

Autonomous Wheelchairs Deployment in Healthcare Facilities: Requirements and Challenges

Mingyu Guo and Weisong Shi, PhD

University of Delaware

Abstract

Deployed autonomous wheelchairs are reliable when autonomy is constrained to bounded domains and predefined destinations, but this design choice leaves key clinical requirements unmet. In health facilities, the primary barriers are not only navigation performance, but also robust interaction with building infrastructure (doors, elevators, access control), socially and operationally appropriate behavior in crowded corridors, and safety-centered fallback when perception or planning degrades. This vision paper characterizes these recurring gaps and distills a requirements agenda for next-generation autonomous wheelchairs: (1) building and workflow compatibility as first-class subsystems; (2) explicit recovery and caregiver-in-the-loop modes; (3) semantic “last-meter” goal grounding for user-referenced destinations; and (4) on-device, timing-predictable architectures aligned with privacy and scalable deployment. To make these requirements actionable, at the CAR lab at UD, we present SWee, a working prototype that emphasizes clinical deployability through an on-device, modular edge-box architecture (VOCAR) integrating onboard sensing and language-grounded goal specification.

Introduction

Autonomous wheelchairs are becoming more important as mobility needs rise while care delivery environments become more operationally constrained. In the United States, the population aged 65 and older is projected to increase from 58 million in 2022 to 82 million by 2050, expanding the number of people who can benefit from reliable mobility assistance and independence-preserving technology.¹ This shift increases routine mobility demand in health facilities and long-term care settings. Transport, escorting, and “last-meter” positioning consume staff time and compete with direct patient care.² Workforce projections further indicate persistent risks of regional and role-specific staffing shortfalls,² making it difficult to absorb growing mobility workload without increasing strain on caregivers. Observational studies also show that clinical work is frequently disrupted by “operational failures”, for example, missing supplies/equipment/information, which fragments attention.³ Together, these pressures motivate autonomous wheelchairs as a way to offload repetitive, low-acuity mobility tasks and return caregiver capacity to higher value clinical work, while preserving patient independence.¹

Real deployments also suggest that autonomy is successful when engineered as a bounded, operationally validated service rather than an open-world capability. WHILL’s commercially deployed autonomous wheelchair service in large public venues is a prominent example, where autonomy is structured around a controlled environment and defined endpoints to deliver practical independence at scale.⁴ In the home domain, DROVE represents a complementary design point: an autonomous wheelchair module that allows users to select pre-defined destinations to traverse tight corridors and doorways with reduced need for continuous fine joystick control.⁵ Together, these systems illustrate a consistent product logic: deployable

autonomy is often achieved by narrowing the operating envelope, formalizing destinations, and prioritizing reliability and repeatability over unconstrained navigation.

However, clinical environments introduce additional institutional and workflow constraints that extend beyond these bounded deployment models. Health facilities expose a gap between how autonomous wheelchairs are commonly evaluated and what deployment requires in practice. While many systems optimize nominal navigation performance, health facilities demand mobility behavior that is safe and legible in crowded and compatible with institutional norms.^{6,7} Even when local motion planning performs nominally, mobility can fail because of institutional constraints, including building access control, workflow structure, and responsibility boundaries across departments.^{8,9}

In practice, autonomous wheelchairs in health facilities encounter recurrent breakdowns that are not captured by nominal navigation benchmarks: stalled traversal at access-controlled doors, inability to summon elevators, ambiguous right-of-way negotiation in dense corridors, localization degradation during long indoor routes, and recovery procedures that require caregiver reset. These failures reveal that deployment success depends less on open-world navigation capability and more on infrastructure-aware reachability, recoverability, and predictable responsiveness.

These observations shift the correctness criterion for autonomy. A clinically useful wheelchair system must deliver mobility with limited and predictable caregiver intervention. Achieving this requires reframing core objectives toward infrastructure-constrained reachability, recoverability under degraded perception or localization, and predictable closed-loop responsiveness as onboard perception and language models scale.¹⁰ Rather than treating institutional failures as edge cases, autonomy must be evaluated against deployment-oriented requirements that determine whether it reduces caregiver burden without introducing operational overhead.

In this paper, we not only translate these deployment constraints into concrete, clinically grounded requirements for next-generation autonomous wheelchairs, but also present SWee as an exploratory prototype that operationalizes elements of this vision through an on-device, modular architecture and language-grounded goal specification. While not a completed clinical deployment, it serves as a concrete architectural instantiation of the proposed requirements and a testbed for examining deployment-oriented autonomy in practice.

Operational Models of Wheelchair Autonomy

Recent commercial autonomy for wheelchairs has converged on two deployable paradigms that trade capability for operational tractability: (i) bounded-domain destination-based navigation and (ii) route-free safety autonomy. In the bounded-domain paradigm, the system operates inside a controlled environment whose geometry, routes, and allowable endpoints are configured and maintained over time. This reduces open-world uncertainty by converting navigation into a repeatable point-to-point execution problem, localization, obstacle-aware path following, and docking within fixed operational boundaries, while shifting residual risk to environment preparation, validation, and ongoing operations.^{5,11} Consequently, the deliverable is not only an autonomy stack but an operational deployment package (mapping, endpoint definition, monitoring, and support).

DROVE instantiates bounded-domain autonomy for the home setting as an autonomous driving module that navigates between predefined destinations.⁵ At the facility scale, WHILL

Autonomous Service operationalizes the same pattern as a fleet service with a structured rollout pipeline that includes facility assessment, mapping and route configuration, system configuration, and live fleet monitoring/support. Airport deployments illustrate the service-oriented operational loop: eligible passengers request the service, transfer at designated stations, the device autonomously traverses a constrained airport segment, and then returns to base to maintain predictable fleet availability.¹² In Europe, Schiphol's continuation into a year-long trial with 10 autonomous wheelchairs starting September 2024 similarly highlights deployment as a sustained operational program rather than a one-off demonstration.¹³

In parallel, route-free safety autonomy targets everyday use without predefined routes by implementing a continuous safety envelope around user-driven motion. Rather than planning full navigation through arbitrary environments, onboard sensing and shared-control policies intervene to reduce incidents such as collisions, drop-offs (curbs/stairs), and tip events, analogous to automotive driver-assistance. This direction has proven practical because it provides immediate benefit across the environments users actually encounter (homes, malls, sidewalks, crowded indoor spaces) while remaining compatible with existing wheelchairs.^{14,15}

Autonomous wheelchair company LUCI is the canonical example of the route-free safety track, marketed as an add-on that brings “smart” capabilities to existing power wheelchairs by emphasizing safety enforcement and independence rather than destination autonomy.^{16,17} Public materials describe a multi-sensor fusion approach that combines stereo vision and radar with additional ranging modalities (e.g., infrared/ultrasonic), mounted onto third-party chairs and powered from the wheelchair battery.^{16,17} A technical case study further describes LUCI as a smart frame with embedded computer incorporating stereo vision, radar, and ultrasonic sensing to enable runtime safety interventions and connected diagnostics.¹⁸ In this framing, “smartness” is expressed as safety enforcement rather than goal navigation, enabling daily use without environment-specific configuration.^{16,18} A lighter-weight variant in the same route-free category is situational awareness augmentation: Braze Mobility positions its product as attachable ultrasonic blind-spot sensing that provides alerts (e.g., light/sound/vibration) to support safer driving while leaving full control with the user.^{19,20}

Clinical-Grade Gaps in the State of the Art

Despite steady progress in research prototypes and commercial systems, autonomous wheelchairs remain limited by recurring barriers to clinical deployment. These barriers are rarely attributable to a single module; instead, they reflect system-level mismatches between what autonomy stacks optimize in nominal settings and what healthcare facilities require in practice. We summarize the most consequential gaps below.

Infrastructure reachability. Current systems can navigate corridors yet fail at reachability interfaces like doors, elevators, access control, and choke points, where success depends on building state and facility policy as much as motion planning. Healthcare deployments report that door/elevator integration can dominate operational reliability, implying that multi-level navigation is often constrained by building cooperation rather than algorithms alone.⁸ In practice, elevators rarely expose robot-accessible APIs due to policy/compliance constraints, and doors are heterogeneous (push/pull, spring-loaded, varied hardware)²¹; secured areas further introduce badge/RFID access without a standard robot interface.

Autonomy supervision and recoverability. Clinical deployments often lack a formally specified notion of recoverability: when autonomy degrades (e.g., localization discontinuities), progress is restored through staff intervention rather than a standardized escalation-and-recovery framework. In conversations with staff involved in the ChristianaCare Moxi deployment, localization failures could lead to repeated turning-in-place and ultimately required rescue by engineers; the deployment blueprint likewise reflects routine human support via a dedicated clinical robot associate.⁸ This exposes a supervision gap: systems may not detect violated assumptions, enforce principled degraded-mode behavior, or restore task progress without routine “rescue” as an implicit subsystem.

Goal expressiveness and map fragility. Many deployed systems (e.g., WHILL and DROVE) primarily support geofenced routes and predefined destinations rather than context-based requests such as “go to the orange chair next to the door.” However, users often need both high-level routing and low-level positioning near specific objects and landmarks; destination-only navigation is frequently insufficient because the destination is not the objective. In facilities with frequent layout changes and moving equipment,²² autonomy that depends on high-resolution maps and curated routes can degrade without continual re-mapping and operational maintenance.¹¹

Assistive task execution and embodied interaction. Most smart wheelchairs focus on collision avoidance and basic navigation rather than end-to-end assistance with reaching, grasping, and manipulation, leaving users dependent on caregivers for routine actions.^{11,16} As a system-level gap, assistance is rarely treated as integrated mobile manipulation: manipulation depends on approach pose, clearance, and viewpoint, while navigation must account for reachability and collision envelopes. As a result, task completion becomes a coupled mobility–manipulation problem in cluttered clinical environments.

Long-tail scene understanding. Although platforms can detect people and obstacles, they often fail to infer context-dependent priorities (e.g., yielding to emergency traffic or urgent equipment movement) in crowded clinical spaces, producing behavior that is physically safe but socially unsafe or operationally disruptive.⁶ Clinically salient events are rare but high consequence, yet current models and evaluations often under-represent these operational norms and right-of-way semantics.

Timing predictability. A key system gap is the timing predictability of the integrated sense–think–act loop.¹⁰ Few platforms bound end-to-end latency and jitter across the pipeline, and the problem worsens when adding heavier models for semantic understanding (e.g., VLMs/VLAs). Cloud offloading can introduce delay under congestion and raise privacy/robustness concerns in sensitive spaces.¹⁰ Consequently, increasing semantic capability can compromise responsiveness, while average-case metrics can obscure tail-latency behavior and overload failures.

Deployment scalability. Sustained deployment requires modular hardware and software, plus operator-facing diagnostics that turn failures into actionable maintenance steps. The system should provide health indicators for sensing, localization, docking/charging, and compute performance. It should support remote triage and structured logs that localize recurring issues to specific locations or infrastructure interfaces. Deployment evaluation should include intervention frequency, time-to-recovery, and maintenance effort over a multi-week operation, rather than single-run success.

Clinical workflow constraints. Clinical transport requires explicit authority and workflow constraints. A patient may request “take me back to my room” while the wheelchair is assigned to supervised transport for imaging; without an authority hierarchy, the system could comply and disrupt care. A clinically integrated authority model allows staff commands to override patient-initiated navigation during active transport episodes while preserving patient autonomy outside those contexts.⁸ As wheelchairs evolve into assistive robotic platforms, constraints must also respect clinical policies and orders such as NPO restrictions; without formal integration, the system may execute commands that conflict with care plans.

Taken together, these gaps reflect three underlying system-level limitations: (i) insufficient cross-layer integration of infrastructure state, user intent, and workflow constraints into a unified operational context; (ii) weak autonomy supervision, including principled recovery and authority arbitration under degraded conditions; and (iii) lack of predictable closed-loop performance as semantic capability increases. Addressing these limitations requires integrated, clinically grounded system design rather than component-level autonomy improvements.

Clinical Grade Requirements and Research Agenda

Next-generation autonomous wheelchairs should be designed for clinical realism and scalable deployment, not only best-case lab autonomy. Based on the system-level gaps above, we view the next generation as requiring (i) *infrastructure-constrained reachability*, (ii) *recoverable autonomy with explicit escalation*, and (iii) *predictable closed-loop performance* as onboard models scale. Figure 1 summarizes our system view.

Required Capabilities

These requirements are not feature additions but architectural properties that must be designed into the autonomy stack from the outset. Each addresses a deployment-critical failure mode observed in clinical environments and defines a system-level constraint that shapes perception, planning, control, and interaction.

Infrastructure-Constrained Reachability and Context-Grounded Goals

Clinical mobility fails most often at infrastructure interfaces and at the semantic boundary between “destination reached” and “task accomplished.” This category therefore focuses on physical reachability and goal grounding within live, policy-constrained environments.

Infrastructure-compatible reachability. Next-generation autonomy should treat doors, elevators, and access control as part of the navigation problem rather than as add-on integrations. Route planning should explicitly represent the required infrastructure interactions along each candidate path and maintain a runtime estimate of whether those interactions are feasible under current facility constraints. The system should accommodate common door behaviors in hospitals (e.g., push/pull operation, spring-loaded closure, variable hold-open timing) and should not assume that elevators expose robot-accessible APIs. When digital interfaces are unavailable, robotic-arm actuation is one optional mechanism for physical interaction (e.g., button pressing, door operation), but it must operate under bounded force/torque and contact-safety constraints, and provide simple failure recovery such as bounded retry or escalation to staff.

Expressive goals beyond destinations. The wheelchair should support goal specifications beyond predefined endpoints, including context-grounded positioning that depends on approach

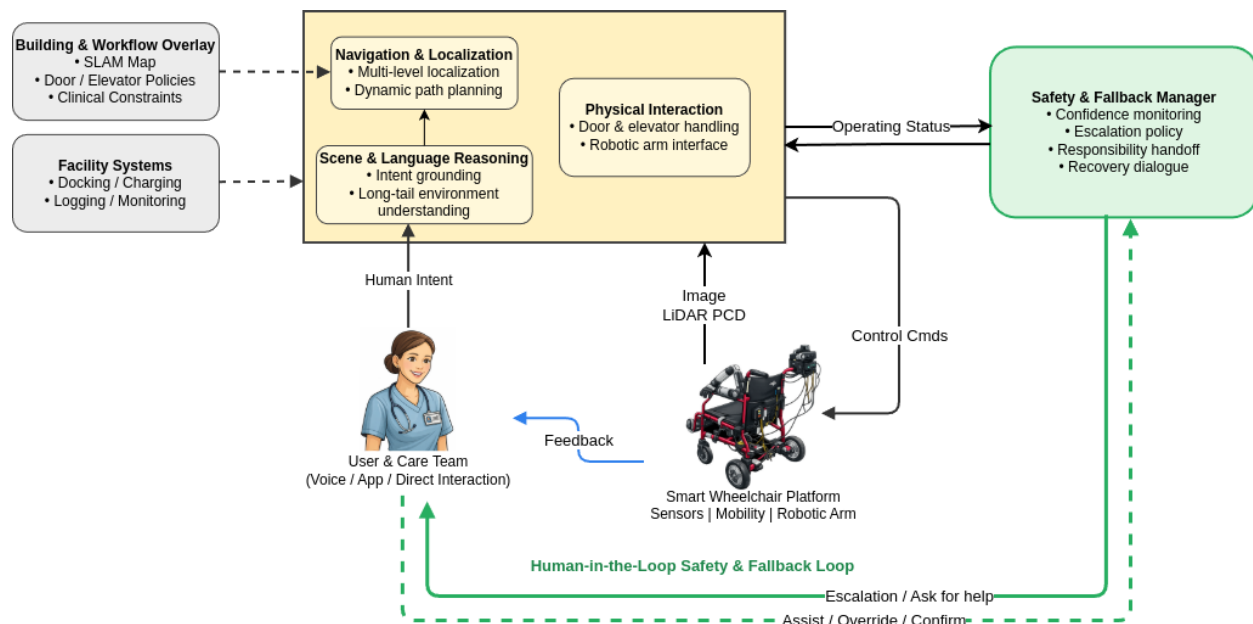
and stance relative to the environment. The navigation stack should therefore couple semantic grounding with local motion generation to produce reliable last-meter behavior near people, objects, and landmarks. Robustness should not depend on static, curated maps: the system should tolerate nonstationary layouts and partial/stale spatial context by relying on online perception and local refinement where needed.

Recoverable Autonomy with Explicit Escalation and Workflow Awareness

Even with correct reachability modeling, clinical deployment requires autonomy that degrades predictably and restores progress systematically. This category addresses supervision, authority, and recovery under real-world variability.

Supervised autonomy with defined fallback and recovery. Clinical operation requires an explicit supervision model that specifies how the system detects degraded autonomy, what behaviors are allowed under degradation, and how task progress is restored. Degradation should be triggered by measurable signals such as localization inconsistency, repeated turning-in-place, persistent planner infeasibility, or prolonged lack of progress. Under these conditions, the platform should enter bounded behaviors that prioritize predictability, and it should expose recovery options that nearby non-expert staff can execute. During recovery, the system should provide a brief, easy-to-understand explanation of why progress has stalled so a rescuer can intervene efficiently, translating internal failure states into plain-language causes, for example, “localization lost,” “path blocked,” etc.

Figure 1. Closed-Loop Clinical Autonomy Pipeline for a Next-Generation Wheelchair



Human intent and onboard sensing are processed on an edge onboard computing stack that combines scene–language reasoning with navigation and multi-level localization to generate mobility control commands.

Workflow and authority-aware operation. Command interpretation should be conditioned on the care context so that behavior matches clinical workflows. The autonomy stack should represent transport episodes and authority relationships among the patient, assigned staff, and facility operators, and use these structures to resolve conflicting directives. This includes the

ability to defer, reject, or request confirmation for commands that conflict with an active transport episode or facility policy, while preserving patient autonomy outside supervised contexts. The same conditioning should apply to assistive actions, not only navigation, so task execution respects operational constraints.

Task-level assistance via integrated mobile manipulation. The system should extend beyond transportation to task-level assistance through integrated mobile manipulation, where base motion and interaction are planned jointly. The platform should support task execution that depends on approach pose, clearance, and viewpoint, and should explicitly handle the coupling between navigation uncertainty and interaction feasibility. Capability should be evaluated by end-to-end task completion and robustness in realistic scenes, rather than isolated grasp or navigation success.

Predictable Closed-Loop Performance and Scalable Deployment

Clinical autonomy must remain stable as semantic capability increases and as systems operate over weeks or months. This category therefore addresses timing guarantees, edge execution, and long-term operational maintainability.

Edge-resident autonomy. Core perception, intent interpretation, and decision-making should run on-device to preserve privacy and avoid dependence on variable connectivity. The platform should be architected so that safety-critical closed-loop behavior remains functional under connectivity loss, while optional services degrade gracefully without destabilizing control. On-device operation should also constrain data handling: sensitive sensor streams should remain local by default, with explicit policy governing any off-device transmission.

Predictable timing under multi-model workloads. The platform should provide predictable end-to-end responsiveness when multiple perception and reasoning workloads co-run on embedded hardware. This requires system-level characterization and control of latency and jitter across the sense–think–act loop, including tail behavior under contention. The autonomy stack should preserve a safety-critical control path whose responsiveness is not compromised by best-effort semantic inference, enabling safe behavior under overload rather than silent timing collapse.

Maintainability and operational scalability. The autonomy package should be modular, serviceable, and upgradable without constant on-site engineering. Deployment should rely on a standardized configuration and validation workflow, where site-specific adaptation is limited to lightweight parameterization (e.g., infrastructure interaction templates, calibration, and policy constraints) rather than bespoke integration. Evaluation should treat intervention burden, reconfiguration effort, and time-to-operability over weeks to months as deployment-facing performance metrics.

Long-tail understanding and legible interaction. The wheelchair should support context-dependent behavior in socially structured spaces, including behavior that is operationally appropriate and interpretable to nearby people. The interface should make state and intent legible through concise cues (status, rationale for pauses, and next action), and should support brief dialogue when needed. For accessibility, the system should support assistive description features that improve situational awareness for blind and low-vision users.

Research Questions

The capabilities above motivate six research questions spanning autonomy, systems, and clinical integration: (RQ1) how to represent access feasibility and policy constraints in a form that can be optimized and validated at the episode level; (RQ2) how to define degraded operation signals and recovery policies that restore progress without deadlock, and how to benchmark intervention burden; (RQ3) how to encode clinical context and authority hierarchies for conflict-aware command interpretation with workflow-level correctness metrics; (RQ4) how to couple semantic grounding with local navigation when spatial context is partial or drifting; (RQ5) how to bound latency/jitter under multi-model contention on embedded hardware and report tail behavior under overload; and (RQ6) how to unify mobility and manipulation into end-to-end task policies with evaluation tied to caregiver-burden reduction.

The SWee Prototype

We introduce Smart Wheelchair (SWee), a prototype designed by the CAR Lab at the University of Delaware to operationalize the deployment-oriented requirements above. SWee emphasizes (i) edge-resident autonomy suitable for safety-critical indoor operation, (ii) goal expressiveness through visually grounded user commands, and (iii) a modular form factor that supports iteration and maintainability. While several capabilities remain under active development, the current system serves as an integrated testbed for studying clinically realistic mobility under resource and workflow constraints.

Current Prototype Capabilities


At a systems level, SWee executes the autonomy stack locally on the Voice-Controlled Autonomous Robot Kit (VOCAR),²³ avoiding reliance on external servers and reducing privacy exposure in clinical spaces. VOCAR integrates a smart microphone, a Jetson Nano computing module,²⁴ and a perception suite consisting of a Unitree L1 3D LiDAR²⁵ and an Intel RealSense D435 depth camera.²⁶ The module is mounted above the user to increase the field-of-view and reduce near-field occlusion, providing consistent sensing geometry during motion. The prototype configuration is shown in Figure 2.

SWee supports operation in dynamic indoor environments without requiring pre-defined high-resolution maps or geofenced routes for fine-grained navigation. For long-range travel to known destinations, SWee can use a lightweight global map to plan at the corridor-and-room level. However, SWee does not depend on that map for local execution: it can navigate in a map-free mode using onboard LiDAR and camera sensing to build a local occupancy map online and generate collision-free motion. This enables fine-grained, context-grounded goals. By fusing object detection, depth estimation, and a VLM, SWee can interpret requests such as “take me to the vacant chair on the right side of the bed” and translate them into actionable local navigation targets. This framing aligns with prior work on language-mediated semantic representations and dialog-grounded goal specification for assistive mobility.²⁷⁻²⁹

SWee is a multi-module autonomy stack that includes voice recognition, object detection, navigation, and other co-running components on resource-constrained embedded hardware. Because many of these components are GPU-heavy DNN workloads, uncontrolled concurrency can create compute contention and reduce end-to-end predictability. To keep behavior responsive while using heavyweight perception and language components, SWee separates decision-making

into two pipelines with different real-time requirements.³⁰ A low-frequency grounding pipeline processes the user's command together with camera frames to identify and localize the referenced destination, running on demand because it only needs to produce target updates when the user issues a command or re-grounding is required. A high-frequency tracking and safety pipeline then performs obstacle detection and target tracking in real time at 30 Hz by fusing YOLOv8 visual detections³¹ with 3D LiDAR point clouds. Once the target is established, this tracking layer maintains a continuously updated target position and safe motion behavior, even when the target moves, without requiring repeated heavyweight VLM inference. This decoupling keeps heavyweight reasoning off the tight control loop and supports more predictable reactions during motion.

Figure 2: Smart Wheelchair Prototype Setup



SWee's map-free operation mode reduces the need for site-specific route curation in clinical settings where layouts change frequently. The VOCAR edge-box form factor also supports a modular deployment model compared to fully integrated wheelchair products, enabling the autonomy stack to be updated and serviced as an attachable unit. This simplifies maintenance workflows: when a sensing component fails or calibration degrades, the autonomy module can be swapped or serviced without requiring invasive modification of the base wheelchair. SWee is also designed to make environment setup more accessible by allowing a floor to be mapped through simple driving, without requiring specialized robotics expertise, and then reused across a fleet.

Developing Capabilities

SWee is currently developing task-level assistance through integration of a robotic arm to support actions such as picking up dropped objects, retrieving items from different heights, opening doors, and pressing elevator buttons. In parallel, SWee's longer-term direction includes richer long-tail scene understanding and user-relevant interaction, consistent with broader smart wheelchair research emphasizing semantic understanding and human-centered interaction beyond geometric navigation.¹¹ The goal is to move beyond purely geometric navigation toward

context-aware behaviors and more human-like assistance, including spoken descriptions and conversational interaction when appropriate, particularly for blind and low-vision users.

Conclusion

In conclusion, the central lesson from prior deployments and research prototypes is that autonomy becomes clinically meaningful only when it is designed around the realities that health facilities impose. This paper therefore frames the problem as a requirements agenda: we identify the recurring gaps that prevent reliable adoption—fragile interaction with doors and elevators, weak accommodation of workflow and shared-space norms, limited semantic “last meter” goal execution, and the absence of explicit recovery and escalation behavior when autonomy degrades and argue that these are not peripheral details but the main determinants of safety and usability. We translate these gaps into concrete must-haves for next-generation systems, emphasizing building-compatibility as a first-class subsystem, formally engineered fallback and caregiver-in-the-loop pathways, workflow-aware authority and constraints, and on-device, timing-predictable operation appropriate for safety-critical indoor environments. Finally, we present a working prototype as a test bed to make these requirements actionable: it provides a platform for iterating on clinically grounded behaviors, validating design trade-offs under realistic constraints, and enabling future evaluation that moves beyond route-only autonomy toward deployable assistive mobility.

The authors may be contacted at myuguo@udel.edu and weisong@udel.edu.

References

1. Mender. (n.d.). *LUCI case study*. Retrieved from <https://mender.io/resources/case-studies/luci-case-study>
2. Zachariae, A., Plahl, F., Tang, Y., Mamaev, I., Hein, B., & Wurl, C. (2024). Human-robot interactions in autonomous hospital transports. *Robotics and Autonomous Systems*, 179, 104755. <https://doi.org/10.1016/j.robot.2024.104755>
3. Zhong, R., Tian, Z., Wang, Q., Guo, M., & Shi, W. (2025). Design and implementation of a voice controlled indoor autonomous robot kit. In *2025 IEEE 3rd International Conference on Mobility, Operations, Services and Technologies (MOST)* (pp. 89–99). IEEE. <https://ieeexplore.ieee.org/document/11071445>
4. Intel. (2018). *Intel RealSense D400 series: Datasheet*. Retrieved February 5, 2026, from <https://cdrdv2-public.intel.com/841984/Intel-RealSense-D400-Series-Datasheet.pdf>
5. Control Bionics. (n.d.). *DROVE*. Retrieved from <https://www.controlbionics.com/products/drove/>
6. ChristianaCare. (n.d.). *Blueprint for charting nurse-led innovation: Lessons from deploying collaborative robots in a complex health system*. Retrieved from <https://christianacare.org/us/en/for-health-professionals/nursing/nurse-led-innovation>
7. Tucker, A. L., & Spear, S. J. (2006, June). Operational failures and interruptions in hospital nursing. *Health Services Research*, 41(3p1), 643–662. <https://doi.org/https://doi.org/10.1111/j.1475-6773.2006.00502.x>

8. WHILL, Inc. (2022). *WHILL announces autonomous mobility service in North America*. Retrieved from <https://whill.inc/us/whill-announces-autonomous-mobility-service-in-north-america>
9. Omer, K., & Monteriù, A. (2025, July 31). Multi-layer robotic controller for enhancing the safety of mobile robot navigation in human-centered indoor environments. *Frontiers in Robotics and AI*, 12, 1629931. <https://doi.org/10.3389/frobt.2025.1629931> PubMed
10. Bernhard, L., Schwingenschlögl, P., Hofman, J., Wilhelm, D., & Knoll, A. (2024). Boosting the hospital by integrating mobile robotic assistance systems: A comprehensive classification of the risks to be addressed. *Autonomous Robots*, 48(1), 1. <https://doi.org/10.1007/s10514-023-10154-0>
11. Schwartz, D., Kondo, K., & How, J. P. (2025). *Efficient navigation in unknown indoor environments with vision-language models*. arXiv (arXiv:2510.04991). <https://arxiv.org/abs/2510.04991>
12. Unitree Robotics. (2023). *Unitree 4D LiDAR-L1: Product specifications*. Retrieved from <https://oss-global-cdn.unitree.com/static/52b72f707b304d229d4321eea223738f.pdf>
13. Che, K., Okamura, A. M., & Sadigh, D. (2020). Efficient and trustworthy social navigation via explicit and implicit robot–human communication. [IROS]. *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 36(3), 692–707. Retrieved from <https://ieeexplore.ieee.org/document/8967120> <https://doi.org/10.1109/TRO.2020.2964824>
14. WHILL, Inc. (2025). *Unifi Aviation launches WHILL autonomous wheelchair service at Detroit Metro Airport*. Retrieved from <https://whill.inc/us/unifi-aviation-launches-whill-autonomous-wheelchair-service-at-detroit-metro-airport/>
15. Patki, S., Fahnestock, E., Howard, T. M., & Walter, M. R. (2019). Language-guided semantic mapping and mobile manipulation in partially observable environments. In *Proceedings of the Conference on Robot Learning (CoRL)*. Retrieved from <https://ttic.edu/ripl/assets/publications/patki19a.pdf>
16. Future Travel Experience. (2024). *Schiphol testing 10 autonomous wheelchairs from WHILL as part of ongoing efforts to make the airport accessible to everyone*. Retrieved from <https://www.futuretravelexperience.com/2024/09/schiphol-testing-10-autonomous-wheelchairs-from-whill-as-part-of-ongoing-efforts-to-make-the-airport-accessible-to-everyone/>
17. Braze Mobility. (n.d.). *Braze Mobility (Blind-spot and proximity sensing for wheelchairs)*. Retrieved from <https://brazemobility.com/>
18. LoPresti, E. F., Sharma, V., Simpson, R. C., & Mostowy, L. C. (2011). Performance testing of collision-avoidance system for power wheelchairs. *Journal of Rehabilitation Research and Development*, 48(5), 529–544. <https://doi.org/10.1682/JRRD.2010.01.0008> PubMed
19. Health Resources and Services Administration. (2025). *Nurse workforce projections, 2023–2038 (Factsheet)*. Retrieved from <https://bhw.hrsa.gov/sites/default/files/bureau-health-workforce/data-research/nursing-projections-factsheet.pdf>

20. Gao, Y., & Huang, C.-M. (2022, January 12). Evaluation of socially-aware robot navigation. *Frontiers in Robotics and AI*, 8, 721317. Retrieved from <https://pmc.ncbi.nlm.nih.gov/articles/PMC8791647/>
<https://doi.org/10.3389/frobt.2021.721317> PubMed
21. Fierce Electronics. (2025). *Blind spot sensors make wheelchairs smarter, safer at Braze Mobility*. Retrieved from <https://www.fiercesensors.com/sensors/blind-spot-sensors-make-wheelchairs-smarter-safer-braze-mobility>
22. Population Reference Bureau. (2024). *Fact sheet: Aging in the United States*. Retrieved from <https://www.prb.org/resources/fact-sheet-aging-in-the-united-states/>
23. Yaseen, M. (2024). *What is YOLOv8: An in-depth exploration of the internal features of the YOLOv8 model*. arXiv (arXiv:2408.15857). <https://arxiv.org/abs/2408.15857>
24. Xi, L., & Shino, M. (2020). Shared control of an electric wheelchair considering environment information. *International Journal of Environmental Research and Public Health*, 17(15), 5502. Retrieved from <https://pmc.ncbi.nlm.nih.gov/articles/PMC7432419/>
<https://doi.org/10.3390/ijerph17155502> PubMed
25. NVIDIA. (2019). *Jetson Nano: Technical specifications*. Retrieved February 5, 2026, from <https://developer.nvidia.com/embedded/jetson-nano>
26. Hemachandra, S., & Walter, M. R. (2014). Learning semantic maps through dialog for a voice-commandable wheelchair. *Robotics: Science and Systems (RSS) Workshop / Technical Report version*. Retrieved February 5, 2026, from <https://ttic.edu/ripl/assets/publications/hemachandra14a.pdf>
27. Shi, W., Dong, Z., & Zhou, P. (2026). Physical intelligence on the edge: A vision for the decade ahead. *Journal of Computer Science and Technology*, 1–16.
<https://doi.org/10.1007/s11390-026-6292-8>
28. Protolabs. (n.d.). *LUCI wheelchair (Partnership story)*. Retrieved from <https://www.protolabs.com/resources/partnerships/luci-wheelchair/>
29. Kim, Y., Velamala, B., Choi, Y., Kim, Y., Kim, H., Kulkarni, N., & Lee, E.-J. (2023). *A literature review on the smart wheelchair systems*. arXiv (arXiv:2312.01285).
<https://arxiv.org/abs/2312.01285>
30. LUCI. (2020). *LUCI adds collision avoidance and anti-tip tech to powered wheelchairs*. Retrieved from <https://luci.com/2020/12/luci-adds-collision-avoidance-and-anti-tip-tech-to-powered-wheelchairs/>
31. Huang, C., Mees, O., Zeng, A., & Burgard, W. (2022). *Visual language maps for robot navigation*. arXiv (arXiv:2210.05714). <https://arxiv.org/abs/2210.05714>

Copyright (c) 2026 Delaware Academy of Medicine and Public Health.

This is an Open Access article distributed under the terms of the Creative Commons Attribution Non-Commercial License (<https://creativecommons.org/licenses/by-nc-nd/4.0/>) which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.