robot collaboration in context. In *Proceedings of the 2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 34–44.
- [64] Lars Lindemann and Dimos V. Dimarogonas. 2018. Control barrier functions for signal temporal logic tasks. *IEEE Control Systems Letters* 3, 1 (2018), 96–101.
- [65] Lars Lindemann and Dimos V. Dimarogonas. 2019. Control barrier functions for multi-agent systems under conflicting local signal temporal logic tasks. *IEEE Control Systems Letters* 3, 3 (2019), 757–762.
- [66] Oded Maler and Dejan Nickovic. 2004. Monitoring temporal properties of continuous signals. In *Proceedings of the International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems*. Yassine Lakhnech and Sergio Yovine (Eds.), *Lecture Notes in Computer Science*, 152–166.
- [67] Martti. 2025. Martti: On-Demand Medical Interpretation Services. Retrieved from April 12, 2025 <https://www.martti.io/>
- [68] Alexandra Mateescu and Madeleine Elish. 2019. AI in Context: The Labor of Integrating New Technologies. Retrieved from <https://apo.org.au/node/217456>
- [69] Sachiko Matsumoto, Pratyusha Ghosh, Rabeya Jamshad, and Laurel D. Riek. 2023. Robot, uninterrupted: Telemedical robots to mitigate care disruption. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 495–505.
- [70] Sachiko Matsumoto, Sanika Moharana, Nimisha Devanagondi, Leslie C. Oyama, and Laurel D. Riek. 2021. Iris: A low-cost telemedicine robot to support healthcare safety and equity during a pandemic. In *Proceedings of the International Conference on Pervasive Computing Technologies for Healthcare*. Springer, 113–133.
- [71] Sachiko Matsumoto and Laurel D Riek. 2022. Shared control in human robot teaming: Toward context-aware communication. arXiv:2203.10218. Retrieved from <https://arxiv.org/abs/2203.10218>
- [72] Nora McDonald, Sarita Schoenebeck, and Andrea Forte. 2019. Reliability and inter-rater reliability in qualitative research: Norms and guidelines for CSCW and HCI practice. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
- [73] Sanika Moharana, Alejandro E. Panduro, Hee Rin Lee, and Laurel D. Riek. 2019. Robots for joy, robots for sorrow: Community based robot design for dementia caregivers. In *Proceedings of the 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 458–467.
- [74] Jakob Nielsen. 1989. Coordinating user interfaces for consistency. *ACM SIGCHI Bulletin* 20, 3 (1989), 63–65.
- [75] Adam Pacheck and Hadas Kress-Gazit. 2023. Physically feasible repair of reactive, linear temporal logic-based, high-level tasks. *IEEE Transactions on Robotics* 39 (2023), 4653–4670.
- [76] Fabio Paternò and Carmen Santoro. 2019. End-user development for personalizing applications, things, and robots. *International Journal of Human-Computer Studies* 131 (2019), 120–130.

- [77] Nir Piterman, Amir Pnueli, and Yaniv Sa'ar. 2006. Synthesis of reactive (1) designs. In *Proceedings of the 7th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI '06)*. Springer, 364–380.
- [78] Ashwini Pokle, Roberto Martin-Martin, Patrick Goebel, Vincent Chow, Hans M. Ewald, Junwei Yang, Zhenkai Wang, Amir Sadeghian, Dorsa Sadigh, Silvio Savarese, et al. 2019. Deep local trajectory replanning and control for robot navigation. In *Proceedings of the 2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 5815–5822.
- [79] David Porfirio, Evan Fisher, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2019. Bodystorming human-robot interactions. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*, 479–491.
- [80] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2018. Authoring and verifying human-robot interactions. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*, 75–86.
- [81] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2019. Computational tools for human-robot interaction design. In *Proceedings of the 2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 733–735.
- [82] David Porfirio, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2020. Transforming robot programs based on social context. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–12.
- [83] David Porfirio, Laura Stegner, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2023. Sketching robot programs on the fly. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 584–593.
- [84] David J. Porfirio, Laura Stegner, Maya Cakmak, Allison Sauppé, Aws Albarghouthi, and Bilge Mutlu. 2021. Figaro: A tabletop authoring environment for human-robot interaction. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–15.
- [85] Sue Holland Pyles and Phyllis Noerager Stern. 1983. Discovery of nursing gestalt in critical care nursing: The importance of the Gray Gorilla syndrome. *Image: The Journal of Nursing Scholarship* 15, 2 (1983), 51–57.
- [86] Laurel D. Riek. 2017. Healthcare robotics. *Communications of the ACM* 60, 11 (2017), 68–78.
- [87] Laurel D. Riek and Lilly Irani. 2025. The future is Rosie?: Disempowering arguments about automation and what to do about it. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–14.
- [88] Hug Robot. 2024. Top Page. Retrieved from <http://hug.fuji.co.jp/en/>
- [89] Diligent Robotics. 2024. Moxi—Diligent Robotics. Retrieved October 1, 2024 from <https://www.diligentrobots.com/moxi>
- [90] Silvia Rossi, Silvano Junior Santini, Daniela Di Genova, Gianpaolo Maggi, Alberto Verrotti, Giovanni Farello, Roberta Romualdi, Anna Alisi, Alberto Eugenio Tozzi, and Clara Balsano. 2022. Using the social robot NAO for emotional support to children at a pediatric emergency department: Randomized clinical trial. *Journal of Medical Internet Research* 24, 1 (2022), e29656.
- [91] Jess L. Rush, Jamil Ibrahim, Kenneth Saul, and Robert T. Brodell. 2016. Improving patient safety by combating alert fatigue. *Journal of Graduate Medical Education* 8 (2016), 620–621.
- [92] Miguel Sarabia, Noel Young, Kelly Canavan, Trudi Edginton, Yiannis Demiris, and Marcela P. Vizcaychipi. 2018. Assistive robotic technology to combat social isolation in acute hospital settings. *International Journal of Social Robotics* 10 (2018), 607–620.
- [93] Neil Savage. 2022. Robots rise to meet the challenge of caring for old people. *Nature* 601, 7893 (2022), 8–10.
- [94] Yasaman S. Sefidgar, Prerna Agarwal, and Maya Cakmak. 2017. Situated tangible robot programming. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 473–482.
- [95] Yasaman S. Sefidgar and Maya Cakmak. 2018. End-user programming of manipulator robots in situated tangible programming paradigm. In *Proceedings of the Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 319–320.
- [96] Revanth Krishna Senthilkumaran, Mridu Prashanth, Hrishikesh Viswanath, Sathvika Kotha, Kshitij Tiwari, and Aniket Bera. 2023. ARTEMIS: AI-driven robotic triage labeling and emergency medical information system. arXiv:2309.08865. Retrieved from <https://arxiv.org/abs/2309.08865>
- [97] Yang Shen, Dejun Guo, Fei Long, Luis A. Mateos, Houzhu Ding, Zhen Xiu, Randall B. Hellman, Adam King, Shixun Chen, Chengkun Zhang, et al. 2020. Robots under COVID-19 pandemic: A comprehensive survey. *IEEE Access: Practical Innovations, Open Solutions* 9 (2020), 1590–1615.
- [98] Susan Leigh Star and Anselm Strauss. 1999. Layers of silence, arenas of voice: The ecology of visible and invisible work. *Computer Supported Cooperative Work* 8 (1999), 9–30.
- [99] Samantha Stedtler. 2025. Staging, accommodating or caring: Reviewing the human labor involved in shaping robots into agents. In *Proceedings of the 2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 1650–1654.
- [100] Leila Takayama. 2011. Toward making robots invisible-in-use. In *New Frontiers in Human-Robot Interaction*, Vol. 2. Joe Saunders and Kerstin Dautenhahn (Eds.).

- [101] Angelique Taylor, Hee Rin Lee, Alyssa Kubota, and Laurel D. Riek. 2019. Coordinating clinical teams: Using robots to empower nurses to stop the line. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [102] Angelique Taylor, Sachiko Matsumoto, and Laurel D. Riek. 2020. Situating robots in the emergency department. In *Proceedings of the AAAI Spring Symposium on Applied AI in Healthcare: Safety, Community, and the Environment*.
- [103] Angelique Taylor, Michele Murakami, Soyoon Kim, Ryan Chu, and Laurel D. Riek. 2022. Hospitals of the future: Designing interactive robotic systems for resilient emergency departments. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–40.
- [104] Angelique Taylor, Tauhid Tanjim, Huajie Cao, and Hee Rin Lee. 2024. Towards collaborative crash cart robots that support clinical teamwork. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 715–724.
- [105] Angelique Taylor, Tauhid Tanjim, Michael Joseph Sack, Maia Hirsch, Kexin Cheng, Kevin Ching, Jonathan St. George, Thijs Roumen, Malte F. Jung, and Hee Rin Lee. 2025. Rapidly built medical crash cart! Lessons learned and impacts on high-stakes team collaboration in the emergency room. In *Proceedings of the 2025 ACM/IEEE International Conference on Human-Robot Interaction (HRI '25)*. IEEE Press, 501–510.
- [106] Angelique M. Taylor, Sachiko Matsumoto, Wesley Xiao, and Laurel D. Riek. 2021. Social navigation for mobile robots in the emergency department. In *Proceedings of the 2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 3510–3516.
- [107] Harold Thimbleby. 2013. Technology and the future of healthcare. *Journal of Public Health Research* 2, 3 (2013), jphr-2013.
- [108] Andrea Thomaz. 2023. Robots in real life: Putting HRI to work. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 3.
- [109] Mari Tietze and Susan McBride. 2020. *Robotics and the Impact on Nursing Practice*. American Nurses Association.
- [110] Stephen Timmons. 2003. Nurses resisting information technology. *Nursing Inquiry* 10, 4 (2003), 257–269.
- [111] Stephanie Tulk Jesso, Christopher Greene, Shiqi Zhang, Amy Booth, Michael DiFabio, Grace Babalola, Adeola Adegbemijo, and Sumantra Sarkar. 2024. On the potential for human-centered, cognitively inspired AI to bridge the gap between optimism and reality for autonomous robotics in healthcare: A respectful critique. In *Proceedings of the International Symposium on Human Factors and Ergonomics in Health Care*, Vol. 13. SAGE Publications, 106–112.
- [112] Aimee Van Wynsberghe. 2020. Designing robots for care: Care centered value-sensitive design. In *Machine Ethics and Robot Ethics*. Wendell Wallach and Peter Asaro (Eds.), Routledge, 185–211.
- [113] Tiffany C. Veinot, Jessica S. Ancker, Heather Cole-Lewis, Elizabeth D. Mynatt, Andrea G. Parker, Katie A. Siek, and Lena Mamykina. 2019. Leveling up: On the potential of upstream health informatics interventions to enhance health equity. *Medical Care* 57 (2019), S108–S114.
- [114] Katie Winkle, Donald McMillan, Maria Arnelid, Katherine Harrison, Madeline Balaam, Ericka Johnson, and Iolanda Leite. 2023. Feminist human–robot interaction: Disentangling power, principles and practice for better, more ethical HRI. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, 72–82.
- [115] Kai Weng Wong and Hadas Kress-Gazit. 2017. From high-level task specification to robot operating system (ROS) implementation. In *Proceedings of the 2017 First IEEE International Conference on Robotic Computing (IRC)*. IEEE, 188–195.
- [116] Tichakorn Wongpiromsarn, Ufuk Topcu, and Richard M. Murray. 2012. Receding horizon temporal logic planning. *IEEE Transactions on Automatic Control* 57, 11 (2012), 2817–2830.
- [117] James Wright. 2023. *Robots Won't Save Japan: An Ethnography of Eldercare Automation*. Cornell University Press.
- [118] Luhuan Yang, Yunhong Lei, Dongmei Chu, Jiawei Jiang, Zifeng Li, Yanhua Tang, and Abd Rahman Anita. 2024. Exploring antecedents and outcomes of burnout among emergency department staff using the job demands-resources model: A scoping review protocol. *PloS One* 19, 3 (2024), e0300040.

Received 2 April 2025; revised 18 December 2025; accepted 2 February 2026