2026 IEEE International Conference on
Advanced Robotics and Its Social Impacts (ARSO)
Vienna, Austria. June 10th-12th 2026

Human Recognition–Driven Interaction with Hello Robot Stretch for Supporting Older Adults

Aakriti Upadhyay, Allen Marshall, Rose Gamble
Sarah and John Graves Institute for Robotics and Autonomy
The University of Tulsa
800 S. Tucker Dr., Tulsa, OK 74104
Email: {aau5917, allen-marshall, gamble}@utulsa.edu

Abstract— **Supporting older adults in aging safely and independently at home is a key application area for service and assistive robotics. While commercial mobile manipulators provide a promising foundation, their ability to autonomously perceive and locate humans in unconstrained indoor environments, particularly for floor-level detection, remains an underexplored engineering challenge. In this work, we present a unified framework for autonomous monitoring and person-finding using the Hello Robot Stretch 3. We evaluate the robot's autonomous search and perception performance in distinguishing a person lying on the floor from a person who is seated or standing in its frontal view, enabling appropriate initiation of interaction or assistance. Experiments are conducted in a realistic, cluttered indoor environment show 90–100% success rate with random robot starting positions. The results demonstrate the feasibility of using the Stretch 3 for assistive service tasks such as fall detection and in-home user monitoring, while also revealing perception challenges in cluttered scenes that must be addressed for robust real-world deployment. Our findings provide practical insights into leveraging commercial service robots for assistive aging-in-place applications.**

I. Introduction

Robotic systems have transitioned from industrial automation to mobile platforms capable of navigating complex human environments. Within the domain of socially assistive robotics (SAR), there is a critical need for systems that provide autonomous monitoring to support the safety of older adults living independently. While the broader field of assistive robotics encompasses physical rehabilitation and functional aid, a foundational requirement for any such system is the ability to reliably locate and assess the status of a human inhabitant without constant caregiver intervention [1], [2], [3].

The global increase in the aging population has led to a rise in age-related mobility impairments and balance issues, significantly elevating the risk of domestic falls [4]. In-home supervision is often intermittent, leaving *silent* windows where emergencies may go undetected. While existing monitoring solutions often rely on static sensors [5] or wearables [6], which are limited by field-of-view and user compliance, mobile monitoring robots offer a dynamic alternative [7]. However, for a mobile robot to be effective, it must do more than move; it must intelligently explore a home to *find* a person who may be in non-standard poses (e.g., lying on the floor after a fall).

Current engineering challenges in domestic monitoring involve the trade-off between exploration efficiency and detection accuracy. Many existing frameworks assume the human is always within a frontal camera view or in an upright position [8]. There remains a gap in integrated software models that allow a mobile platform to autonomously perform frontier-driven exploration [9] specifically optimized for human search-and-rescue within the constraints of a domestic layout.

In this work, we address this gap by presenting a targeted monitoring framework implemented on the Hello Robot Stretch 3. Moving beyond general assistive claims, this study focuses strictly on the technical integration of active perception and navigation for autonomous person-finding. We present an assistive robotic software model for monitoring older adults in home-like cluttered environment while the robot actively explores and interacts with its surroundings (Figure 1). The key contribution of this paper is a unified framework that enables the robot to:

- Execute frontier-driven exploration to systematically search indoor environments.
- Utilize active head motion and confidence-based detection to distinguish between standing poses and floor-level poses (indicative of falls).
- Initiate a navigation and interaction protocol once the person is located.

By focusing on the monitoring and search phase of assistive care, this study evaluates how the Stretch's unique hardware configuration can be leveraged to solve the *hidden person* problem in domestic safety.

II. Related Work

Researchers have increasingly explored the use of mobile robots for autonomous monitoring and domestic safety to support day-to-day living of older adults. In this section, we examine state-of-art approaches to perception-driven exploration and socially assistive monitoring, focusing on how mobile platforms address the limitations of traditional, static monitoring systems in identifying object-of-interest.

The Hello Robot Stretch 3 has emerged as a platform for indoor, human-centered research due to its compact footprint and compliant manipulator design [10]. While large-scale industrial robots present significant safety challenges in home settings and small vacuum robots are constrained by *blind* sensing and limited mobility, the Stretch 3 offers a type of

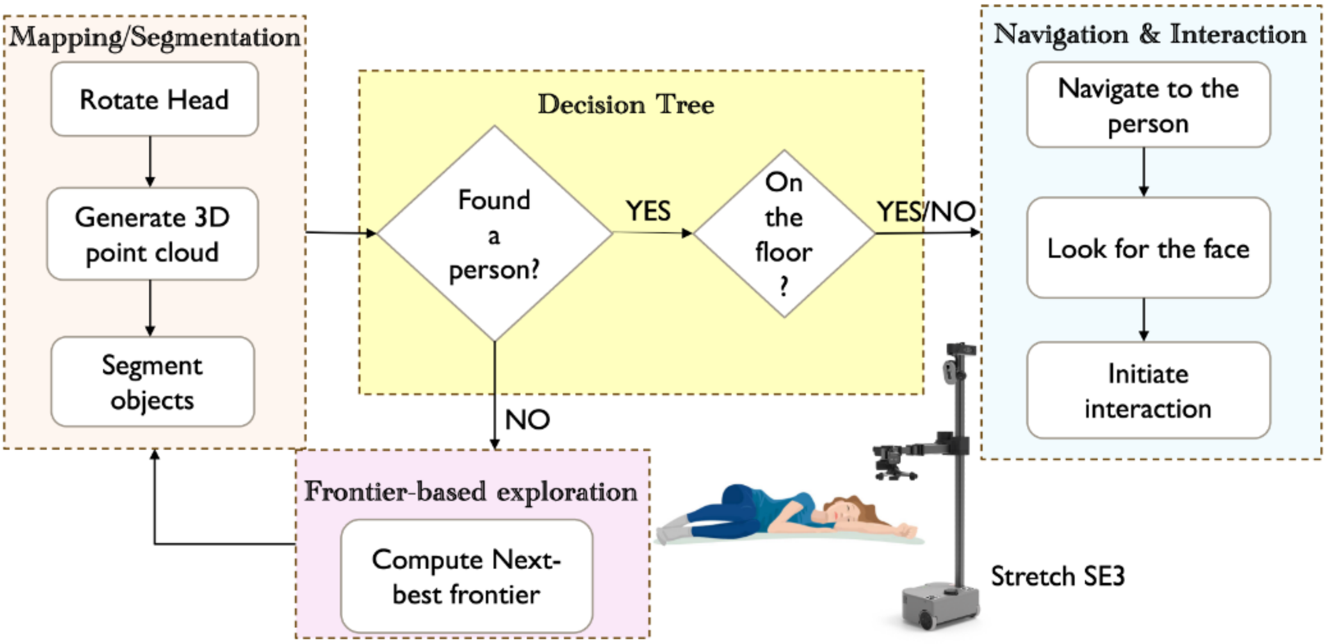

Fig. 1. The figure illustrates the different phases of our framework. In the Mapping/Segmentation phase, the robot rotates its head to scan the environment and identify objects. The Decision Tree phase determines whether a person is detected and verifies their position (on the floor or in front of the robot). Once confirmed, the Navigation & Interaction phase enables the robot to move toward the user and initiate assistive verbal communication. If no person is detected, the robot explores a new frontier location near the map boundary.

middle ground. Previous deployments have predominantly explored its utility in discrete, task-oriented scenarios. For instance, researchers have demonstrated its potential for object manipulation [11] and social interaction [12], including letter writing [13] and proxemic-aware delivery tasks [14]. While these studies establish the robot's hardware viability, they typically rely on either manual teleoperation interfaces [15] or predefined interaction goals that assume a known, task-centric environment.

In contrast, our work shifts the focus from task execution to environment awareness. While existing literature validates the robot's ability to manipulate objects, there remains a critical gap in utilizing its mobility for continuous, autonomous monitoring. By integrating advanced perception with frontier-driven exploration, our approach enables the system to transition from a teleoperated or task-specific tool into an autonomous agent capable of detecting and monitoring users in unconstrained domestic settings.

A. Semantic-aware Active Exploration

Autonomous object-centric exploration has emerged as a significant research direction, particularly for indoor inspection and semantic mapping [16], [17]. While traditional exploration focuses on maximizing map coverage, semantic-aware exploration prioritizes regions of interest based on the likelihood of finding target objects [8], [18]. Current state-of-the-art approaches across aerial and ground platforms often rely on Rapidly-exploring Random Trees (RRTs) for path generation [19], [8] and leverage pretrained 2D neural networks (e.g., You Only Look Once (YOLO) or Segment Anything) to project semantic labels into 3D space via LiDAR or stereo vision. These systems typically optimize motion for efficiency and safety, using multi-phase planning to transition between global exploration and local search.

Despite these advancements, standard YOLO-based detection and RRT-based planning are generally optimized for upright human subjects, leaving viewpoint optimization for floor-level inhabitants, such as prone or supine fall victims, largely unexplored. Our work addresses this limitation by modifying the active perception loop to optimize both the robot's path and camera pitch for *atypical* human signatures. By adapting the exploration strategy to distinguish between standard interactions and emergency floor-level poses in cluttered environments, we target a critical gap in autonomous safety monitoring.

B. Fall Detection and Monitoring Systems

Research in geriatric care has increasingly turned to context-aware sensors [20] and wearable devices [21] to mitigate the risks of undetected falls. Ambient systems (e.g., acoustic or camera-based) provide non-intrusive monitoring but are inherently limited by fixed fields-of-view and environmental occlusions, creating blind spots where emergencies may go unnoticed [22]. In contrast, wearable devices incorporating accelerometers and gyroscopes enable continuous tracking but often suffer from low user compliance; older adults may forget to wear or charge them, reducing reliability during critical events [23].

Mobile robotic platforms, such as Hobbit [7], have been proposed to address these monitoring gaps through dynamic observations. However, many existing socially assistive robotics (SAR) frameworks emphasize social interaction or high-level activity recognition, often assuming that users remain within sensing range and in upright positions [24]. This leaves an important technical challenge unaddressed: how a robot should autonomously navigate and orient its sensors to locate a person who is both out of view and potentially in a low-visibility, floor-level pose.

To overcome this limitation, our approach incorporates pose-agnostic detection through active, low-angle perception. Unlike prior systems constrained by fixed-tilt sensors or static bases, it leverages the extensible hardware of the Stretch 3 to systematically sweep domestic environments, targeting the detection of non-standard poses in cluttered spaces.

III. Stretch Configurations

Our approach uses the Stretch SE3 mobile manipulator developed by Hello Robot [10] to investigate the potential for commercial robot usage in this domain. Stretch SE3 has a telescoping arm that extends 52 cm horizontally and is attached to a prismatic lift that reaches 110 cm vertically. The arm has 3 degrees of freedom (DoFs) dexterous wrist and 1 DOF gripper. The movement of the arm is orthogonal to the movement of the differential drive base. Stretch has three cameras: a realsense D435if camera and wide-angle camera attached to a pan-tilt head, and a realsense D405 camera mounted on the gripper. The limits of movement for each joint and their corresponding units are reported in Table I. The mobile base has 3 DOFs (planar rotational) and a 2D lidar (Slamtec RPLIDAR A1) mounted on it. This lidar has a range of 0.15-12 m and an angular resolution of 1 degree. Furthermore, Stretch SE3 is interoperable with ROS2 (Robot Operating System).

TABLE I
STRETCH JOINT LIMITS

Joints	Type	Bounds		Unit
		Lower	Upper	
Head Pan	R	-4.04	1.72	rad
Head Tilt	R	-2.03	0.47	rad
Lift	P	0	1.1	m
Arm	P*	0	0.52	m
Wrist Roll	R	-2.91	2.91	rad
Wrist Pitch	R	-1.57	0.45	rad
Wrist Yaw	R	-1.28	4.49	rad
Gripper	R	-0.13	0.23	rad

* R: Revolute; P: Prismatic.

IV. Methodology

The objective of this work is to evaluate whether the Stretch robot can detect a person in its surroundings and determine whether the person is lying on the floor (for example, due to a fall or dizziness) or positioned in front of the robot so that it can initiate communication. In Algorithm 1, we design a human detection and navigation framework that generates a three-dimensional (3D) point cloud of the surrounding floor area during the robot's head motion. This point cloud is represented as a voxel grid map and is analyzed to detect the presence of a person in the space (lines 5–7). The method operates under two main scenarios based on the orientation of the robot's head-mounted camera (lines 8-16).

*a) **Detection of a Person in Frontal View**:* In the first scenario, the robot assumes that a person may be present in the frontal field of view, such as a standing or seated individual. The head-mounted camera rotates horizontally while being tilted 30° downward.

*b) **Detection of a Person on the Floor**:* If no person is detected in the first scenario, the robot switches to the second scenario. In this case, the head-mounted camera is tilted an additional 30° downward to search for a person lying on the ground.

In both scenarios, a detection is considered successful only when the confidence score exceeds 60%, a threshold chosen to be clearly above the 50% random-chance level and to avoid uncertain detections. Scores near 50% indicate ambiguity and are insufficient for triggering robot navigation and human interaction. By requiring greater than 60% confidence, the system reduces false positives and ensures that the robot approaches and communicates only when a human is detected with sufficient reliability.

Algorithm 1: Human Detection and Navigation

Input: F: set of feasible frontier poses, d_t: minimum distance threshold, P_{robot}: robot's current position, V: set of visited frontier poses, P_{human}: detected human position, O: set of detected objects, M: environment map

1 $V \leftarrow \{\phi\}$
2 $F \leftarrow F \cup P_{robot}$
3 $success \leftarrow False$
4 **while** $F \neq \{\phi\}$ **do**
5 $P_{human} \leftarrow NULL$
6 $HeadScan(M)$ // Generates the point-cloud map
7 $IdentifyObj(O)$
8 **if** *person* $\in O$ **then**
9 $P_{human} \leftarrow GetPosition(M)$ // Gets human position on the map
10 $PlanPath(P_{robot}, P_{human})$ // Navigates to the person
11 $V \leftarrow V \cup P_{robot}$
12 $F \leftarrow F \setminus P_{robot}$
13 **if** *Face Detected* **then**
14 $success \leftarrow True$ // person is in frontal view
15 **else if** *person on the floor* **then**
16 $success \leftarrow True$
17 $N \leftarrow FrontierPoses(P_{robot})$ // Find new frontiers from current pose
18 **foreach** $p \in N$ **do**
19 **if** $p \notin V$ ***and*** $\|p, \{\forall v \in V\}\| > d_t$ **then**
20 $T \leftarrow T \cup Score(p)$ // Ref. equation 1
21 $F \leftarrow F \cup SelectNxtPose(T)$ // Gets the next best pose
22 $N \leftarrow \phi$
23 $T \leftarrow \phi$
24 **return** *success* /* Initiates interaction on True */

A. Frontier Scoring and Parameterization

Each scenario includes a person detection task. If the task fails in both scenarios, meaning that no person is detected from the robot's current position, the robot selects a new search location. This location is chosen from frontier positions near unexplored boundaries of the environment map, lines 17-23 in Algo. 1.

The selected frontier position must satisfy the following conditions:

- It is not too close to the robot's current position.
- It is sufficiently far from previously visited locations.
- It maintains a safe clearance from obstacles.

The best frontier position is selected using the following scoring function:

$$score = \alpha \cdot d_{robot} + \beta \cdot d_{visited} - \gamma \cdot d_{obs}, \tag{1}$$

where d_{robot} represents the distance from the robot's current position, $d_{visited}$ is the distance from previously visited positions, and d_{obs} is the distance from nearby obstacles. The constants α, β, and γ control the relative importance of each term.

Given the robot's maximum sensing range R, candidate frontiers are generated within this range and all distance terms are normalized by R to ensure scale invariance across environments. The frontier score combines three components in Eq. 1. The parameter α weights the distance-from-robot term, encouraging progress toward frontiers farther from the current pose, while β weights the distance-from-visited-regions term, prioritizing exploration of unexplored space. The parameter γ penalizes frontiers that are close to obstacles, promoting safe navigation.

To assess the sensitivity of the system to exploration-driven behavior, we maintain a constant weight for robot-centric objectives (α) while modulating the exploration parameter β. The obstacle penalty $\gamma \in (0, 1]$ acts as a regularizer: small values weaken collision avoidance, while excessively large values may suppress valid frontiers in narrow passages. In practice, γ primarily affects feasibility rather than frontier ordering and is therefore treated as a secondary parameter.

B. Head Motion and Coverage Analysis

During each search attempt, the robot performs a full head rotation covering an azimuthal angle of ϕ at head tilt θ. We consider two tilt angles here and take into consideration the tilt with maximum sensing range for coverage area computation. The effective coverage area for a single search attempt is given by:

$$A_{\text{trial}} = A_\theta = \frac{\phi}{2\pi} \cdot \pi R_\theta^2 = \frac{\phi}{2} \cdot R_\theta^2, \tag{2}$$

where R_θ is the maximum sensing range possible at the tilt angle θ. Because multiple search attempts may observe overlapping regions, an overlap factor $\lambda \in [0, 1]$ is introduced. The total effective coverage area after N attempts is:

$$A_{\text{total}}(N) = N \cdot A_\theta \cdot (1 - \lambda). \tag{3}$$

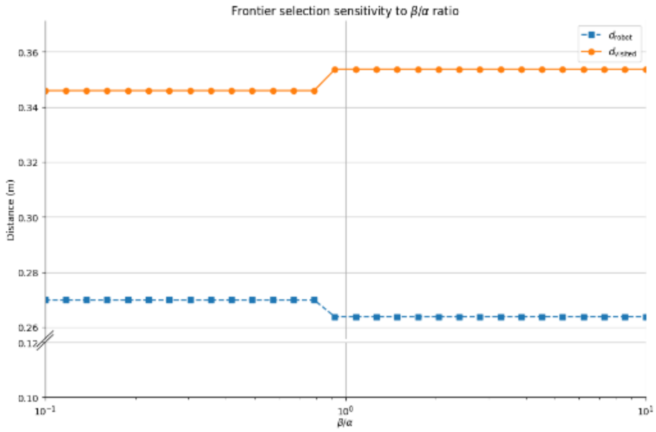

Fig. 2. Exploration behavior as a function of β/α (log-scale x-axis) with $\alpha = 1$. Distances to the robot (d_{robot}) and to visited regions (d_{visited}) normalized by the sensing range R_{30}. The metrics exhibit a transition to a stable regime for $\beta \gtrsim 1$, where frontier selection becomes largely insensitive to further increases in β.

This formulation provides a realistic estimate of the area explored by the robot during repeated search attempts at distinct and multiple frontier poses. [1]

V. Experimental Results

All experiments are conducted on a Dell Precision 7780 workstation with 64 GB RAM, running Ubuntu 22.04 LTS, using CPU-based execution. The robot platform has 8 GB RAM. We use YOLOE [25] for object detection and construct an online 3D point-cloud map using realsense active depth camera (D435if) to support segmentation analysis. Robot navigation toward the user is planned using RRT-Connect [26]. Face detection is performed with MediaPipe [27], and our framework is built on top of the Stretch AI repository [28].

We consider two head-tilt configurations: a floor-scanning view with tilt angle $\theta = 60^\circ$ and a frontal view with $\theta = 30^\circ$. Under these configurations, the robot achieves a maximum sensing range of $R_{30} = 2.08$ m for the 30° head tilt and $R_{60} = 1.7$ m for the 60° head tilt. The weighting parameter $\lambda = 0.3$ and the minimum distance threshold is set to $d_t = 1.45\,\text{m}$, equal to $(1 - \lambda)R_{30}$, to ensure sufficient separation between previously visited poses and newly selected frontier poses. The perception and navigation framework runs on a local workstation and communicates with the Stretch robot via a ROS2 bridge server.

A. Effect of β/α on Frontier Selection

Fig. 2 illustrates how frontier selection evolves as β increases from 0.1 to 10 (with $\alpha = 1$). For small β, the robot exhibits locally greedy behavior by preferring nearby frontiers. As β increases, the selection shifts toward frontiers further from previously visited areas, reaching a stable performance plateau for $\beta \gtrsim 1$.

[1] Source code repository: https://github.com/aakupadhyay/Human-Recognition-Driven-Interaction

Based on these results, we select $\alpha = 1.0$ and $\beta = 1.5$ to ensure thorough environmental coverage while avoiding myopic frontier selection. This balance is critical for fall detection, where a person might be obscured from a single viewpoint. By preventing the robot from staying too local, this configuration reduces blind spots and increases the probability of detecting a fallen individual from an optimal vantage point.

B. Area Coverage and Search Duration

We observe that the robot cannot perform more than one search attempt when operating on battery power. Consequently, we conduct all experiments with the robot connected to a charging cable. We therefore consider a tethered robot scenario, in which the robot navigates its environment while connected to a 3.96 m cable. We evaluate the tethered robot under two settings with 5 and 10 search attempts, measuring the total area covered and total search time. In each attempt, the robot scans the environment at head-tilt angles of 30° and 60°. Hence, the reported search time includes both floor-level and frontal-view scanning. Due to the head-mounted camera being positioned 125 cm above the ground, detections between $(0 - 30)^\circ$ head-tilt are unreliable, as the robot assumes the person is floating and cannot localize them for navigation. For each tilt angle, the head performs a full azimuthal sweep of $\phi = 300^\circ$ by rotating horizontally in 30° increments over 10 steps.

TABLE II
TIME AND COVERAGE AREA

Metrics	Search attempts		
	1	5	10
Coverage Area (m^2)	8.7	43.5	87
Total Time (s)	397.96	2168.42	4983.38

Table II reports the total search time and area covered in an environment without a person present. The coverage area is computed using the 30° head-tilt configuration, as the coverage region at 60° is fully contained within that at 30°. The search terminates if no person is found after 5 or 10 attempts or if no feasible frontier remains.

C. Person Detection in a Cluttered Environment

We evaluate the person-detection performance under two viewing conditions: *floor view* and *frontal view*. For each condition, we conduct 20 trials, with each trial allowing up to 10 search attempts at distinct frontier positions. In all experiments, the robot starts from a random position in a cluttered office environment containing large medical robots and multiple desktop workstations. Detection performance is measured using *precision*, which evaluates the robot's ability to correctly identify a human and to distinguish between a person lying on the floor and a person sitting/standing in front of the robot. Precision is computed using True Positives (TP) and False Positives (FP) as:

$$\text{Precision} = \frac{TP}{TP + FP}. \tag{4}$$

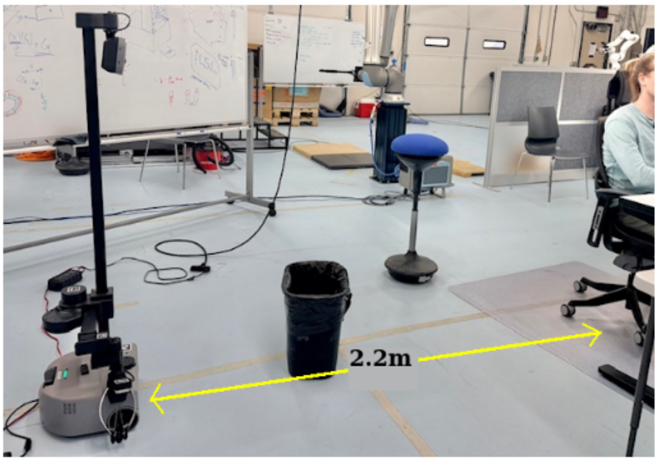

Fig. 3. Test Scenario: A user was seated 2.2 m away from the Stretch 3 robot. The robot took 245.37 s to detect the user, navigate to them, and initiate interaction in a cluttered space.

a) Floor view scan: Table III reports the detection precision and the average time required to identify a human. The robot successfully detected a person lying on the floor in all 20 trials, achieving a 100% success rate with an average detection time of 370 s. In all trials, the distance between the robot and the user ranged from 1.76 to 3.52 m.

b) Frontal view scan: Figure 3 illustrates a test scenario using the Stretch SE3 robot, in which the user is initially outside the robot's maximum sensing range of 2.08 m. The robot is initialized at different locations, detecting the human at varying search attempts depending on its random starting position and the selected frontier poses. For the frontal view, the distance between the robot and the user ranged from 2.2 to 7.23 m. Across all 20 trials, the robot successfully detected the human in 18 trials with an average total detection time of 1121 s.

Although the robot correctly distinguished between a person lying on the floor and a seated or standing person, it occasionally misclassified manipulator robotic arms as humans, resulting in false positives (Table III). This issue can be mitigated by adjusting the instance confidence threshold to improve detection performance.

VI. DISCUSSION AND FUTURE WORK

The experimental results demonstrate a high success rate in identifying atypical human poses; however, several limitations must be addressed to transition this system from a controlled laboratory setting to a functional domestic environment.

A. Computational and Software Constraints

The current software computes offline trajectories, providing stability and predictability under static conditions but failing to account for dynamic obstacles such as moving individuals. Future iterations will integrate real-time dynamic obstacle avoidance and adaptive trajectory planning [29], [30] to enhance safety in shared human–robot environments.

TABLE III
HUMAN DETECTION SUCCESS RATE

Head rotation	Total trials	TP	FP	Precision	Avg. Time (s)
Floor view scan	20	20	0	100%	370.8±18
Frontal view scan	20	18	2	90%	1121.1±28.9

To transition the system from a research prototype to a viable real-world utility for emergency monitoring, future work will prioritize reduction in search and processing times, which currently average 1121.1 s for frontal scans. This latency is primarily induced by a sequential, CPU-bound pipeline in which point cloud processing, filtering, voxel updates, and map projection are executed on the host processor, limiting real-time quick responsiveness. To meet the stringent requirements of critical fall detection and emergency response, we will systematically migrate the mapping and inference pipeline to the NVIDIA GeForce RTX 5000 GPU equipped in the local workstation. This includes the development of CUDA-accelerated kernels for parallel spatial processing, along with TensorRT-based optimization for inference acceleration, thereby eliminating major host-side computational bottlenecks and reducing memory transfer overhead.

B. Hardware and Physical Limitations

Several hardware characteristics of the Hello Robot Stretch 3 introduce operational constraints. The base-mounted LiDAR, positioned at approximately 15 cm above the ground, may fail to detect low-lying obstacles, potentially leading to unintended contact during navigation. Additionally, the telescopic arm lacks dedicated collision sensing, increasing the risk of interaction with surrounding structures such as table stands.

Moreover, the robot's limited battery capacity necessitated a tethered power configuration during experimentation, restricting its operational range and limiting true autonomy. While this setup enabled controlled evaluation, it does not reflect realistic deployment conditions. In practical scenarios, this limitation could be addressed through autonomous docking stations and improved power management strategies to enable long-term, untethered operation.

VII. CONCLUSION

We evaluated the Hello Robot Stretch 3 for autonomous search, detection, and navigation toward a human to initiate interaction. Experimental results show that the robot reliably detects a person on the floor or within its frontal field of view and navigates to them from randomly initialized starting positions, achieving success rates of 90–100%. These results validate the effectiveness of the proposed perception–navigation framework and demonstrate the potential of commercial mobile manipulators for human-centered assistive tasks. Future work will focus on improving real-time performance through a GPU-parallelized implementation and evaluating the system in a high-fidelity mock home environment with realistic domestic constraints. We will also extend the framework with dynamic obstacle avoidance and enhanced sensor fusion for improved robustness and safety.

ACKNOWLEDGMENTS

This research was supported by Award 60NANB24D221 from the U.S. Department of Commerce's National Institute of Standards and Technology (NIST).

REFERENCES

[1] W.-Y. G. Louie, D. McColl, and G. Nejat, "Acceptance and attitudes toward a human-like socially assistive robot by older adults," *Assistive technology*, vol. 26, no. 3, pp. 140–150, 2014.

[2] L. M. Beuscher, J. Fan, N. Sarkar, M. S. Dietrich, P. A. Newhouse, K. F. Miller, and L. C. Mion, "Socially assistive robots: measuring older adults' perceptions," *Journal of gerontological nursing*, vol. 43, no. 12, pp. 35–43, 2017.

[3] E. Ackerman, "A robot for humanity: How robots can empower people who need them the most," *IEEE Spectrum*, vol. 60, no. 11, pp. 22–29, 2023.

[4] S. Giovannini, F. Brau, V. Galluzzo, D. A. Santagada, C. Loreti, L. Biscotti, A. Laudisio, G. Zuccala, and R. Bernabei, "Falls among older adults: screening, identification, rehabilitation, and management," *Applied Sciences*, vol. 12, no. 15, p. 7934, 2022.

[5] D. Cook, "Sensors in support of aging-in-place: The good, the bad, and the opportunities," *National Academies of Sciences, Engineering, and Medicine; Division of Behavioral and Social Sciences and Education*, 2020.

[6] X. Wang, J. Ellul, and G. Azzopardi, "Elderly fall detection systems: A literature survey," *Frontiers in Robotics and AI*, vol. 7, p. 71, 2020.

[7] D. Fischinger, P. Einramhof, K. Papoutsakis, W. Wohlkinger, P. Mayer, P. Panek, S. Hofmann, T. Koertner, A. Weiss, A. Argyros, *et al.*, "Hobbit, a care robot supporting independent living at home: First prototype and lessons learned," *Robotics and autonomous systems*, vol. 75, pp. 60–78, 2016.

[8] H. Wang, G. Wang, H. Zhu, Z. Du, and W. Wang, "A semantic-aware measurement-driven autonomous target search framework for ugv in complex unknown environments," *Measurement Science and Technology*, vol. 36, no. 7, p. 076313, 2025.

[9] I. Lluvia, E. Lazkano, and A. Ansuategi, "Active mapping and robot exploration: A survey," *Sensors*, vol. 21, no. 7, p. 2445, 2021.

[10] C. C. Kemp, A. Edsinger, H. M. Clever, and B. Matulevich, "The design of stretch: A compact, lightweight mobile manipulator for indoor human environments," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 3150–3157.

[11] M. F. Ali, D. C. Nchekwube, O. Genova, A. Freddi, and A. Monteriù, "An assistive robot in an indoor scenario: the stretch hello robot as environment organizer," in *2023 IEEE International Conference on Metrology for eXtended Reality, Artificial Intelligence and Neural Engineering (MetroXRAINE)*. IEEE, 2023, pp. 676–681.

[12] R. Sharma, A. Marshall, D. Montgomery, and R. Gamble, "Towards adaptive assistance: A preliminary architecture for dynamic user profiling in social robots," 2025.

[13] L. G. Troconis, O. Genova, A. Freddi, and A. Monteriù, "Human-like robotic handwriting with inverse kinematic using a mobile telescopic arm," in *2024 20th IEEE/ASME International Conference on Mechatronic and Embedded Systems and Applications (MESA)*. IEEE, 2024, pp. 1–6.

[14] Y.-L. Tsai, R. J. Mattas, S. A. Olatunji, and W. A. Rogers, "Assistive robots for home healthcare: Exploring proxemic factors for older adults," in *Proceedings of the International Symposium on Human Factors and Ergonomics in Health Care*, vol. 14. SAGE Publications Sage CA: Los Angeles, CA, 2025, pp. 21–25.

[15] V. Ranganeni, V. Dhat, N. Ponto, and M. Cakmak, "Accessteleopkit: A toolkit for creating accessible web-based interfaces for tele-operating an assistive robot," in *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*, 2024, pp. 1–12.

[16] L. Zheng, C. Zhu, J. Zhang, H. Zhao, H. Huang, M. Niessner, and K. Xu, "Active scene understanding via online semantic reconstruction," in *Computer Graphics Forum*, vol. 38, no. 7. Wiley Online Library, 2019, pp. 103–114.

[17] X. Xu, X. Zhang, Y. Liu, H. Wang, X. Zhang, and Y. Zhuang, "Semantic-driven informed planning and 3d reconstruction for the quadrotor unmanned aerial vehicle," *IEEE Transactions on Vehicular Technology*, vol. 74, no. 3, pp. 3843–3853, 2024.

[18] L. Lu, Y. Zhang, P. Zhou, J. Qi, Y. Pan, C. Fu, and J. Pan, "Semantics-aware receding horizon planner for object-centric active mapping," *IEEE Robotics and Automation Letters*, vol. 9, no. 4, pp. 3838–3845, 2024.

[19] T. Dang, C. Papachristos, and K. Alexis, "Autonomous exploration and simultaneous object search using aerial robots," in *2018 IEEE Aerospace Conference*. IEEE, 2018, pp. 1–7.

[20] P. Vallabh and R. Malekian, "Fall detection monitoring systems: a comprehensive review," *Journal of Ambient Intelligence and Humanized Computing*, vol. 9, no. 6, pp. 1809–1833, 2018.

[21] A. Ramachandran and A. Karuppiah, "A survey on recent advances in wearable fall detection systems," *BioMed research international*, vol. 2020, no. 1, p. 2167160, 2020.

[22] Y. S. Delahoz and M. A. Labrador, "Survey on fall detection and fall prevention using wearable and external sensors," *Sensors*, vol. 14, no. 10, pp. 19 806–19 842, 2014.

[23] P. Bet, P. C. Castro, and M. A. Ponti, "Fall detection and fall risk assessment in older person using wearable sensors: A systematic review," *International journal of medical informatics*, vol. 130, p. 103946, 2019.

[24] P. Asgharian, A. M. Panchea, and F. Ferland, "A review on the use of mobile service robots in elderly care," *Robotics*, vol. 11, no. 6, p. 127, 2022.

[25] A. Wang, L. Liu, H. Chen, Z. Lin, J. Han, and G. Ding, "Yoloe: Real-time seeing anything," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 24 591–24 602.

[26] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.

[27] M. Nazarkevych, V. Lutsyshyn, H. Nazarkevych, L. Parkhuts, and M. Kostiak, "Methods of face recognition in video sequences and performance studies," in *CEUR Workshop Proceedings*, vol. 3421, 2023, pp. 246–253.

[28] hello-robot, "Stretch AI: Intelligent behaviors for the stretch 3 mobile manipulator," https://github.com/hello-robot/stretch_ai, 2025, accessed: 2025-12-30.

[29] A. Nasrinahar and J. H. Chuah, "Intelligent motion planning of a mobile robot with dynamic obstacle avoidance," *Journal on Vehicle Routing Algorithms*, vol. 1, no. 2, pp. 89–104, 2018.

[30] V. D. Cong *et al.*, "Path following and avoiding obstacle for mobile robot under dynamic environments using reinforcement learning," *Journal of Robotics and Control (JRC)*, vol. 4, no. 2, pp. 157–164, 2023.